International Journal of Creative and Open Research in Engineering and Management
ISSN: **3108-1754 (Online)**
Volume 02 Issue 05 May-2026 | Impact Factor: 3.5

Mental Health Risk Prediction for College Students Using Machine Learning

Sharanya K
Department of CSE
R V Institute of Technology and Management
Bangalore, India
rvit23bcs271.rvitm@rvei.edu.in

Sharvari G
Department of CSE
R V Institute of Technology and Management
Bangalore, India
rvit23bcs315.rvitm@rvei.edu.in

Sheethal V
Department of CSE
R V Institute of Technology and Management
Bangalore, India
rvit23bcs017.rvitm@rvei.edu.in

Shreya BP
Department of CSE
R V Institute of Technology and Management
Bangalore, India
rvit23bcs314.rvitm@rvei.edu.in

Dr. Savitha G
Department of CSE
R V Institute of Technology and Management
Bangalore, India
savithag.rvitm@rvei.edu.in

***Abstract*—**Most college kids find their heads spinning when school work piles up, friendships shift, or home feels too far away. When life gets like that, tension builds - sometimes tipping into worry or sadness strong enough to mess with sleep, focus, meals.Spotting those shifts fast matters because reaching out early tends to lighten the load before it grows heavier. Help at the right moment changes how things unfold later on.

Nowadays more schools turn to number-focused methods. Computers try spotting kids who might struggle using math tools hidden inside programs. One way uses past choices to guess future ones another draws lines between types of behavior one sorts outcomes by closeness while others grow like trees from decisions made before. Some work better than others depending on what you need found out. How well each does its job shows which fits best when results matter.

One way this research worked began by testing several prediction methods for student mental health concerns. Not until the data passed close review did any modeling start - only then could analysis move forward. Effectiveness of every method came into view through common evaluation metrics, each revealing different results.

Surprisingly, Logistic Regression stood out with better accu-racy than fancier models. That hints at how straightforward approaches might work just as well when spotting possible mental health issues. With an eye on obvious trends instead of heavy math, the method stays quick and easy to follow.

One thing stands clear: machines might help spot trouble before it grows. Spotting signs earlier means schools can step in faster - less guesswork, more care where it counts. Help arrives when needed, easing strain on student minds everywhere

Index Terms—College students; Mental health; Machine learn-ing; Logistic Regression; Early detection; Predictive modeling

I. Introduction

Lately, more college students face tough emotions like sadness, worry, or constant pressure. Life at university brings heavy workloads, shifting friendships, because adjusting takes effort. Trouble inside the mind shows up in grades, how someone talks with peers, even future job paths [1], [3]. Research after research finds many learners struggle quietly, often without help close by. Spotting signs early matters, since support can shift outcomes slowly. Heavy thoughts pile up when no one notices, especially during busy semesters filled with change.

Looking back, old-school ways to spot mental struggles lean heavily on questionnaires filled out by people themselves, talks with specialists, along with mind-focused evaluations. Even though such techniques hold worth, they take up much time, depend too much on personal views, plus sometimes miss catching issues before they grow. Another angle - more folks need help now than ever, yet experts who handle emotional well-being remain few in number, making fast support hard to reach [6], [11]. Because of this shift, interest grows around systems built to process information automatically, ones that might flag warning signs earlier through patterns found in data.

Among tools born from artificial intelligence, machine learning stands out when it comes to sifting through vast amounts of data, spotting subtle links tied to mental health. Because researchers now lean heavily on this approach, pre-dictions about psychological states often rely on factors like age, school results, daily routines, or how people act [7], [10]. Though methods differ, some models—like SVM, RF, LR, and certain neural networks—show consistent skill in forecasting outcomes, backed by multiple investigations into their accuracy [2], [8].

