Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Stress-Level Classification for Students with Recommendation Outputs

Chandrasekar Adhithya Harsha Pasumarthi
Panther Creek High School, Frisco, TX, USA
Corresponding Author: Chandrasekar Adhithya Harsha Pasumarthi, **E-mail**: ca.pasumarthi@gmail.com

| ABSTRACT

Student mental health has deteriorated measurably and continuously over the past decade. A 2025 survey by SRM University AP, published in the Asian Journal of Psychiatry, found that nearly 70 percent of 1,628 college students across eight major Indian cities reported moderate to high anxiety levels, while over 60 percent exhibited clinically significant depression symptoms. These numbers are not outliers. Cross-national data shows Japan at 80 percent stress prevalence, the United States at 67 percent, and Pakistan at 71 percent. What unites these figures is not geography but the structure of modern academic life: excessive workload, unrelenting performance pressure, financial strain, social comparison amplified by digital platforms, and institutional support systems that were simply not built for this scale of need. The clinical response to this crisis has been slow and resource-constrained. Counseling services are stretched thin. Most students who need help do not seek it, largely because of stigma, lack of awareness, or long wait times. Only 7 percent of college students in the United States access professional mental health support in any given year (EssayPro, 2025). This gap between need and access defines the problem this article addresses. This article presents a machine learning-based classification framework that predicts student stress levels across three tiers: low, medium, and high. The model draws on academic, lifestyle, physiological, and psychometric data collected through structured surveys. It evaluates five classification algorithms, including Logistic Regression, Random Forest, XGBoost, and ensemble methods, with accuracies ranging from 86.2 to 97.92 percent in peer-reviewed studies. The framework also includes a recommendation engine that generates targeted, evidence-based intervention guidance for each stress tier. The goal is to give educational institutions a practical, scalable tool for identifying students at risk before their stress reaches crisis level.

| KEYWORDS

Stress-Level Classification; Students; Recommendation Outputs; academic life

| ARTICLE INFORMATION

ACCEPTED: 01 April 2026 **PUBLISHED:** 15 April 2026 **DOI:** 10.32996/jcsts.2026.8.6.1

1. Introduction

Stress, in controlled doses, is not harmful. Physiologically, it is a survival mechanism. It sharpens attention, mobilizes energy, and prepares the body to respond to challenge. But the stress most university students experience today bears little resemblance to that adaptive function. It is not episodic. It does not resolve after an exam or a deadline. For a growing proportion of students, stress is the baseline state of being a student, and it is producing measurable damage to mental health, academic performance, and long-term wellbeing.

The statistics published in recent years tell a consistent story. In the United States, 37.1 percent of college students have been diagnosed with stress-related conditions, 31.1 percent with anxiety disorders, and 20.5 percent with depression (EssayPro, 2025). In India, a study spanning eight Tier-1 cities found that more than two-thirds of students reported emotional distress severe enough to affect their daily functioning (India Today, 2025). In Japan, 80 percent of students report experiencing stress on a daily basis. These are not marginal populations experiencing isolated difficulties. This is a majority phenomenon.

What makes this particularly difficult to address at institutional scale is the mismatch between need and available infrastructure. University counseling centers in most countries operate with counselor-to-student ratios far above recommended levels. The American Psychological Association recommends no more than one counselor per 1,500 students; actual ratios at many universities are double or triple that figure. In India, the situation is worse. Dr. Neetu Tiwari of NIIMS Medical College has argued that no counselor should manage more than 500 students, yet the current reality in many Indian institutions falls far short of that standard (India Today, 2025). The result is a system where students in genuine distress face weeks-long wait times for appointments, where preventive mental health work is simply not possible, and where the first time a student receives professional attention is often after a crisis has already occurred.

This is the context in which machine learning has begun to receive serious attention as a mental health support tool. The argument is practical, not utopian. ML models do not replace counselors. They do not diagnose clinical conditions. What they can do, when trained on the right data, is identify patterns of risk across large student populations quickly and cheaply. They can flag individuals who are likely to be experiencing high stress before that stress becomes a presenting crisis. They can help institutions allocate their limited counseling resources toward students who most need them. And through recommendation systems built on top of classification outputs, they can deliver personalized guidance to the majority of students whose stress levels do not warrant immediate clinical intervention but who would benefit from structured support.

This article builds on three recently published empirical studies to examine how that process works in practice. It covers the data these systems require, the algorithms that perform best on student stress datasets, the features that carry the most predictive weight, and the kinds of recommendations that are both evidence-based and operationally realistic for educational institutions to deliver.

2. Understanding Student Stress: Definitions, Dimensions, and Causes

2.1 Defining Stress in an Academic Context

Stress is a physiological and psychological response to perceived demands that exceed available coping resources. Hans Selye, who first systematized the biological study of stress in the 1930s, distinguished between eustress (productive, motivating stress) and distress (harmful, depleting stress). In academic settings, the distinction matters because not all pressure is harmful. Deadline pressure, competitive grading, and high expectations can motivate effort and improve outcomes. The problem arises when demands consistently and substantially exceed a student's capacity to cope, producing chronic activation of the stress response without adequate recovery.

Chronic stress has well-documented neurobiological consequences. Prolonged elevated cortisol levels impair hippocampal function, weakening memory consolidation and reducing the capacity for new learning. The prefrontal cortex, which governs executive function, attention regulation, and decision-making, becomes less effective under sustained stress. For students, these are not abstract neurological events. They translate directly into difficulty concentrating during lectures, reduced retention of course material, impaired performance during examinations, and a diminishing sense of self-efficacy that compounds the original stressor.

In the taxonomy used by Yeler and Sürücü (2025), student stress sources fall into three broad categories. Psychological sources include test anxiety, low self-esteem, depressive episodes, and a chronic fear of failure that many high-achieving students carry quietly throughout their academic careers. Physiological sources include sleep deprivation, headaches, gastrointestinal complaints, and elevated blood pressure, all of which are both symptoms of stress and independent sources of further stress. Environmental sources include academic workload, financial constraints, the quality of living conditions, family expectations, social dynamics, and institutional culture.

2.2 What the PSS-10 Measures and Why It Matters

The Perceived Stress Scale, specifically the 10-item version developed by Cohen, Kamarck, and Mermelstein (1983), is the primary psychometric instrument used in the ML studies reviewed in this article. The PSS-10 measures the degree to which individuals perceive their lives as unpredictable, uncontrollable, and overwhelming over the preceding month. It consists of ten questions rated on a five-point scale from never (0) to very often (4), with four positively framed items reverse-scored.

Composite scores are interpreted as follows: 0 to 13 indicates low perceived stress, 14 to 26 indicates moderate perceived stress, and 27 to 40 indicates high perceived stress. These thresholds have been validated across multiple demographic groups and in cross-cultural studies. The PSS-10 is particularly well-suited for student populations because it captures perceived stress rather than objective stressor frequency, making it sensitive to the subjective and contextual nature of academic pressure. A student managing five courses and a part-time job may score lower than a student carrying three courses if the former has stronger coping resources, social support, and self-efficacy.

