A MULTIMODAL APPROACH TO MENTAL HEALTH PREDICTION: MACHINE LEARNING PERSPECTIVE ON LIFESTYLE AND BEHAVIOURS

A Thesis Submitted
In Partial Fulfillment of the Requirements for the degree of

MASTERS OF TECHNOLOGY
in
Data Science

by
Renuka Sutone
(2K23/DSC/07)

Under the Supervision of

Prof. Ruchika Malhotra

Head of Department, Department of Software Engineering

Delhi Technological University

DEPARTMENT OF SOFTWARE ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi 110042

May, 2025

DEPARTMENT OF SOFTWARE ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi 110042

ACKNOWLEDGEMENT

We wish to express our sincerest gratitude to Prof. Ruchika Malhotra for her continuous guidance and mentorship that she provided us during the project. She showed us the path to achieve our targets by explaining all the tasks to be done and explained to us the importance of this project as well as its industrial relevance. She was always ready to help us and clear our doubts regarding any hurdles in this project. Without her constant support and motivation, this project would not have been successful.

Place: Delhi
Date: 20.05.2025

Renuka
Renuka Sutone
(2K23/DSC/07)

DEPARTMENT OF SOFTWARE ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of
Engineering) Bawana Road,
Delhi 110042

CANDIDATE'S DECLARATION

I, Renuka Sutone, 2K23/DSC/07 students of M.Tech (Data Science), hereby certify that the work which is being presented in the thesis entitled "A Multimodal Approach to Mental Health Prediction: Machine Learning Perspective on Lifestyle and Behaviour". in partial fulfilment of the requirements for the award of degree of Master of Technology, submitted in the Department of Software Engineering, Delhi Technological University is an authentic record of my own work carried out during the period from Jan 2025 to May 2025 under the supervision of Prof. Ruchika Malhotra.

The matter presented in the thesis has not been submitted by me for the award of any other degree of this or any other institute.

Renuka

Candidate's Signature

This is to certify that the student has incorporated all the corrections suggested by the examiners in the thesis and the statement made by the candidate is correct to the best of our knowledge.

Ruchika Malhotra

Signature of Supervisor (s)

DEPARTMENT OF SOFTWARE ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of
Engineering) Bawana Road,
Delhi 110042

CERTIFICATE

I hereby certify that the Project Dissertation titled "A Multimodal Approach to Mental Health Prediction: Machine Learning Perspective on Lifestyle and Behaviour" which is submitted by Renuka Sutone, Roll No – 2K23/DSC/07, Department of Software Engineering, Delhi Technological University, Delhi in partial fulfilment of the requirement for the award of the degree of Master of Technology, is a record of the project work carried out by the students under my supervision. To the best of my knowledge this work has not been submitted in part or full for any Degree or Diploma to this University or elsewhere.

Ruchika Malhotra

Place: Delhi

Prof. Ruchika Malhotra
Head of Department

Date: 20.05.2025

Department of Software Engineering, DTU

ABSTRACT

Depression in particular, is now a major and increasing concern for students everywhere. These problems are usually made worse by students having to deal with pressures from school, society, and changes in their life, all of which leave them more at risk. The use of self-report surveys and periodic counseling tends to make it challenging to notice depressive symptoms that are not easy to detect. To deal with the drawback explained above, the current analysis presents an advanced system that relies on different machine learning (ML) models. Random Forest, XGBoost, Logistic Regression, LightGBM, Support Vector Classifier (SVC), and K-Nearest Neighbors (KNN) are some of the models used in this system to detect depression in students.

To build the predictive models, we used large numbers of distinct features coming from demographics, academic results, and daily routines. To tackle the class balance problem, SMOTE was used, making sure the data in the two classes was more even. Carrying out many preprocessing tasks helped the models to perform higher quality work. To do this, missing values were taken care of through appropriate imputation, duplicate or inconsistent data was removed, filtering based on values was used to keep the important parts, and statistical methods were used on outliers to reduce errors in the dataset.