Take Hasan and team [1], they used various machine learning approaches to forecast depression, anxiety, and stress in college students, showing how social background and daily habits matter. In a similar way, Vaishnavi's work [2] tested several algorithms - ensemble ones stood out with better results. Some researchers turned to powerful tools like XGBoost, CNNs, or mixed-model setups aiming at stronger predictions [5], [12]. Even so, hurdles stick around: uneven data splits, poor model transfer to new groups, and difficulty understanding what intricate systems are actually doing [4], [8].

Spotting issues early turns out to matter a lot, according to research. Instead of waiting, schools might use machine learning to flag students who struggle, offering help ahead of time. Yet most current approaches lean on heavy computing power along with massive data piles - tough to manage where tools and info are limited [6].

II. Literature Review

More students at universities face mental health issues now than before, pushing scientists to explore how data can spot problems sooner. Depression, anxiety, and stress show up often in surveys, weighing down grades and daily living. Instead of only asking people how they feel or waiting for doctors' appointments, researchers look elsewhere. These usual ways sometimes miss signs because feelings differ and help comes late. Machines that learn patterns find clues faster, noticing risks before things get worse. Early warnings could come from behavior shifts machines detect without anyone saying a word. Machine learning has gotten better at handling tough, layered data about mental health. Because they're good at sorting and forecasting psychological states accurately, methods like Support Vector Machines, Random Forest, Logistic Regression, and K-Nearest Neighbors see common use. Features ranging from daily habits to school results, age, and actions help these systems spot links between personal traits and emotional well-being.

Most studies look at how different machine learning models stack up against one another when making predictions. Not every method wins all the time - results shift based on data traits and what features are picked [1], [5]. Instead of relying on just one, some approaches mix several basic predictors to get better results; Random Forest and boost-like strategies fit here. These blended systems tend to handle new data more reliably. On their own, straightforward tools like Logistic Regression still work well, especially with organized information - not flashy, yet clear and fast.

In addition to model selection, the identification of key influencing factors plays a crucial role in improving prediction accuracy. Research indicates that mental health conditions among students are influenced by a combination of academic, social, and behavioral factors, including stress levels, family relationships, financial conditions, physical activity, and sleep patterns [3], [7], [12]. Incorporating these features into predictive models enables a more comprehensive understanding of mental health dynamics and improves classification performance.

What stands out in recent studies is how cleaning data matters, especially when one group has far fewer examples than others. Picking key traits, scaling numbers evenly, then boosting smaller groups - say with SMOTE - often lifts a model's accuracy [4]. On top of that, stacking smart layers like neural nets helps untangle messy patterns hidden across information streams. Systems using grid-like filters or memory cells manage better guesses, mainly if data piles are huge, mixed up, yet full of signals [8], [10].

Even with progress, problems still show up when using machine learning to predict mental health outcomes. Heavy reliance on massive data sets and complex calculations often gets in the way of practical use outside labs. On top of that, understanding how decisions are made by models stays cloudy - this confusion grows worse when data varies between groups. When systems do not clearly reveal their reasoning, hesitation follows, especially where privacy and care meet.

Facing these issues, clearer methods are emerging - ones that weigh accuracy against ease of understanding. This research steps in using a side-by-side analysis of several algorithm types, judging each through measures like ROC-AUC. Instead of chasing complexity, it leans toward models that deliver clear results without heavy computation. Through this lens, student mental health forecasts become more usable in real school environments. Progress here supports earlier help systems where timing matters most. What counts is not just speed but how well insights can be grasped by those applying them.

A. Research Gaps and Objectives

Based on the review of existing literature, the following research gaps are identified:

- Even with big steps forward in using machine learning to predict mental health issues, some key holes still exist in current studies. Because of these missing pieces, better models are needed - ones that work well, make sense to people, and can actually be used in real settings.
- Right off the bat, a bunch of current research leans heavily on intricate systems like Random Forest, XGBoost, CNNs, along with combined setups to boost how well predictions turn out [5], [8], [12]. Though they tend to deliver solid results, these approaches usually depend on massive amounts of data plus serious computing power - making them tough to deploy practically, particularly within settings where resources are tight, say schools or colleges.
- Most research tries boosting precision using combined models along with complex approaches - yet skips clear explanations of how decisions are made [2], [6]. When dealing with areas such as mental health, seeing which

factors matter most builds confidence while supporting useful conclusions. Without transparency, professionals struggle to grasp why a result was reached, limiting real-world use in choices that affect people.