In the Bangladesh study by Biswas, Mim, and Ferdous (2026), PSS-10 responses were collected from undergraduate students at public, private, and national universities. The Total PSS-10 score emerged as the single strongest predictor of stress classification in the full-feature Random Forest model, with a feature importance score of approximately 0.60, substantially higher than any other individual feature. This is consistent with the theoretical literature: self-perceived stress, precisely because it integrates objective stressors with subjective appraisal and coping capacity, provides the most complete and clinically meaningful signal.

2.3 The Social and Structural Dimensions of Student Stress

It would be a mistake to treat student stress purely as an individual psychological phenomenon. The SRM University study (India Today, 2025) gathered data from students in eight Indian cities and found patterns that are structural rather than personal. Delhi students showed the highest rates of anxiety and depression, which the researchers attributed to a combination of high urban living costs, intense academic competition at elite institutions, and the social isolation experienced by students who migrate from smaller cities to metropolitan campuses. These students face not just academic pressure but the challenge of constructing an adult identity in an unfamiliar environment, often without the family and community networks that previously buffered their stress.

The study also found consistent gender differences. Female students reported higher levels of emotional distress and lower overall wellbeing than male students across all eight cities. This finding aligns with international literature. Gao, Ping, and Liu (2020), in a longitudinal study of Chinese college students published in the Journal of Affective Disorders, found that female students reported higher depression, anxiety, and stress scores than male students at every measurement point over a two-year period. The mechanisms are complex and include differences in socialization around emotional expression, differential exposure to social comparison pressures via digital platforms, and structural inequities in how academic and professional expectations are applied to women.

Socioeconomic status adds another layer. The Essex University study (Biswas et al., 2026) specifically examined whether socioeconomic factors could predict student stress. Family income was the strongest socioeconomic predictor, with an importance score of 0.70 in the SES-only model. Students from lower-income families face a different stress profile from those in higher-income brackets: financial worry about tuition and living costs, reduced access to academic resources, greater likelihood of working part-time alongside studies, and in some cases, family pressure to support other household members financially. When these socioeconomic variables were the only features available to the model, accuracy dropped to 46 percent, confirming that socioeconomic status contributes meaningfully to stress but does not determine it. Stress is produced by the interaction of financial, academic, social, and psychological factors, and effective classification requires data from all of these domains.

3. The Scale of the Student Mental Health Crisis

3.1 Global Prevalence

The data presented in Figure 1 below comes from a 2025 cross-national survey compiled by EssayPro. The figures show that student stress is elevated in every country studied, but the severity varies considerably.

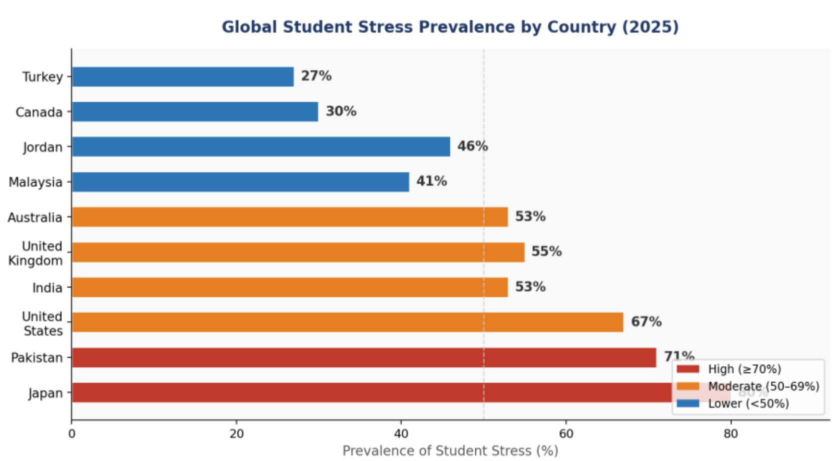

Figure 1: Global Student Stress Prevalence by Country (EssayPro, 2025)

Japan, at 80 percent, leads partly because of structural features of the Japanese education system: high-stakes university entrance examinations, intense social pressure around academic performance, and cultural norms that discourage public expressions of distress or help-seeking. Pakistan, at 71 percent, reflects a different set of pressures, primarily socioeconomic: students navigating financial hardship, limited institutional resources, and in many cases, the weight of being the first in a family to pursue higher education. The United States, at 67 percent, sits in a context of rising tuition costs, competitive labor markets, and documented declines in campus mental health infrastructure relative to student need.

Canada's comparatively lower figure of 30 percent and Turkey's 27 percent may partly reflect methodological differences in how stress was measured across studies, and partly reflect genuine differences in academic culture, social support systems, and the role of family in buffering student stress. The cross-national comparison is useful not because it produces a definitive ranking but because it confirms that student mental health strain is not a product of any single culture or economic system. It is a near-universal feature of how formal education is organized in the early twenty-first century.

3.2 The Indian Picture in Detail

The SRM University AP study, published in the Asian Journal of Psychiatry and reported by India Today in October 2025, provides the most detailed recent picture of student mental health in India. The sample of 1,628 students aged 18 to 29, drawn from eight Tier-1 cities with a gender distribution of 52.9 percent female and 47.1 percent male, produced findings that Dr. Neetu Tiwari described as potentially indicating the early signs of a systemic mental health crisis on Indian campuses. Figure 2 presents the core findings.

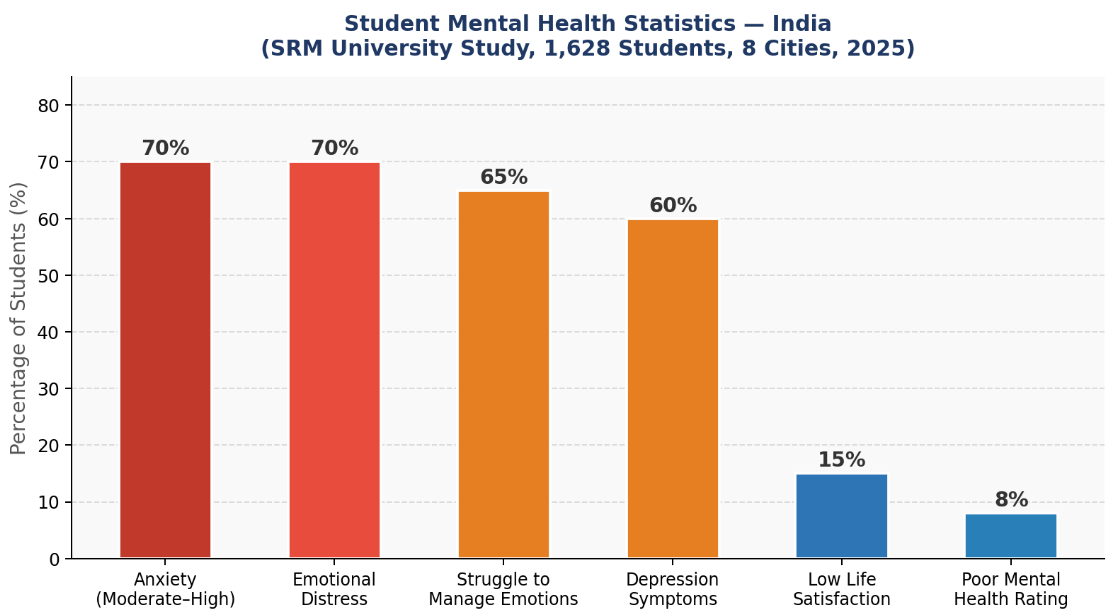

Figure 2: Student Mental Health Statistics, India (SRM University Study, India Today, 2025)