Evaluating how accurate the model performed showed that both XGBoost and LightGBM were 97% accurate. The study proves that machine learning algorithms can help identify early mental health problems, even when the symptoms are not very noticeable. The results supports the use of AI-based approaches to early identify mental health risks and offer the relevant guidance for supported interventions to help students that are affected by depression.

More studies will attempt to connect digital behavior information with AI algorithms to make the predictions more accurate. Moreover, XAI elements will be used to make the model transparent and instill trust in its users. Testing the system among students from different cultures and countries is essential to find out if it can be used by people the world over.

TABLE OF CONTENT

List of Figures

List of Tables

CHAPTER 1

INTRODUCTION

In light of the worldwide mental health problem and the shortcomings of the tools we currently use to identify problems, it becomes really important to figure out what the main challenges are that make it hard to spot and treat mental health issues among students early on. The following problem statement explains the exact areas that this research is trying to fix.

1.1 Problem Statement

Despite more people knowing about the problem and schools trying to help, students with depression are often not noticed and not given enough help because the usual ways of checking mental health have some limits. Screening tools in use now are inflexible to culture, do not keep progress over time, and depend on people's opinions, which explains why many cases are not identified. Resources available in schools through the clinical system are not sufficient to monitor the mental health of large groups. It is vital to have data-based, scalable, and easy-to-understand models that identifies the early signs of depression in students from different backgrounds. This research aims to fill that gap by using machine learning to build a clear and useful way of making diagnoses that works well for the changing needs and situations students face in school and life.

1.2 Global Mental Health Burden

Mental health disorders are now one of the biggest health issues today, especially among young people, causing a lot of problems for us from a social and financial point of view. According to the World Health Organization, around 970 million people around the world were dealing with these problems in 2019. These conditions, mainly depression and anxiety, have become more common in teenagers and young people, which makes them a difficult problem for society and for doctors to deal with. Data from the WHO shows that in 2019, a significant number of people (970 million) worldwide were struggling with mental health disorders . As the COVID-19 pandemic and global inequality have become more serious, the number of psychological problems faced by people in low- and middle-income countries has increased, since access to healthcare there is limited [1].

WHO anticipates that in the future, depression and anxiety will be key factors causing more than 12% of global disability, most severely affecting teens under 20 . The consequences of untreated mental health issues aren't limited to people's health and well-being, but extend to significant wealth losses due to lower productivity, frequent absenteeism, and raising the need for This global crisis shows just how much mental health support is needed, especially for students and young people, and it's really important that people use methods that really work and can be used on a bigger scale.

The WHO's 2013–2030 Mental Health Action Plan notes that about 90% of people living in poorer countries who need help for mental health problems don't get enough care mostly because there aren't enough trained mental health workers and the health care system isn't set up well enough. For this reason, public health bodies must choose innovative ways to identify and address health concerns that are flexible for use with everyone.

1.3 Limitations of Traditional Screening Tools

Most screenings and diagnostics for mental health uses questionnaires filled out by patients like the PHQ-9 or the GAD-7.In any case, these tools are commonly accepted and used across hospitals and universities, perhaps because they serve some purposes.

It has been observed that both the PHQ-9 and GAD-7 are not very effective with ethnically mixed or non-Western populations, with sensitivities reported between 58% and 65% [2]. Thus, those who have significant problems may not be treated right away, since they are often viewed as perfectly healthy. The PHQ-9 was found to have little use in a Peruvian study since its best cut-off score for depression was found to be seven, which may not meet global expectations.

These screening tools only provide a single picture of a person's mental state, rather than tracing the development of symptoms over time. This means they can't provide a complete view of mental health changes over a long period, so they aren't good for taking action in advance. Patients may not accurately report their complaints for different reasons such as the wish to look better or forgetting to mention some details [3].