- One big problem seen in past research? Models often depend too much on uneven data. Some papers use tiny collections of information. Others pull details only from certain places or hospitals - making results hard to apply widely [1], [4]. Even when methods like SMOTE try fixing skewed samples, they pop up only now and then across experiments. That spotty use might tilt predictions without anyone noticing.
- One more thing - results are hard to compare because each study measures success differently. Some rely on unique data sets, others pick separate ways to test accuracy. Because of this mix, spotting the best method for predicting mental health outcomes gets tricky. Without common rules, progress stays scattered.

To address these gaps, the objectives of this work are defined as follows:

- Start clean. Missing bits get filled smart. Categories turn into numbers instead. Pick only what matters most here. Stronger signals show up then. Noise fades behind. Each choice shapes better guesses later. Quality jumps when steps stick together. Outcomes grow clearer each pass. Predictions gain trust this way.
- Looking at how stress, anxiety, depression, burnout, and mental health connect might reveal hidden links. Patterns start showing up when these factors are viewed together. Instead of focusing on one alone, seeing them side by side brings clarity
- One way to predict mental health issues is by using Logistic Regression, while another tries Support Vector Machines. A different path takes Random Forest into account, whereas K-Nearest Neighbors follows patterns nearby. Each method gets tested side by side so results can show which works best. Performance differences become clear only after running them all through similar cases.

III. System Architecture and Methodology

A design meant for predicting mental health trends uses a step by step flow. Starting with raw information entering the system, it moves into cleaning and adjusting that data. Then comes a close look at patterns hiding inside. Important clues get pulled out next. These guide the learning phase where models adapt through examples. Testing checks how well they perform. Results appear as forecasts ready for practical judgment. Each stage links smoothly to keep outcomes sharp, useful when applied beyond theory

Right away, information about learners flows into the system - things like background details, school performance, along with signs of stress, how anxious they feel, levels of sadness, emotional exhaustion, time spent studying, and hours asleep. Such factors help make sense of a student's inner well-being, having shown up often in earlier research [3], [7].

After arriving, the data moves into cleanup. Missing pieces get filled or dropped. Errors inside it vanish one by one. Words turn into numbers so math can follow. Everything lines up straight now. Ready for patterns to show. Machine tools need things this way.

After cleaning comes a close look at how variables connect. Patterns start showing up when numbers talk through charts. Some pieces of data stick together more than others. Spotting these links guides what gets picked later. Choices made here shape what the system learns down the line.

Later on comes model training - this step puts chosen features to work. Instead of just one method, several take turns: Logistic Regression tries its way, then Decision Tree gives it a go, followed by Random Forest joining in. Every approach studies past examples differently. From those lessons, each builds its own sense of how mental health cases might sort out.

After training, each model gets checked with scores like ROC-AUC to see how well it works. Though several methods are tested, one stands out when results are lined up side by side. Because it handles predictions better - plus makes decisions easier to follow - Logistic Regression ends up being the top choice.

Fig. III.1. Proposed machine learning pipeline for student mental health prediction..

A fresh look at spotting student mental health issues emerges through machine learning tools. Starting with raw data, the process moves step by step - first shaping information, then digging into patterns hidden beneath. Instead of rushing ahead, it pauses to highlight key factors that matter most. After dividing data wisely, models take form through careful setup and adjustment. Each stage builds trust in results while echoing earlier work done in this space. Solid structure guides every choice, linking today's method to past insights found in similar studies.