The most striking figure is that more than 70 percent of respondents reported emotional distress that was disrupting their daily functioning. This is not a measure of occasional bad days. It refers to sustained emotional strain that affects the capacity to study, maintain relationships, and manage basic responsibilities. Sixty-five percent reported specific difficulty regulating their emotions or behavior. Taken together, these figures suggest that the majority of students in India's largest cities are currently operating under a level of psychological load that, in a clinical assessment, would warrant active support.

The disparities within the data are as important as the aggregate numbers. Central University students reported higher rates of depression than students at private institutions. Students from government arts and science colleges recorded the lowest life satisfaction scores. Dr. Sabine Kapasi, CEO of Enira Consulting and a UN advisor, noted that the Indian educational system has structured success almost entirely around grades and social affirmation, creating an environment where emotional development is treated as secondary at best and irrelevant at worst. The consequences of that structural choice are visible in these numbers.

3.3 The Help-Seeking Gap

The most practically significant gap in the student mental health landscape is not between students who are stressed and students who are not. It is between students who are stressed and students who seek help for it. Figure 3 presents this disparity.

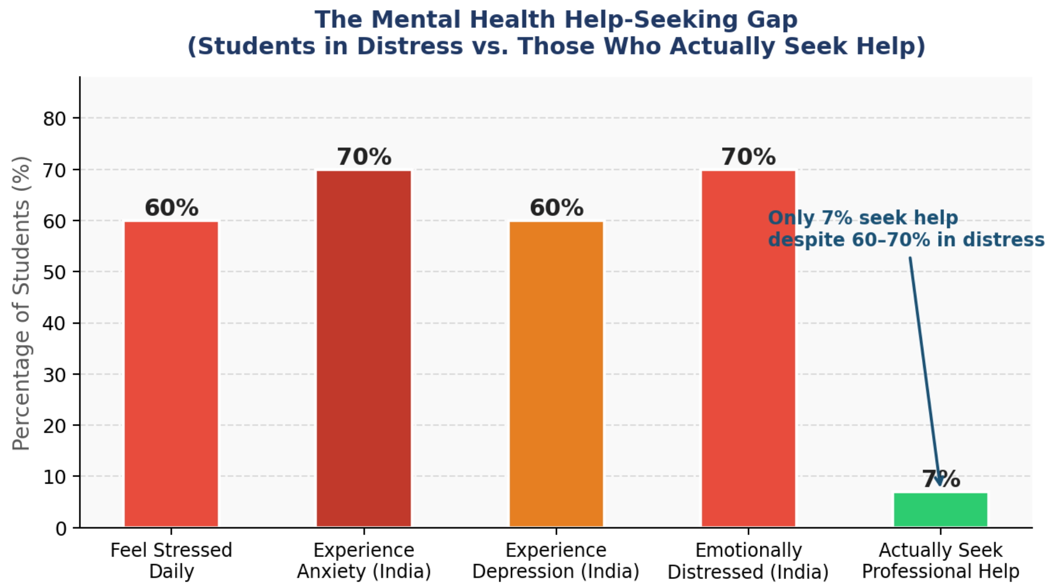

Figure 3: The Mental Health Help-Seeking Gap (EssayPro, 2025; India Today, 2025)

Only 7 percent of college students in the United States access mental health professionals for stress or depression in any given year (EssayPro, 2025). In India, the figure is likely lower. The barriers are well-documented. Stigma is the most frequently cited: in many communities, acknowledging mental health difficulties is equated with personal weakness, and students fear academic or social consequences from being seen as struggling. Lack of awareness is a secondary factor: many students experiencing anxiety or depression for the first time do not recognize what they are experiencing, attributing concentration difficulties, fatigue, and emotional numbness to personal inadequacy rather than to a recognizable and treatable condition.

Institutional barriers compound the problem further. Kala Soni, a school counselor at Witty International School in Mumbai, noted that students frequently hesitate to use available counseling services because they fear being judged, or because accessing those services is itself stigmatizing in a campus culture that treats mental health as a private matter. Even at institutions that have invested in expanded counseling provision, such as the IIT campuses that now operate 24-hour helplines and mandatory first-year counseling sessions, demand consistently exceeds capacity. The result is a system in which the majority of students who need structured support either receive none or receive it only after a crisis.

This is the practical argument for automated early detection. The goal is not to replace counselors but to identify, before crisis, the students who most need their attention.

4. Why Traditional Assessment Methods Are Not Enough

The standard approach to assessing student mental health at an institutional level relies on some combination of voluntary self-referral to counseling services, periodic large-scale surveys such as annual wellbeing checks, faculty reporting of students who appear to be struggling, and reactive intervention after a crisis event. Each of these mechanisms has well-known weaknesses.

Self-referral systematically underidentifies students in distress, for the stigma and awareness reasons discussed above. Large-scale surveys provide useful population-level data but are typically administered once per year, produce results months after data collection, and cannot identify at-risk individuals in real time. Faculty reporting is unreliable because instructors are not trained clinical assessors, because many students in distress maintain adequate academic performance for considerable periods before their difficulties become visible, and because the faculty-student relationship varies enormously across institutions and disciplines. Reactive post-crisis intervention addresses harm after it has occurred rather than preventing it.

The limitations of subjective assessment methods are also discussed by Yeler and Sürücü (2025), who note that self-report tools like campus surveys can be affected by individual differences in self-awareness, cultural norms around disclosure, and deliberate manipulation of responses by students who want to be seen as managing well. The authors argue that objective, data-driven methods reduce these biases by deriving stress classifications from patterns across multiple variables rather than from a single self-reported assessment.

Machine learning does not eliminate these problems entirely. The models reviewed in this article still depend substantially on self-reported data, particularly the PSS-10 scale. But they process that data differently from a single-item assessment: they identify how PSS-10 scores interact with sleep patterns, academic performance, social engagement, and financial circumstances to produce a classification that is more robust than any of those variables in isolation. And they can do this at scale, across thousands of students, in time frames that allow for timely intervention rather than retrospective analysis.