It is only when care is advanced that depression is recognized which limits the useful interventions doctors can provide. Because of shorter appointments, insufficient mental health screenings and a lack of mental health skills among primary care doctors, up to 30% of at-risk patients may be overlooked. Therefore, these issues give rise to the need for new and flexible approaches, and this is where machine learning can be useful.

1.4 Machine Learning Opportunities in Predictive Psychiatry

Machine learning is a possible solution that can improve mental health screening tools. By using complex algorithms that can find hidden patterns in lots of different kinds of data, ML models can look at big amounts of information about how people act, how they do in school, and how they live day-to-day, to spot warning signs of mental health problems early on.

Unlike the other traditional methods, ML frameworks use data that was collected for a longer period and spot any changes that may happen, allowing us to understand someone's mental health more precisely. As a result, it is now possible to diagnose care patients early and create personal plans for their treatment. Incorporating ML in psychiatry allows doctors to use more objective information from smartphones and wearables and less on what people verbally describe about their feelings.

Though ML could help in predictive psychiatry, the field is still just beginning to use it. Gaining enough suitable data to train a model, severe privacy concerns and humans have difficulty explaining how the machine learning models work are some of the problems remaining in the field. Since the workings of deep learning models are not transparent to clinicians, their trust is often reduced and they are not used as much in clinical settings.

1.5 Research Gaps

Two significant issues are still present in the present studies on machine learning for mental health prediction:

• **Longitudinal Modeling:** Existing ML frameworks for mental health data do not take into account the changes that can develop gradually with time. This makes it hard for them to spot early signs of depression or predict a person's decline over time [3].

• **Interpretability Trade-offs:** While deep learning is useful for making correct predictions, it is less likely to be used in clinics due to its lack of clarity. The interpretability of AI results is important for winning the trust of participants and ensures that AI is used ethically.

1.6 Contribution

As a result, this study helps compensate for these gaps and supports the progress of AI-based mental health diagnosis.

• **Feature-Driven Predictive Framework:** The authors suggest a way to predict trends using several forms of machine learning. Student responses from various areas were used to point out several key aspects, including their sleep, difficulty with their schoolwork and their physical health. Research revealed that a student's sleeping and learning stress are the most important factors, with levels of 22.1% and 18.7% respectively, in determining whether someone is more likely to suffer from depression which supports what scientists had already discovered.

• **Comparative Model Evaluation:** We use seven types of machine learning algorithms such as XGBoost, Random Forest, Logistic Regression, LightGBM, SVC, KNN, and ensemble techniques. The accuracy, precision, recall, and F1-score were thoroughly checked for all of the models. XGBoost turned out to be the highest performer, reaching an accuracy rate of 85%. Using the SMOTE technique proved to be effective in handling the imbalances in depression labels. Stratified cross-validation was used to guarantee that the model would work equally well for students from many different areas in India.

• **Practical Implementation Design:** It is designed to smoothly blend into any type of IT system learning organizations use now. The models are set up to analyze real-time data from a student's daily or weekly activities. Furthermore, we add SHAP (SHapley Additive exPlanations) values to help make our results easier to understand, so teachers and school leaders can clearly see why the model predicted what it did. It is designed to adhere to ethical AI and can be used as a model for conferences and papers published in Scopus and IEEE.

The findings of the studies that are suggested are helpful for integrating machine learning into how we deal with student's mental health. By bridging this gap between predictive analytics and actionable support mechanisms, it equips educational institutions with a scalable and ethically approach to identify at-risk students and deploy timely interventions.

CHAPTER 2

RELATED WORK

Several important strides have been made in artificial intelligence for mental health through the use of machine learning (ML), hybrid learning tools, and tools based on natural language processing (NLP). Because of these technologies, supporting students' mental health happens in a different way now compared to before. In many studies, when data and good computer algorithms are put together, they can detect warning signs of mental health problems which aids in forming better support systems for people.