A. Data Description

A collection of details about students makes up the data here - things like background traits, school performance, feelings of stress, scores on anxiety, how down they feel, signs of burnout, time spent studying, amount of sleep, plus overall mind wellness. Because these pieces shape how mentally sound learners might be, they're pulled into view. Past research has pointed out their weight when guessing emotional states [3], [7].

B. Data Preprocessing

Before any modeling begins, raw data gets cleaned up so it works better with algorithms. Missing pieces are filled in or adjusted, odd entries get tossed out, while word-based categories turn into numbers. Cleaning things early cuts down on confusion later. When data flows smoothly into models, results tend to come out sharper - especially when guessing outcomes in mental health cases [4].

C. Exploratory Data Analysis (EDA)

Looking at the data helps spot how each feature connects to mental health. Instead of just listing numbers, a color-coded map reveals links between variables. Figure 1 displays clear patterns - stress ties closely to anxiety, which also pulls along depression and burnout. On the flip side, higher stress levels match up with lower scores on the mental health index.

Further, pairwise relationships between features are visualized using a pairplot, as illustrated in Fig. 2. This helps in identifying patterns, distributions, and dependencies among variables.

Scatter plots help spot patterns between key variables. Looking at Fig. 3, longer study times link to rising burnout scores. Then again, Fig. 4 reveals more sleep connects to stronger mental well-being. Results like these echo earlier research [7], [12]

float placeins

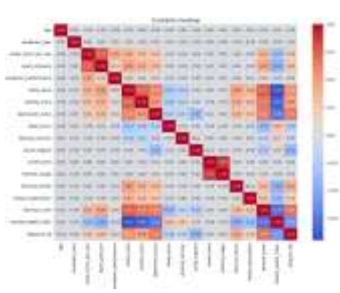

Fig. III.2. Correlation heatmap showing relationships among features in the mental health dataset.

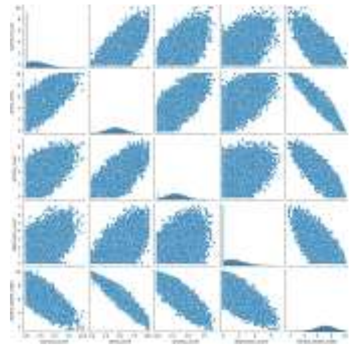

Fig. III.3. Pairwise relationships among key features including burnout, stress, anxiety, depression, and mental health index.

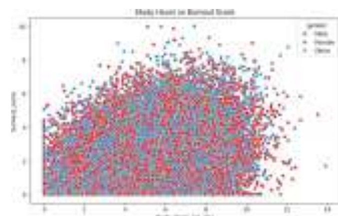

Fig. III.4. Scatter plot showing the relationship between study hours per day and burnout score.

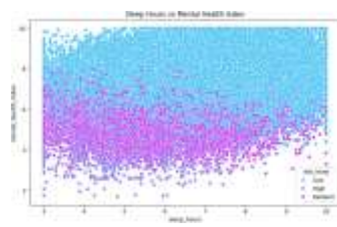

Fig. III.5. Scatter plot illustrating the relationship between sleep hours and mental health index.

D. Feature Selection

Picking key traits happens first when forecasting mental well-being. Stress levels show up clearly, so they stay in. Anxiety scores matter just as much, yet some numbers get left behind if they repeat others. Study time links closely to outcomes, which makes it useful. Sleep duration stands out during review, even though not every habit counts. Cutting noise sharpens the model without making it too complex. Burnout measurements add weight, while weak signals fade away.

E. Train-Test Split

Most times, folks set aside part of the data to check how well the model works later. One chunk trains the system, while another checks if it learned right. Splitting things like eighty percent here, twenty there happens pretty often. Seeing how it handles fresh examples helps tell if it truly gets the pattern. That way, when new stuff shows up, it does not fall apart.

F. Model Pipeline

From raw data to final predictions, each stage flows into the next without gaps. Moving through cleaning, then choosing key traits, keeps noise out early. Training begins only after patterns are sharpened and inputs set right. Results stay steady because every run follows identical logic paths. Skipping steps isn't an option when outputs must match across tests.