5. Machine Learning Methodology

5.1 Dataset Characteristics

The primary dataset used in the studies reviewed here is the Student Stress Monitoring Dataset, shared on the Kaggle platform and described in detail by Ovi et al. (2025). It contains 1,100 student observations across 21 independent features and a three-class target variable (stress_level: 0 = Low, 1 = Medium, 2 = High). The dataset is notable for its balanced class distribution, with approximately equal numbers of students in each stress tier, which reduces the class imbalance problem that commonly degrades ML performance in health datasets. Figure 4 shows the class distribution.

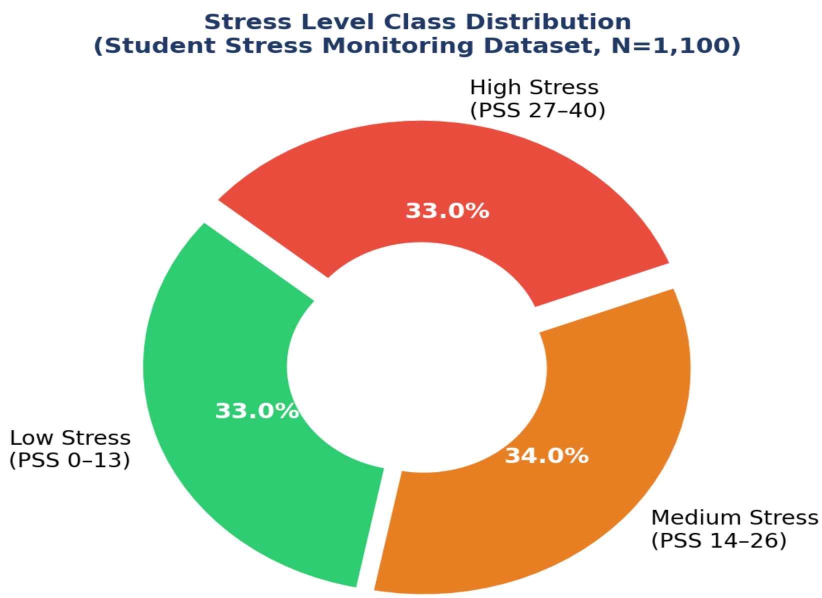

Figure 4: Stress Level Class Distribution, Student Stress Monitoring Dataset (N=1,100)

The 21 features span the three stress domains identified in the theoretical literature. Psychological features include anxiety score, mental health history, self-esteem rating, and depression indicators. Physiological features include sleep quality, sleep duration, headache frequency, and blood pressure rating. Environmental features include academic performance measures, study load, extracurricular involvement, social support scores, peer pressure ratings, noise level in living conditions, and future career concerns. This multi-domain feature structure is what gives the models their predictive power: stress is a systemic phenomenon, and its classification requires data from across the system.

The Bangladesh study by Biswas et al. (2026) used a different, smaller dataset of 150 students from three types of Bangladeshi universities (public, private, and national), collected via Google Forms between January and February 2025. This dataset included 14 demographic and socioeconomic questions alongside the ten PSS-10 questions. Despite its smaller size, this dataset produced strong model performance when the full feature set was used, and it provides the most detailed available evidence about how socioeconomic factors interact with other stress predictors.

5.2 Data Preprocessing

Before training, both datasets underwent several preprocessing steps that are standard in ML pipelines for health and behavioral data.

Missing value checking was performed using Python's isnull().sum() method. Both datasets contained no missing values, eliminating the need for imputation. This is worth noting because missing values are common in survey data collected from student populations, where respondents may skip questions they find sensitive or irrelevant, and their absence here may reflect careful survey design.

Feature standardization was applied using StandardScaler, transforming each numerical feature to a mean of zero and a standard deviation of one. This step is particularly important for algorithms that are sensitive to feature scale, including Logistic Regression and Support Vector Machine. Without standardization, a feature measured on a large numerical scale (such as monthly family income) would dominate a feature measured on a small scale (such as a PSS-10 sub-score), introducing spurious weighting effects.

Categorical variables were encoded numerically using label encoding. Gender was coded as 0 for male and 1 for female. Relationship status was encoded as 0 for single and 1 for married. University type, socioeconomic class, and scholarship status were encoded similarly. This encoding allows categorical information to be incorporated into ML algorithms that require numerical input.

The train-test split followed a stratified 80/20 ratio in the Yeler and Sürücü (2025) study, meaning that 880 observations were used for training and 220 for testing, with the class distribution maintained proportionally across both sets. The Biswas et al. (2026) study used a 60/40 split, which is more conservative given the smaller dataset size. Stratified splitting prevents the test set from accidentally containing a disproportionate number of one stress class, which would produce misleadingly high or low accuracy metrics.

5.3 Classification Algorithms: How Each One Works

Logistic Regression

Logistic Regression is a linear classifier extended to multi-class problems through what is known as a multinomial or one-vs-rest formulation. For each observation, the model computes a weighted sum of the input features, passes it through a sigmoid function to produce a probability, and assigns the observation to the class with the highest probability. Despite being labeled a regression technique, it is fully a classification algorithm when applied to categorical outcomes.

The key advantage of Logistic Regression in this context is interpretability. Its coefficients can be examined directly to understand how each feature contributes to the classification, making it the most transparent of the algorithms tested. In the Yeler and Sürücü (2025) study, it achieved the highest accuracy of 88.1 percent, with a precision of 0.91 and an F1-score of 0.91 for the high-stress class. The authors attribute this performance to strong linear separability in the feature space once standardization was applied. When the relationships between features and stress class follow roughly linear patterns, Logistic Regression extracts the signal efficiently without the overhead of more complex models.

Random Forest

Random Forest is an ensemble method that constructs multiple decision trees on bootstrapped samples of the training data, with each tree considering a random subset of features at each split. The final prediction is determined by majority vote across all trees. This randomization process reduces the overfitting that plagues individual decision trees, which tend to memorize training data rather than learn generalizable patterns.

In the studies reviewed here, Random Forest achieved 86.2 percent accuracy on the Kaggle dataset (Yeler and Sürücü, 2025) and 93.75 percent on the Bangladesh dataset (Biswas et al., 2026). The difference in performance across datasets likely reflects the difference in feature composition: the Bangladesh dataset included richer socioeconomic and demographic information that contributed additional signal. Random Forest also provided the feature importance scores discussed in Section 6, making it valuable not only as a classifier but as a tool for understanding which variables carry the most predictive weight.

XGBoost

XGBoost, short for Extreme Gradient Boosting, is a tree-based ensemble algorithm that builds models sequentially. Each new tree is trained to correct the errors of the ensemble of trees built before it. The process is governed by the gradient of the loss function, which indicates in which direction the model should adjust its predictions. A learning rate parameter controls how aggressively each new tree adjusts the ensemble.