Yang and colleagues [4] came up with a new approach for highlighting psychological stress in students during health emergencies. Using ES-ANN in combination with LOF helped them achieve much better results than when they used models like Random Forest and Decision Tree classifiers. The value here is that, unlike normal surveys, actions become the basic evidence rather than people's claimed actions. This research pointed out the need to detect any unusual mental health behaviors promptly. That is the main reason why this study gives importance to powerful preprocessing and a large collection of samples, as they help discover meaningful patterns in mental health assessments among students.

Baba and Bunji [5] turned to LightGBM to look at university student mental health survey results collected in several years. Although they conducted their research by looking at response times, it still confirmed the traditional answers about student anxiety and future academic problems. According to their findings, LightGBM did well in the study because it handles much data efficiently and without frequently overfitting. Moreover, researchers found that such models can help maintain their usefulness, allowing digital mental health screenings to be used often.

Also, Goutam et al. [6] examined using hybrid ML methods, including SHAP which achieved a high accuracy of predicting anxiety disorders, PTSD and depression at 89.6%. They recommend that AI developed for use in healthcare should offer effective predictions as well as be easy to understand. Because of SHAP analysis, they could understand the way sleep, difficulties at school and habits impact the predictions the model made. In line with future aims, the methodology ensures that people affected by the decisions rely on and understand the results from using the diagnostic tool in mental health.

The authors of [7] used a shared model that included EEG patterns as well as other non-EEG features, along with CatBoost and XGBoost to train their ensemble models. It is shown that using two types of information improves the accuracy of emotion recognition by about 93.1%. It explains that combining different data types can help improve mental health diagnostics for students, so this approach should be considered in future systems involving wearables and biometrics.

Ku and Min [8] analysed how predictions made by machine learning models can be affected by uncertainty in human replies in the field of mental health predictors. They found that among the five algorithms, CNNs were the most resistant to changes and noise in the input information. It is especially useful since the data that has been used from students can be subjective and change from student to student responses. Results from their research support why algorithm stability should be important, so both Logistic Regression and more advanced models were used in the current study.

Now, transformer-based methods in Natural Language Processing are helping us to improve mental health prediction analysis by detecting both feelings and sentiments. According to Ibitoye et al. [9], CETM (Contextual Emotion-based Transformer Model) was introduced to analyze user generated content by giving emotional attention to each part of the content. RoBERTa implemented by Baidu was found to be more effective than baseline BERT, reaching 94.5% accuracy. They are powerful models because they can understand the subtitles in emotionally intense language. Although no NLP was used in this study, it shows how transformer models could be used in future systems looking at student feedback or textual responses to identify signs of distress.

Chung and Teo [10] in their systematic literature review grouped the applications of ML models in illnesses like schizophrenia, bipolar disorder, and PTSD. Even though supervised learning is used more in medical fields, their work suggests trying out unsupervised and reinforcement learning as well. It becomes especially vital to use these perspectives as mental health data becomes richer and includes more unlabeled behavior.

The contribution of the DeprMVM model, outlined by researchers in [11], stands out among the others. To address class imbalance, it uses support vector machines, multilayer perceptrons and also SMOTE and clustering methods. Its accuracy of 99.39% and F1-score of 99.51% suggest hybrid learners play a vital role in main diagnostic work. The research's findings agree with the method used, since it associates many types of models and depends on resampling to improve prediction.

Wang et al. [12] studied mental stress in healthcare workers during the COVID-19 pandemic by using a neural network combined with 32 different demographic and psychological factors. This shows that it is vital to take into consideration the unique challenges employees face and design machine learning models that fit their daily jobs. Such an approach could also benefit schools wanting to look after student mental well-being.

Gopalakrishnan et al. [13] also looked at AI-supported mobile health systems utilizing EMA, sensors worn on the body, and chatbot interfaces. They found that now, mental health is being closely observed in real time based on the person and their situation. The study on mental health prediction is working towards implementing active and cellphone-based mental health services for students in the future.

They promote ML for mental health by offering help on both theoretical and practical aspects. They point out changes happening in hybrid modeling, live data reporting, illustrating what the models do and the use of multiple types of data. Like prior studies, this research recommends that any system intended to assist students in mental health needs to be accurate, capable of learning and easy for people to understand.