G. Models Used

A handful of standard machine learning methods get tested - Logistic Regression, Decision Tree, then Random Forest. These approaches pop up often when sorting data into categories, especially in studies about mental well-being [2], [5], [8]. Training happens on one portion of the data, while results come from checking scores like accuracy and ROC-AUC.

Out front, Logistic Regression edges out the rest when it comes to ROC-AUC scores across tested models. Because it's straightforward, easy to follow, and light on processing needs, this method fits well into actual use for predicting mental health outcomes.

Putting it all together, the method links cleaning data, studying features, training models, then checking results - building a clear, working tool for predicting mental health outcomes.

IV. RESULTS AND DISCUSSION

Here comes how well the machine learning tools worked when guessing student mental health issues. Tools tested were Logistic Regression, Decision Tree, along with Random Forest. Most attention went to ROC-AUC scores - those numbers show separation power in yes-or-no predictions [4], [12].

A. A. ROC-AUC Analysis

Starting off, each model gets checked through ROC curves - these show how well the true positives rise while false positives shift, depending on where you set the line. Instead of just one point, we look at the whole sweep. What stands out is that ROC-AUC sums up how sharply a model tells classes apart, no matter the cutoff chosen.

According to Fig. 5 below, the ROC curve comparison of the following classifiers is provided, namely Logistic Regression, Decision Tree, and Random Forest. As can be seen from Fig. 5 above, the ROC curve corresponding to Logistic Regression gives the highest AUC value. Consequently, it can be concluded that the classifier having the highest ability to classify mental illnesses is Logistic Regression. Moreover, it can be argued that the data used in this study contains characteristics that are best represented using a linear model. This assumption is consistent with previous research [4].

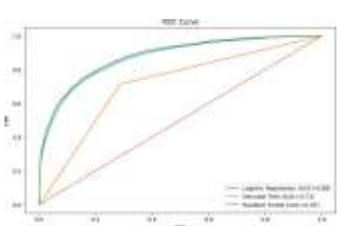

Fig. IV.1. ROC curves of Logistic Regression, Decision Tree, and Random Forest models.

B. Prediction Distribution Analysis

Looking beyond ROC analysis, the forecasted mental health results get examined to see how the model sorts various conditions. Grouped by level of risk, each prediction shows where individuals fall mentally.

Looking at Fig. 6, we see how mental health risk levels are spread out. The predictions cover multiple groups, which means the model picks up on differences in student well-being. Because results appear across categories, one might say it avoids favoring just one outcome. Balance shows through - no category dominates the output. That evenness hints the system treats each level with similar weight.

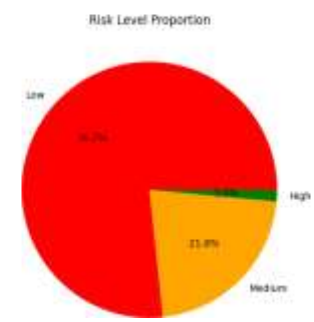

Fig. IV.2. Distribution of predicted mental health risk levels.

C. Discussion

Out front, Logistic Regression takes the lead when it comes to ROC-AUC scores - beating both Decision Tree and Random Forest. Even though methods like Random Forest usually shine by handling tricky patterns, here the data's clear layout gives an edge to less complicated approaches.

Logistic Regression stands out in how well it tells apart student groups when checked through ROC analysis. Instead of just lumping everyone together, its predictions spread clearly across varying mental health risk tiers. Earlier studies back this up - what matters most is the data itself, not how fancy the method looks. Simpler models can win if the dataset fits right. Complexity doesn't always mean better results.

Most times, clarity matters more than complexity - especially when it comes to mental health tools. A straightforward method like Logistic Regression makes it easier to see how decisions are formed. Because of this transparency, teams can trust and apply the outcomes with more confidence. Simplicity here does not mean weakness; instead, solid results come through clear design. Even without intricate math, accuracy stays strong. What works well often stands out by being understandable first.