XGBoost is widely regarded as one of the most effective algorithms for structured tabular data, which is precisely the format of the datasets used here. It achieved 86.8 percent accuracy in the Yeler and Sürücü (2025) study, marginally ahead of Random Forest. In practice, XGBoost tends to require more careful hyperparameter tuning than Random Forest to reach its performance ceiling, but it is also more flexible in handling complex, non-linear feature interactions. For stress prediction tasks where the relationships between features are likely to be non-linear and interaction-dependent, XGBoost's sequential correction mechanism provides a useful advantage.

Decision Tree, CART, AdaBoost, and Bagging

The Biswas et al. (2026) study tested a wider set of algorithms than the Yeler and Sürücü (2025) study, including Decision Trees, CART (Classification and Regression Trees), AdaBoost, and Bagging. The first two are single-tree methods that split the feature space recursively to create homogeneous nodes, using Gini impurity as the splitting criterion. The latter two are ensemble methods that combine multiple weak classifiers to produce a stronger aggregate prediction.

Decision Tree, CART, AdaBoost, and Bagging all achieved 97.92 percent accuracy in the Biswas et al. (2026) study, which is the highest figure reported across any study in this review. This remarkable performance is likely explained by a combination of factors: the relatively small dataset size (150 observations) reduces the generalization challenge, the PSS-10 total score is a highly informative feature that gives tree-based models a reliable primary split, and the balanced class distribution prevents the class-imbalance performance degradation common in health datasets. The results are promising but should be interpreted cautiously because high accuracy on small datasets does not always generalize to larger, more heterogeneous populations.

Support Vector Machine

Support Vector Machine (SVM) finds the hyperplane in feature space that maximizes the margin between classes. When class boundaries are not linearly separable, SVM uses a kernel function to project the data into a higher-dimensional space where linear separation becomes possible. In the Biswas et al. (2026) study, SVM achieved only 56.25 percent accuracy, performing considerably worse than tree-based and ensemble methods. The authors attribute this to the highly non-linear relationships between socioeconomic and behavioral features and stress class outcomes. SVM is included here as a methodological baseline rather than a competitive candidate for deployment.

6. Model Performance and Feature Importance

6.1 Accuracy Across Algorithms and Studies

Figure 5 compares classification accuracy across the algorithms and datasets reviewed in this article. The spread from 56.25 to 97.92 percent reflects both genuine differences in algorithm capability and the influence of dataset characteristics, particularly size, feature richness, and class balance.

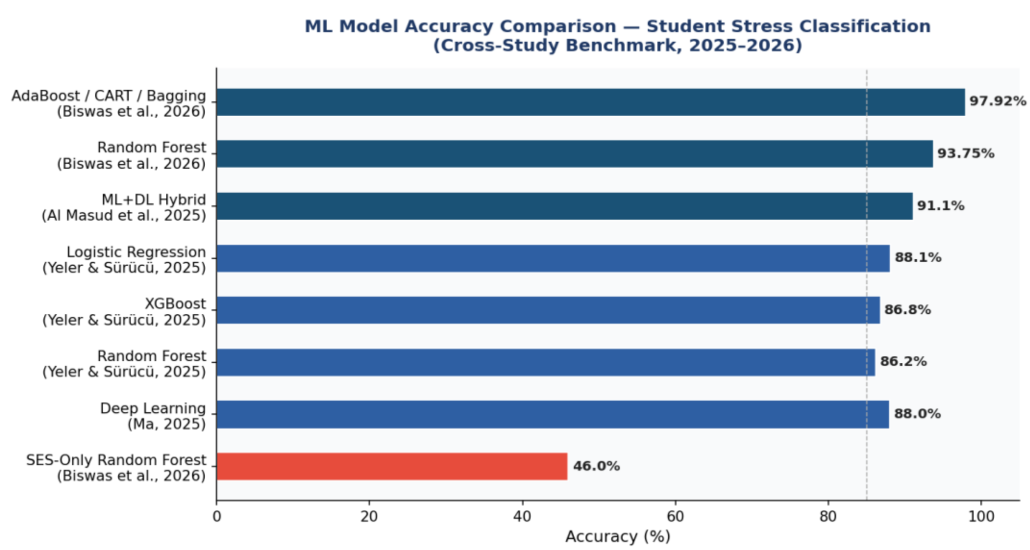

Figure 5: ML Model Accuracy Comparison (Yeler & Sürücü, IJANSER 2025; Biswas et al., Essex Student Journal 2026)

The most practically relevant figures for institutional deployment are probably those from the Yeler and Sürücü (2025) study, which used the larger and more methodologically standardized Kaggle dataset. Logistic Regression at 88.1 percent accuracy, with a 0.91 precision and F1-score for the high-stress class, represents a strong baseline for a deployable system. The near-equal performance of XGBoost (86.8 percent) and Random Forest (86.2 percent) suggests that any of these three algorithms could anchor a reliable production classification system.

The per-class performance data from the Yeler and Sürücü (2025) confusion matrix analysis shows that all three algorithms maintained consistent performance across all three stress classes, with no significant imbalance between them. This is an important property for a mental health classification system, where underidentifying high-stress students is a more serious error than incorrectly classifying low-stress students as medium-stress. A system that achieves high aggregate accuracy by performing well on the majority class but poorly on the critical minority class would not be acceptable for clinical support applications.

6.2 What the Features Tell Us

Figure 6 presents feature importance from two models in the Biswas et al. (2026) study: the full-feature Random Forest and the SES-only Random Forest. The comparison is informative beyond its classification accuracy implications.