CHAPTER 3

RESEARCH METHODOLOGY

Mental health disorders among studentsare becoming more common, leading to major health concerns for the community. Traditional diagnostic methods, which mostly depend on people sharing their own experiences, talking in interviews, and taking tests, aren't always able to catch mental health problems early because there can be problems with how honest people are, how someone's view of things might influence the results, or not having enough doctors and other professionals to help. Instead, machine learning allows for monitoring and discovering fine and vital patterns that standard algorithms usually miss. As a result, ML is now a key part of modern systems used for predicting mental health.

By studying what happened in the past, machine learning algorithms discover how being stressed at school, altering daily habits and aspects of a person's demographic influence the chances of depression. Being able to automate the process allows healthcare workers to address less cases and catch issues earlier. Sometimes, these devices operate non-stop, so learners and teachers can observe the outcomes online through their school or college's health platform in real time.

ML is very useful for mental health since it can deal with complex and different types of data. Conditions like sleep quality, one's eating habits, school marks, time spent on screens, socializing, and daily stressors all play a role in shaping someone's mental state. Older statistical procedures usually have difficulties dealing with so many types and modeling relations that are not simple linear ones. A model like Zoo Bird can still learn about these relationships and update as new data is added.

Certainly, ML allows one to predict changes while reacting to them when necessary. Rather than only helping students when they get into serious mental health troubles, schools and universities can check their mental health using ML and intervene early. They use SVMs and ensemble methods to spot slight variations in students' behavior in class and choose the ones who might be at-risk.

Another important benefit is that it lets a company make things specifically for each customer. Machine learning models can be adjusted to look at the things that make different student groups special, so they can give predictions and suggestions that fit those groups better. This is especially helpful in schools or workplaces with people fthat come from a different backgrounds, since what sets off mental health issues can differ depending on where they come from, how much money they have, or their culture.

Thanks to things like deep learning and TabNet, transformers and other improvements in ML, the technology is now more advanced and more capable. They assist in predicting what will happen and understanding the main factors causing a mental disorder. This level of transparency help people trust AI-assisted treatments more and makes it easier to create treatments that fit specific problems.

Furthermore, linking ML with live data sources like phone apps, smartwatches, and what people do online can help keep mental health checks going regularly and make their results more reliable. This makes it possible to change and adjust teaching methods as students' needs change.

In summary, machine learning has made a very huge difference in how we predict mental health problems by making it easier and more accurate to spot them early on and in larger groups. It makes it possible to assess risks on a personal level, read the conclusions, and use the technology in an ideal way in institutions. As ML technologies keep developing, using them in mental health support programs is expected to help students get better results and make things a bit easier for counselors who work with them.

3.1 Framework Overview

The main goal of our framework is to work with a robust machine learning method designed specifically for handling psychologically sensitive information. A student's age, gender, pressure from academics, CGPA, their sleeping habits, their favorite meals and the family's mental health are all pieces of information gathered in the beginning. Data must include various measurements to allow the system to accurately represent the causes of depression in school. After gathering the raw data, it is cleaned and preprocessed to remove any errors or missing information to improve the model's efficiency.

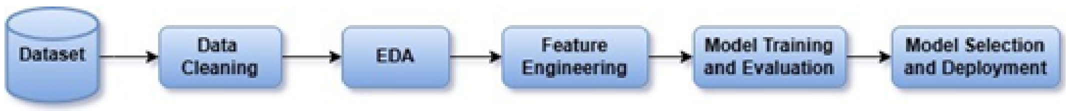

Fig. 3.1. Proposed Framework for identification of mental disorder

The second part of our framework idealy deals with preparing the data to train the used machine learning algorithms. Each categorical feature has been turned into a number using the label encoding, and SMOTE is used to correct the present class imbalance so that both types of cases are evenly present. You can then train random forest data, XGBoost, LightGBM, Deep Neural Networks, and TabNet on the same dataset once it's divided properly. With GridSearchCV, each model gets the right settings so it performs well, and we assess its performance using accuracy, precision, recall, F1-score.