Apart from slight flaws, the model handles mental health prediction well - showing promise when applied to students through smart algorithms. Its ability to spot issues early comes through clearly, thanks mainly to data patterns caught on time. Not perfect, yet strong enough to suggest tech can play a helpful role where support is thin.

V. CONCLUSION

Looking at how well machines can guess student mental health issues, this work tests several ways to sort people into groups. After cleaning numbers and checking trends, one method stood out clearly - Logistic Regression scored highest on ROC-AUC. Instead of jumping to conclusions fast, it weighed each clue carefully before deciding. Other models tried hard but missed subtle signals that mattered most. What made the difference was not speed, yet precision when handling messy real-world answers. Each step - from shaping raw inputs to picking key traits - shaped better guesses down the line. Accuracy didn't spike by chance; careful tuning held everything together.

Surprisingly, basic models often match complex ones on structured data. Though many now rely on tools like Random Forest for mental health forecasts, here Logistic Regression holds its ground - sometimes outperforms them - using far less computing power. Oddly enough, what matters most isn't how fancy a model looks. Performance ties closely to the nature of the data itself, not layers of complexity - a point earlier work already hinted at.

Notably, looking at how predictions play out shows the model can sort students by mental health risk level - opening doors for earlier help. In this area, applying machine learning brings a method shaped by data, one that grows with demand, tackling student well-being issues more directly, much like earlier work pointed out.

Even though outcomes look good, the research faces some constraints - reliance on existing data stands out, along with a narrow range of features. Moving ahead, efforts might shift toward using broader, richer datasets instead of sticking to current ones. Advanced methods for shaping inputs could replace basic approaches, opening room for better results. Hybrid systems or deep learning paths may offer stronger predictions than today's models. Real-time information streams might slot into frameworks where static data once ruled. On top of that, simpler interfaces could help schools adopt these tools without friction.

Machine learning proves useful for predicting mental health outcomes - this study shows why picking the right model matters just as much. Performance alone isn't enough when clarity and speed play key roles in real-world use. What works best often depends on how well it explains results without needing heavy computing power. The findings highlight a path forward where practicality shapes choices, not just accuracy. Each decision affects how easily tools can be used outside labs.

VI. FUTURE WORK

Even so, the suggested system handles student mental health predictions well, yet there remain paths forward worth exploring. Progress could take different shapes beyond what has been tested already.

One thing though - the data here isn't very big or wide in range. Down the line, research might pull together broader information from various places and populations instead.

Later on, trying more complex ways to shape data might help models work better. Instead of stopping at basics, adding traits like how someone uses social media, moves throughout the day, or sleeps could reveal hidden trends in mental well-being. Earlier studies have pointed out that mixing these varied signals often leads to sharper predictions.

Later on, instead of sticking to basic algorithms, upcoming work might look into stronger models like combined systems or layered networks. Methods including Random Forest or XGBoost, along with CNNs and LSTM structures, already proved useful when guessing mental health outcomes. Because they handle tangled data better, these tools could lift accuracy a step higher.

One way forward might involve tackling uneven data distribution. Another angle could look at refining how models are tested over time. Instead of just accepting skewed datasets, researchers may try approaches like SMOTE to balance categories. Cross-checking results across different subsets often helps too. These steps tend to make predictions more stable. Outcomes usually show less tilt toward dominant classes. Source points to number four for deeper detail.

Looking ahead, building apps that work instantly matters a lot. When systems go live on phones or browsers, they watch how students feel over time. Spotting shifts early helps guide support fast. Need grows steadily for tools that reach more people without hassle. Reaching wider means making help easier to get.

One step at a time, clearer models matter most when dealing with mental well-being. Down the line, digging into explainable AI could reveal how each detail shapes predictions. This shift might build stronger confidence and real-world use by those on the front lines.