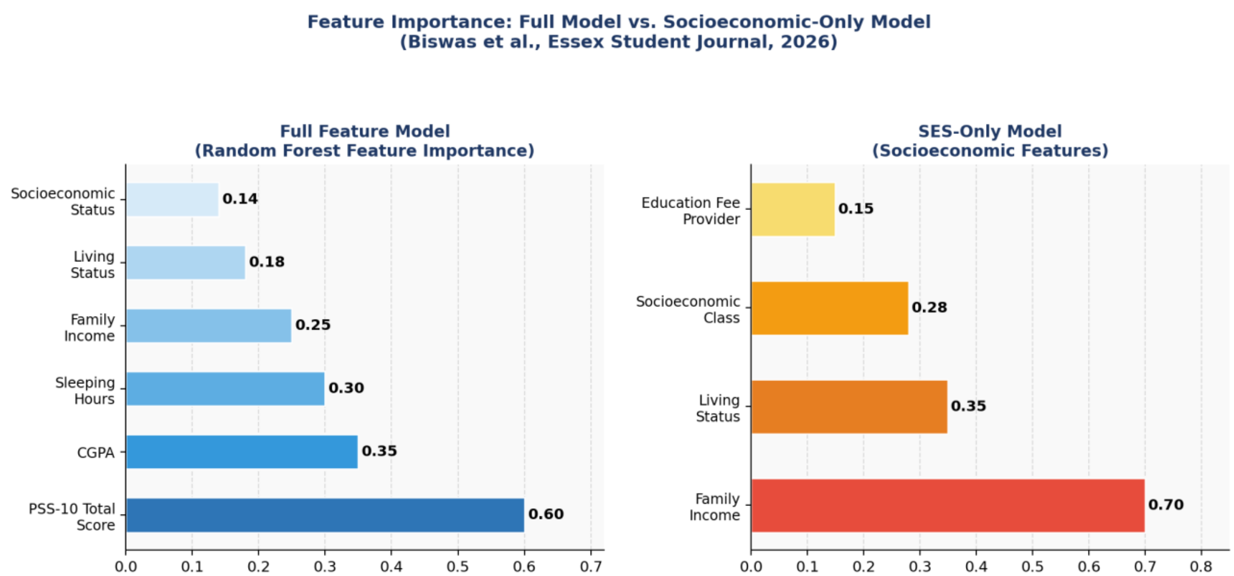

Figure 6: Feature Importance Comparison, Full Model vs. SES-Only Model (Biswas et al., Essex Student Journal, 2026)

In the full-feature model, the PSS-10 Total Score dominates with an importance of approximately 0.60. This is not surprising: the PSS-10 is a validated, multi-question composite of stress perception, and its aggregate score carries more information than any single demographic or behavioral variable. What is more instructive is the ranking of the remaining features. CGPA (academic performance) ranks second at approximately 0.35, followed by Sleeping Hours at 0.30 and Family Income at 0.25. Living Status and Socioeconomic Status contribute additional but smaller amounts.

The sleeping hours finding deserves attention. Among all behavioral variables, sleep quality and quantity are the most consistently identified predictors of both stress and mental health outcomes in the educational psychology literature. EssayPro (2025) reports that 42.6 percent of college students sleep fewer than seven hours on weeknights, and 77 percent of high school students are clinically sleep-deprived. The SRM University study (India Today, 2025) identified sleep as one of the top three stressors. The feature importance data here places it as the third strongest predictor of stress classification, reinforcing what clinicians and researchers have argued for years: sleep is not a lifestyle preference that students can sacrifice for academic productivity. It is a physiological necessity whose deficit is both a symptom and a cause of the stress the classification system is trying to detect.

The SES-only model tells a different but related story. When the model can only see socioeconomic variables, Family Income becomes overwhelmingly dominant at approximately 0.70. Living Status and Socioeconomic Class follow at 0.35 and 0.28 respectively. However, the 46 percent accuracy ceiling means that these variables alone cannot produce a reliable classifier. Stress does not reduce to poverty. A student from a high-income family can be severely stressed by perfectionism, social anxiety, and academic overload, while a student from a low-income family who has strong coping skills, family support, and a manageable course load may report low stress. The full feature set captures this complexity; the SES-only model cannot.

I. 7. The Personalized Recommendation Engine

A classification output that says a student is at high stress risk is useful only if it connects to action. Without a recommendation layer, the system produces a label and nothing else. The recommendation engine described here translates each stress tier into a structured set of evidence-based interventions, drawing on behavioral health literature, clinical guidelines, and the specific stressor profiles identified by the studies reviewed in this article.

The engine operates on a rule-based architecture at its current stage: each stress classification triggers a predetermined recommendation profile. Future iterations could incorporate large language model-driven personalization that adapts recommendations to individual student profiles, as discussed in Section 10. The current rule-based design has the advantage of transparency and clinical auditability. Every recommendation can be traced back to a specific evidence base, and the logic connecting classification to recommendation is visible to administrators and counselors who oversee the system.

7.1 Low Stress Tier: Building Resilience Before It Is Needed

Students classified as low stress are the majority of the population at any given time, and they are typically not the focus of mental health interventions. That is a missed opportunity. Preventive mental health work with students who are currently coping well is both cheaper and more effective than crisis intervention. Building psychological resilience, sleep hygiene, and adaptive coping skills before high-stress periods arrive is the most cost-effective use of institutional mental health resources.

Recommended interventions for this tier focus on maintenance and prevention rather than treatment.

- Continue current study and sleep schedules, and use PSS-10 self-assessment at the start of each semester to monitor stress trends over time.
- Maintain at least 150 minutes of moderate aerobic exercise per week. Research consistently links physical activity to reduced cortisol levels and improved mood regulation.
- Practice brief daily mindfulness exercises of 10 to 15 minutes. Mindfulness-based stress reduction programs have demonstrated measurable reductions in anxiety symptoms in university student populations.
- Engage actively with peer networks and campus communities. Social connection is one of the most robust buffers against stress escalation, and its protective effects are strongest when built during low-stress periods.
- Develop time management structures, particularly for periods of high workload such as mid-semester and examination periods, before those periods arrive rather than in response to them.

7.2 Medium Stress Tier: Targeted Early Intervention

Medium stress students are the most critical group for the purposes of this framework. They are experiencing genuine distress that is affecting their functioning, but they are not yet in crisis. The research literature on stress progression indicates that medium stress, when unaddressed, transitions to high stress substantially more frequently when the academic calendar reaches high-intensity periods such as examination weeks. Early intervention at this stage prevents a much more resource-intensive response later.

Recommendations for this tier are more specific and more behavioral than those for the low stress tier.

- Implement structured time-blocking using the Pomodoro Technique or similar methods: 25-minute focused study periods followed by 5-minute breaks, with longer breaks every four cycles. This approach reduces the cognitive fatigue that accumulates during prolonged, unstructured study sessions.
- Increase aerobic exercise to at least 20 to 30 minutes, three to five times per week. Studies in exercise physiology have documented cortisol reductions of 30 to 40 percent following regular aerobic activity in student populations. This is a pharmacological-scale effect available without prescription.
- Introduce cognitive behavioral journaling: daily written processing of academic stressors and the thought patterns associated with them. CBT-based journaling helps students identify cognitive distortions such as catastrophizing and all-or-nothing thinking, which amplify stress responses disproportionately.

- Prioritize sleep above most other competing demands. Seven to nine hours of sleep per night is not aspirational for this population. It is a clinical recommendation. Students sleeping fewer than six hours are impaired in cognitive processing to a degree comparable to mild intoxication.
- Schedule a proactive counseling check-in through the institution's student wellness platform. The emphasis here is on proactive: the student should not wait until they feel unable to cope. A single 30-minute session with a counselor or peer support specialist at this stage can interrupt stress progression that would otherwise continue.
- Where financial stress is a contributing factor, which feature importance data suggests it often is, facilitate access to scholarship information, emergency bursary funds, and financial counseling services. Financial anxiety has direct physiological stress effects, and addressing its source is more effective than treating its psychological symptoms alone.