As part of the framework, you need to be able to describe the decision made. By using SHAP and LIME, it is easier to explain the predictions generated by the model to others. They help us identify the main components of a feature, making it easier for professionals to understand the best steps to take. During the last phase, lightweight APIs make sure that students' data is handled correctly and right away once the framework is put into use in school settings.

3.2 Dataset Profile

A total of 27,901 student entries from more than 45 cities in India was used in this study. The data that we have collected using mental health surveys, which looked into things like where students stayed come from, what grades and classes they are in, how they feel and think about things, and their daily lives to help understand their mental health disorders better. A record for each individual gives both objective and subjective signs to use for classifying mental health.

- **Source and Composition:** The first dataset was gathered by surveying university students and working professionals about academic stress. The dataset is comprises of 18 features, among which are academic pressure, satisfaction with work or study, sleep hours, suicidal thoughts, and the way people eat. It also looks at things like a person's gender, age, city where they live, job, what they're studying, and if their family has a history of mental illness, because all of these things might help show if a person is more likely to feel depressed. Each record gets a unique ID field, but that ID does not help with predicting anything.

- **Demographic Attributes:** Out of all the respondents, the youngest was 18 and the eldest 34 years old, with an average age of 26.7. The figures showed that 52.4% were male and 47.6% were female, which means there were equal numbers of male and female gender. Individuals joined us from places like Delhi, Bangalore, Jaipur, Varanasi, Srinagar, and a host of other locations, both metro and non-metro. With many kinds of degrees (like B.Sc, MBBS, M.Tech) and professions (mainly students), the campus becomes quite varied.

  The following are the important lifestyle and psychological features:

  - **Sleep Duration** (e.g., 'Less than 5 hours', '5-6 hours', '7-8 hours')
  - **Dietary Habits** (Healthy, Moderate, Unhealthy)
  - **Work/Study Hours**, ranging from low (<2 hours) to high (>8 hours)
  - **Academic Pressure and Work Pressure**, both rated on a 1–5 Likert scale
  - **Financial Stress**, labeled from 1 (low) to 5 (severe)
  - **Suicidal Thoughts**, a binary feature capturing whether participants reported past suicidal ideation

- **Class Distribution:** The Depression target was given a value of 1 for depressed individuals and 0 for people without depression. Fifty-eight and a half percent of the dataset (16,336 people) are listed as depressed, while only 41.5 percent (11,565 people) are unaffected by depression. Therefore, more preprocessing measures were applied, such as SMOTE oversampling, to make the training data more equal. As depression is present in many survey participants, the data is useful and hard to classify at the same time.

- As well as being helpful for making predictions, this data helps in understanding the different psychosocial stresses that affect students in their education. Since it has both numbers and non-number features, it is well-suited for complex machine learning methods that seek to discover depression patterns connected to someone's academics and life.

	id	Gender	Age	City	Profession	Academic Pressure	Work Pressure	CGPA	Study Satisfaction	Job Satisfaction	Sleep Duration	Dietary Habits	Degree	Have you ever had suicidal thoughts ?	Work/Study Hours	Financial Stress	Family History of Mental Illness	Depression
0	2	Male	33.0	Visakhapatnam	Student	5.0	0.0	8.97	2.0	0.0	'5-6 hours'	Healthy	B.Pharm	Yes	3.0	1.0	No	1
1	8	Female	24.0	Bangalore	Student	2.0	0.0	5.90	5.0	0.0	'5-6 hours'	Moderate	BSc	No	3.0	2.0	Yes	0
2	26	Male	31.0	Srinagar	Student	3.0	0.0	7.03	5.0	0.0	'Less than 5 hours'	Healthy	BA	No	9.0	1.0	Yes	0
3	30	Female	28.0	Varanasi	Student	3.0	0.0	5.59	2.0	0.0	'7-8 hours'	Moderate	BCA	Yes	4.0	5.0	Yes	1
4	32	Female	25.0	Jaipur	Student	4.0	0.0	8.13	3.0	0.0	'5-6 hours'	Moderate	M.Tech	Yes	1.0	1.0	No	0