One step ahead, growing the data pool could sharpen results. Not far behind, smarter models might catch patterns now missed. Right alongside, live processing may make responses quicker. Tied into all this, clearer reasoning behind outputs helps users trust what they see.

Notably, looking at how predictions play out shows the model can sort students by mental health risk level - opening doors for earlier help. In this area, applying machine learning brings a method shaped by data, one that grows with demand, tackling student well-being issues more directly.

Even though outcomes look good, the research faces some constraints - reliance on existing data stands out, along with a narrow range of features. Moving ahead, efforts might shift toward using broader, richer datasets instead of sticking to current ones. Advanced methods for shaping inputs could replace basic approaches, opening room for better results. Hybrid systems or deep learning paths may offer stronger predictions than today's models. Real-time information streams might slot into frameworks where static data once ruled. On top of that, simpler interfaces could help schools adopt these tools without friction.

Machine learning proves useful for predicting mental health outcomes - this study shows why picking the right model matters just as much. Performance alone isn't enough when

clarity and speed play key roles in real-world use. What works best often depends on how well it explains results without needing heavy computing power. The findings highlight a path forward where practicality shapes choices, not just accuracy. Each decision affects how easily tools can be used outside labs.

REFERENCES

[1] M. E. Hasan, M. Arif, S. M. R. Hasan, et al., "Prevalence, associated factors, and machine learning-based prediction of depression, anxiety, and stress among university students: A cross-sectional study from Bangladesh," Journal of Health, Population and Nutrition, vol. 44, 2025.

[2] K. Vaishnavi, U. N. Kamath, B. Ashwath Rao, and N. V. Subba Reddy, "Predicting mental health illness using machine learning algorithms," Journal of Physics: Conference Series, vol. 2161, no. 1, 2022.

[3] X.-Q. Liu, Y.-X. Guo, W.-J. Zhang, and W.-J. Gao, "Influencing factors, prediction and prevention of depression in college students: A literature review," World Journal of Psychiatry, vol. 12, no. 7, pp. 860–873, 2022.
.

[4] V. Wijaya and N. Rachmat, "Comparison of SVM, Random Forest, and Logistic Regression performance in student mental health screening," Journal of Electrical Engineering and Computer Sciences, vol. 9, no. 2, pp. 173–184, 2024.

[5] Q. Li and J. Zhou, "A comparative analysis of extreme gradient boosting, decision tree, support vector machines, and random forest algorithm in data analysis of college students' psychological health," Informatica, vol. 49, pp. 127–134, 2025.

[6] J. Chung and J. Teo, "Mental health prediction using machine learning: Taxonomy, applications, and challenges," Applied Computational Intelligence and Soft Computing, vol. 2022, Article ID 9970363, 2022..

[7]] H. A. Rahman, M. Kwicklis, M. Ottom, et al., "Machine learning-based prediction of mental well-being using health behavior data from university students," Bioengineering, vol. 10, no. 5, 2023.

[8] U. Madububambachu, A. Ukpebor, and U. Ihezue, "Machine learning techniques to predict mental health diagnoses: A systematic literature review," Clinical Practice Epidemiology in Mental Health, vol. 20, 2024.

[9] Y. Huang, S. Li, B. Lin, et al., "Early detection of college students' psychological problems based on decision tree model," Frontiers in Psychology, vol. 13, 2022.

[10] Y. Fu, F. Ren, and J. Lin, "Apriori algorithm based prediction of students' mental health risks in the context of artificial intelligence," Frontiers in Public Health, vol. 13, 2025.

[11] U. Madububambachu, U. Ihezue, A. Ukpebor, and S. Ayo, "Predictive machine learning approaches for mental health diagnoses in college students," International Journal on Engineering Technologies and Informatics, vol. 5, no. 3, 2024.

[12] Y. Zhai, Y. Zhang, Z. Chu, et al., "Machine learning predictive models to guide prevention and intervention allocation for anxiety and depressive disorders among college students," Journal of Counseling Development, vol. 103, pp. 110–125, 2025.