7.3 High Stress Tier: Structured Crisis Response

Students classified as high stress require prompt, coordinated institutional response. The recommendation engine at this tier does not replace clinical judgment; it accelerates the referral process and provides a structured framework for the response.

- Immediate referral to licensed mental health professionals. The system should generate a referral notification to the student's institution counseling service, with a suggested priority flag if the student's PSS-10 score is at the upper range of the high-stress classification.
- Temporary academic accommodations coordinated with academic staff, including deadline extensions, leave of absence options, and reduced course load pathways. The research on academic accommodation and mental health outcomes consistently finds that reducing acute academic load during high-stress periods improves both mental health and long-term academic performance.
- Structured crisis intervention planning with scheduled follow-up contacts. A single counseling session is not sufficient for students in the high-stress tier. Regular check-in cadences over a four to six week period are needed to monitor whether stress is reducing, stable, or escalating.
- Introduction to evidence-based digital mental health tools as supplementary support between professional sessions. Woebot, a CBT-based conversational agent with clinical trial evidence supporting its effectiveness in reducing anxiety and depression symptoms, is available at no cost. Headspace for Students and Calm provide guided mindfulness resources. These tools do not replace professional support but can extend its effect between sessions.
- Physical health screening. High-stress students should be assessed for sleep disorders, nutritional deficiencies, and untreated physical health conditions, all of which contribute independently to psychological stress and are often overlooked in purely psychological assessments.
- Community and peer connection. The social isolation identified as a key stressor for Delhi students in the SRM University study (India Today, 2025) is not unique to India. Students in the high-stress tier frequently report reduced social engagement, which compounds their distress. Facilitating connection to peer support groups or campus community programs should be a component of the crisis pathway alongside professional clinical support.

8. System Architecture and Institutional Deployment

The academic value of a classification framework is limited unless it can be deployed in a form that institutions can actually use. This section describes the operational architecture of a production-ready version of the system and the integration requirements that determine whether deployment is feasible.

8.1 Core Components

The system has four core components. The data collection layer consists of the survey instrument, most likely a digital version of the PSS-10 supplemented by demographic, academic, and behavioral questions, delivered through the institution's existing student portal or learning management system. Reducing friction at this stage is important: surveys that require separate login credentials, are buried in menus, or take more than 10 minutes to complete will see low completion rates. The system works best when the survey is integrated into regular institutional touchpoints, such as course enrollment, pre-semester check-ins, or periodic wellbeing modules.

The classification layer is the ML model trained on the available dataset. For new deployments, a pre-trained model from a demographically similar population can be used initially, with fine-tuning as institution-specific data accumulates. The model generates a stress tier classification and a confidence score for each student who completes the survey. High-confidence high-stress classifications trigger an immediate referral flag. Lower-confidence classifications across tiers may generate a follow-up survey or a recommendation for a brief wellness check-in.

The recommendation layer maps classification outputs to the tiered recommendation profiles described in Section 7. These recommendations are delivered to students through the institution's communication platform, typically email or app notification, with links to relevant resources. For high-stress classifications, a parallel notification is sent to the institution's counseling service.

The monitoring and audit layer tracks classification distributions over time, monitors the proportion of high-stress flags that result in counseling contact, and provides aggregate population-level data for institutional reporting. This layer is also where bias auditing occurs: regular checks that classification rates are not systematically different for specific demographic groups in ways that would indicate model unfairness.

8.2 Integration Considerations

The most significant practical barrier to deployment is data privacy. Student mental health data is among the most sensitive personal data an institution holds. In the United Kingdom, processing falls under GDPR Article 9, which imposes strict conditions on sensitive personal data processing. In the United States, FERPA governs student records. In India, the forthcoming Personal Data Protection framework imposes similar requirements. Any deployment must include fully informed consent, clear explanation of how data will be used, anonymization of stored data, and an explicit commitment that classification outputs will not be shared with academic staff who have responsibility for student assessment. The last point is critical: if students believe that disclosing mental health difficulties through a survey could affect their academic standing, they will not complete the survey honestly, and the system will not work.

Institutional buy-in from both academic and pastoral staff is equally important. A classification system that generates referral flags but connects to an overwhelmed counseling service, or one that senior academic leadership treats as a surveillance tool rather than a support mechanism, will produce worse outcomes than no system at all. Deployment should be preceded by stakeholder consultation, staff training, and explicit institutional commitment to the proportionate expansion of counseling capacity that a functional early detection system would demand.

9. Ethical Considerations

Mental health AI systems require more careful ethical scrutiny than most other applied ML applications because the stakes of errors and misuse are high and the affected population is vulnerable. The ethical framework for this system rests on four principles.

9.1 Informed Consent and Data Autonomy

Students must understand what data is being collected, why, how it will be stored and processed, who will have access to the outputs, and what the consequences of participation or non-participation are. Consent must be genuinely voluntary: there should be no academic penalty for declining participation and no mechanism by which non-participation is interpretable as a mental health flag in itself. Students should have access to their own classification data and the ability to challenge or request deletion of it.

9.2 Algorithmic Fairness

The SRM University study (India Today, 2025) found that gender and institution type were significantly associated with stress profile. Biswas et al. (2026) found that socioeconomic class affected both stress levels and the predictive weight of different features. These findings mean that a model trained on one demographic population may perform poorly when applied to another, systematically misclassifying members of underrepresented groups. Before deployment, models must be audited for performance disparities across gender, socioeconomic class, ethnicity, disability status, and nationality. Where disparities exist, the training data must be diversified, or the model must be calibrated separately for different population subgroups.

9.3 Clinical Scope

The system classifies perceived stress, not clinical psychiatric conditions. A high-stress classification from this model is not equivalent to a diagnosis of anxiety disorder, depression, or any other recognized clinical condition. Communicating this clearly to both students and institutional staff is necessary to prevent inappropriate use of the system's outputs. Recommendations at the high-stress tier should always include a referral to clinical professionals rather than presenting the system's own recommendations as a substitute for clinical assessment.

9.4 Preventing Surveillance Creep

There is a meaningful difference between a mental health support tool and a surveillance system. The former is designed to benefit the student. The latter is designed to produce information that others can use to make decisions about the student. This system is designed as the former, and its governance structure must actively prevent it from becoming the latter. Classification outputs should not be accessible to academic staff, employer partners, scholarship committees, or any other entity with a decision-making relationship with the student. The system's purpose is to connect students to support, not to produce a mental health record that follows them through their academic career.

10. Limitations and Future Directions

10.1 Current Limitations

The framework described in this article has several limitations that are worth stating directly. First, the primary datasets are relatively small. The Kaggle dataset has 1,100 observations; the Bangladesh dataset has 150. Performance figures from these datasets, while strong, may not generalize to the full diversity of student populations across different countries, institutional types, and demographic compositions. Larger, more diverse training datasets are needed before widespread deployment.