	id	Gender	Age	City	Profession	Academic Pressure	Work Pressure	CGPA	Study Satisfaction	Job Satisfaction	Sleep Duration	Dietary Habits	Degree	Have you ever had suicidal thoughts ?	Work/Study Hours	Financial Stress	Family History of Mental Illness	Depression
27896	140685	Female	27.0	Surat	Student	5.0	0.0	5.75	5.0	0.0	'5-6 hours'	Unhealthy	'Class 12'	Yes	7.0	1.0	Yes	0
27897	140686	Male	27.0	Ludhiana	Student	2.0	0.0	9.40	3.0	0.0	'Less than 5 hours'	Healthy	MSc	No	0.0	3.0	Yes	0
27898	140689	Male	31.0	Faridabad	Student	3.0	0.0	6.61	4.0	0.0	'5-6 hours'	Unhealthy	MD	No	12.0	2.0	No	0
27899	140690	Female	18.0	Ludhiana	Student	5.0	0.0	6.88	2.0	0.0	'Less than 5 hours'	Healthy	'Class 12'	Yes	10.0	5.0	No	1
27900	140699	Male	27.0	Patna	Student	4.0	0.0	9.24	1.0	0.0	'Less than 5 hours'	Healthy	BCA	Yes	2.0	3.0	Yes	1

Fig 3.2: Dataset

3.3 Exploratory Data Analysis (EDA)

The goal of EDA was to analyze how all the important variables from the dataset are related and of what structure they are. While completing this part, I investigated the data, using graphs and stats, to discover anything that seemed unusual. If I looked at the plots and graphs, I saw how many students were male or female, their anxiety levels at school and the number of those who were depressed, while box plots pointed out anything irregular about CGPA and hours spent studying each day. Using the correlation heatmap, we could see that depression is associated with stress related to academics, not sleeping enough and having thoughts of suicide. Using EDA, researchers found that the data includes class imbalance which influenced the way the data was readied for analysis and modeling.

1. **Pie Plots (for Academic_Pressure, Study_Satisfaction, WrkStdy_Hours, Financial_Stress)**

   Why plot it: They can help us see how much or how little each category makes up in a single type of category. They let us notice the distribution of students' responses for those two aspects related to their studies and finances.

What it signifies: The size of each section indicates the number of students dealing with that form of pressure, with bigger sections indicating more students. It makes it easy to spot which categories are seen the most and least often.

How it's done:

- The code goes through each column name from the given list (pie_cols).
- The value counts help it find the number of times every unique value in the column occurs.
- It draws the pie chart by using matplotlib. It does this using pyplot.pie.
- The labels are created by using the categories on the value counts.
- autopct='%0.2f%%' makes it easier to see how much of the whole is included in each section of the pie chart.
- The 'explode' tool is used to slightly move the slices apart, so they are easier to see.