Second, the models are trained on cross-sectional data: they classify current stress level based on current features but cannot predict how stress will change over time. A student who scores as medium stress in October may be at high stress by December as examination pressure builds. Longitudinal tracking, with repeated measurements over the academic year, would substantially improve the system's ability to identify students whose stress trajectory is worsening.

Third, the recommendation engine is currently rule-based, which limits its ability to account for individual variation within stress tiers. Two students in the high-stress tier may have very different stressor profiles: one driven primarily by academic overload, another by financial anxiety, a third by social isolation. Recommendations that address the specific dominant stressor rather than the general stress tier would be more effective, but producing them requires either more granular classification or a natural language generation layer.

10.2 Promising Future Directions

Several research and development directions would materially improve the system's capability and reach.

Multimodal data integration is the most technically promising avenue. Current studies rely entirely on survey data. Adding passively collected signals from wearables, such as heart rate variability, sleep tracking, and physical activity data, would provide continuous monitoring without survey burden. Research published in IEEE Access by Abdelfattah, Joshi, and Tiwari (2025) demonstrated effective stress detection from multimodal physiological data, and the integration of wearable data with survey data in a unified classification model could substantially improve both accuracy and temporal resolution.

Federated learning architectures would allow models to be trained collaboratively across multiple institutions without any institution's student data leaving its own servers. Each institution's local model learns from its own data; model parameters rather than raw data are shared with a central aggregation server. This approach addresses the data privacy concerns that currently create the most significant barriers to multi-institutional deployment.

Large language model integration for recommendation generation is a near-term possibility. Models such as those used to power conversational AI assistants could generate personalized, contextually appropriate recommendation text based on a student's full feature profile rather than just their stress tier classification. This would allow the recommendation engine to address the specific combination of stressors affecting a particular student rather than providing generic tier-level guidance.

Cross-cultural validation studies are needed before this framework can be responsibly deployed beyond the cultural contexts in which the current training data was collected. A model trained primarily on data from Bangladeshi and Turkish students may not capture the stress patterns of students in Japan, Sub-Saharan Africa, or Latin America with equal fidelity. Culturally calibrated stress classification thresholds, and possibly culturally adapted recommendation profiles, would be needed for global deployment.

11. Conclusion

The case for machine learning-based student stress classification rests on a straightforward empirical foundation. Large proportions of students in every country studied are experiencing stress at levels that impair their functioning. Most of them do not seek help. The institutional systems designed to provide that help are not scaled to meet the demand. Something has to change in how early identification of at-risk students happens, because the current model is not working.

The studies reviewed in this article demonstrate that classification systems trained on survey data can achieve accuracies between 86 and 98 percent using well-established algorithms, and that the features driving those classifications, particularly PSS-10 scores, sleep hours, academic performance, and family income, are precisely the variables that mental health researchers and clinicians would predict to be important. The models are doing what theory suggests they should do, and they are doing it reliably.

The recommendation engine described here is a first-generation system. Its rule-based structure is a limitation, but it has an important virtue: transparency. Every recommendation can be reviewed, critiqued, and updated by clinicians and policymakers. It does not operate as a black box making inscrutable decisions about student mental health. That transparency is a design choice, not a technical constraint, and it should be preserved as the system becomes more sophisticated.

What the research consistently shows, from the SRM University data on Indian students to the Essex study of Bangladeshi undergraduates to the Turkish ML study on the Kaggle dataset, is that student stress is measurable, classifiable, and amenable to structured intervention. The technology to support that intervention at institutional scale is available. What has been missing, in most institutions, is the organizational commitment to use it. Girish Rowjee's observation that mental health support is no longer a nice-to-have but a core institutional responsibility reflects a shift in how educational institutions are being asked to think about their obligations to students (India Today, 2025). A classification and recommendation system of the kind described here is one practical piece of the infrastructure those obligations require.

Students are under pressure that previous generations did not face at this scale or intensity. The tools to help them are better than they have ever been. The gap is in deployment, governance, and institutional will. Those are human problems, not technical ones.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1]. Yeler, K. and Sürücü, S. (2025). Classification of Student Stress Levels Using Machine Learning Methods. International Journal of Advanced Natural Sciences and Engineering Researches (IJANSER), 9(10), 210-214. https://as-proceeding.com/index.php/ijanser.
[2]. Biswas, T., Mim, R. K. and Ferdous, F. S. (2026). Machine Learning-Based Detection of Student Stress Levels Through Socioeconomic Status. Essex Student Journal, 17(1). https://doi.org/10.5526/esj.453.
[3]. India Today Education Desk. (2025, October 16). 3 in 5 students in India's top cities struggle with anxiety. India Today. https://www.indiatoday.in/education-today/latest-studies/story/student-mental-health-3-in-5-students-in-india-top-cities-struggle-with-anxiety-2799172-2025-10-07.
[4]. Navrozashvili, M. (2025, September 19). Pre-Semester Student Stress Statistics: Everything You Need to Know. EssayPro Blog. https://essaypro.com/blog/student-stress-statistics.
[5]. Abdelfattah, E., Joshi, S. and Tiwari, S. (2025). Machine and Deep Learning Models for Stress Detection Using Multimodal Physiological Data. IEEE Access, 13, 4597-4608.
[6]. Cohen, S., Kamarck, T. and Mermelstein, R. (1983). A global measure of perceived stress. Journal of Health and Social Behavior, 24(4), 385-396.
[7]. Chen, T. and Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785-794.
[8]. Breiman, L. (2001). Random Forests. Machine Learning, 45(1), 5-32.
[9]. Gao, W., Ping, S. and Liu, X. (2020). Gender differences in depression, anxiety, and stress among college students: a longitudinal study from China. Journal of Affective Disorders, 263, 292-300.
[10]. Gobena, G. A. (2024). Effects of Academic Stress on Students' Academic Achievements and Its Implications for Their Future Lives. Anatolian Journal of Education, 9(1), 113-130.
[11]. Pascoe, M. C., Hetrick, S. E. and Parker, A. G. (2020). The impact of stress on students in secondary school and higher education. International Journal of Adolescence and Youth, 25(1), 104-112.

[12]. Ovi, M. S. I., Hossain, J., Rahi, M. R. A. and Akter, F. (2025). Protecting Student Mental Health with a Context-Aware Machine Learning Framework for Stress Monitoring. arXiv preprint, arXiv:2508.01105.
[13]. World Health Organization. (2023). Mental health of adolescents. WHO Fact Sheet. https://www.who.int/news-room/fact-sheets/detail/adolescent-mental-health.
[14]. Selye, H. (1956). The Stress of Life. McGraw-Hill.
[15]. Lazarou, E. and Exarchos, T. P. (2024). Predicting stress levels using physiological data: Real-time stress prediction models utilizing wearable devices. AIMS Neuroscience, 11(2), 76-102.