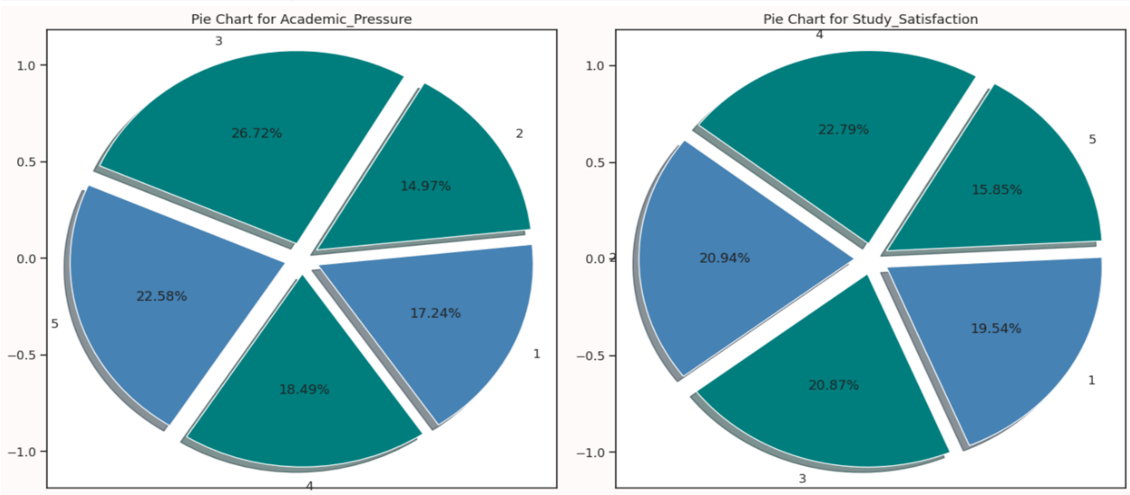

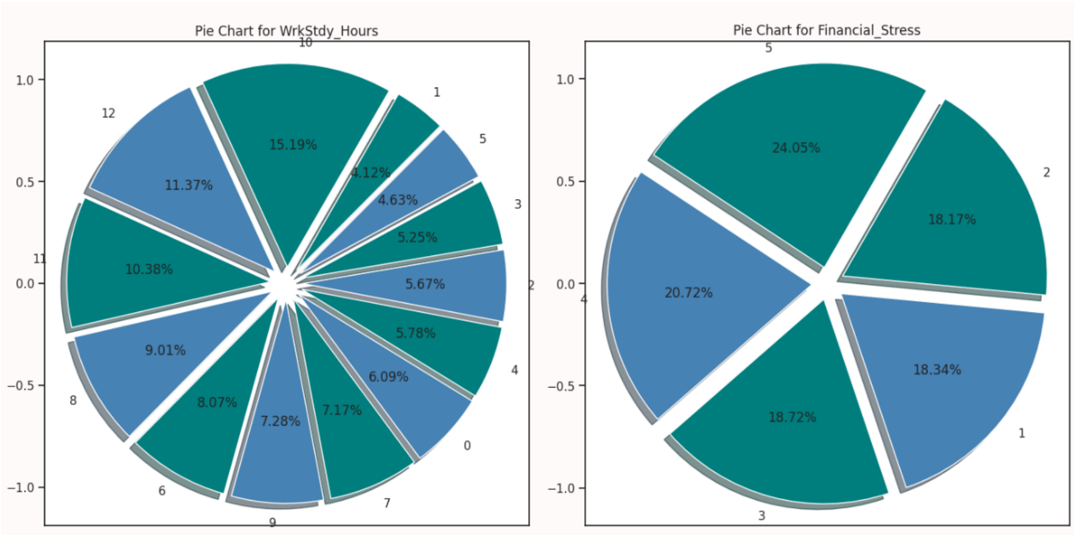

Fig 3.3: Pie plots

2. **Bar Plots (for Gender, Profession, Sleep_Duration, Dietary_Habits, Suicidal_Thoughts, Family_Mental_History)**

   Why plot it: A bar plot is best used when you want to display the amount of times each category in our dataset appears in a categorical variable. They help you swiftly compare the amount of each event within different groups.

   What it signifies: The vertical height of the bar graph signifies the number of observations in a given category. It allows us to understand the range of ages, living habits, and history of mental health in the people included in the data.

   How it's done:

   - The code goes through each name stored in the bar_cols list.
   - Each column gets its own bar plot, made by using seaborn.countplot.
   - data=data specifies the DataFrame.
   - The counts are grouped by the column named by x=i, which appears on the x-axis.
   - Customization of the line drawing's look is achieved with the help of palette, dodge, width, edgecolor, linewidth, and saturation.

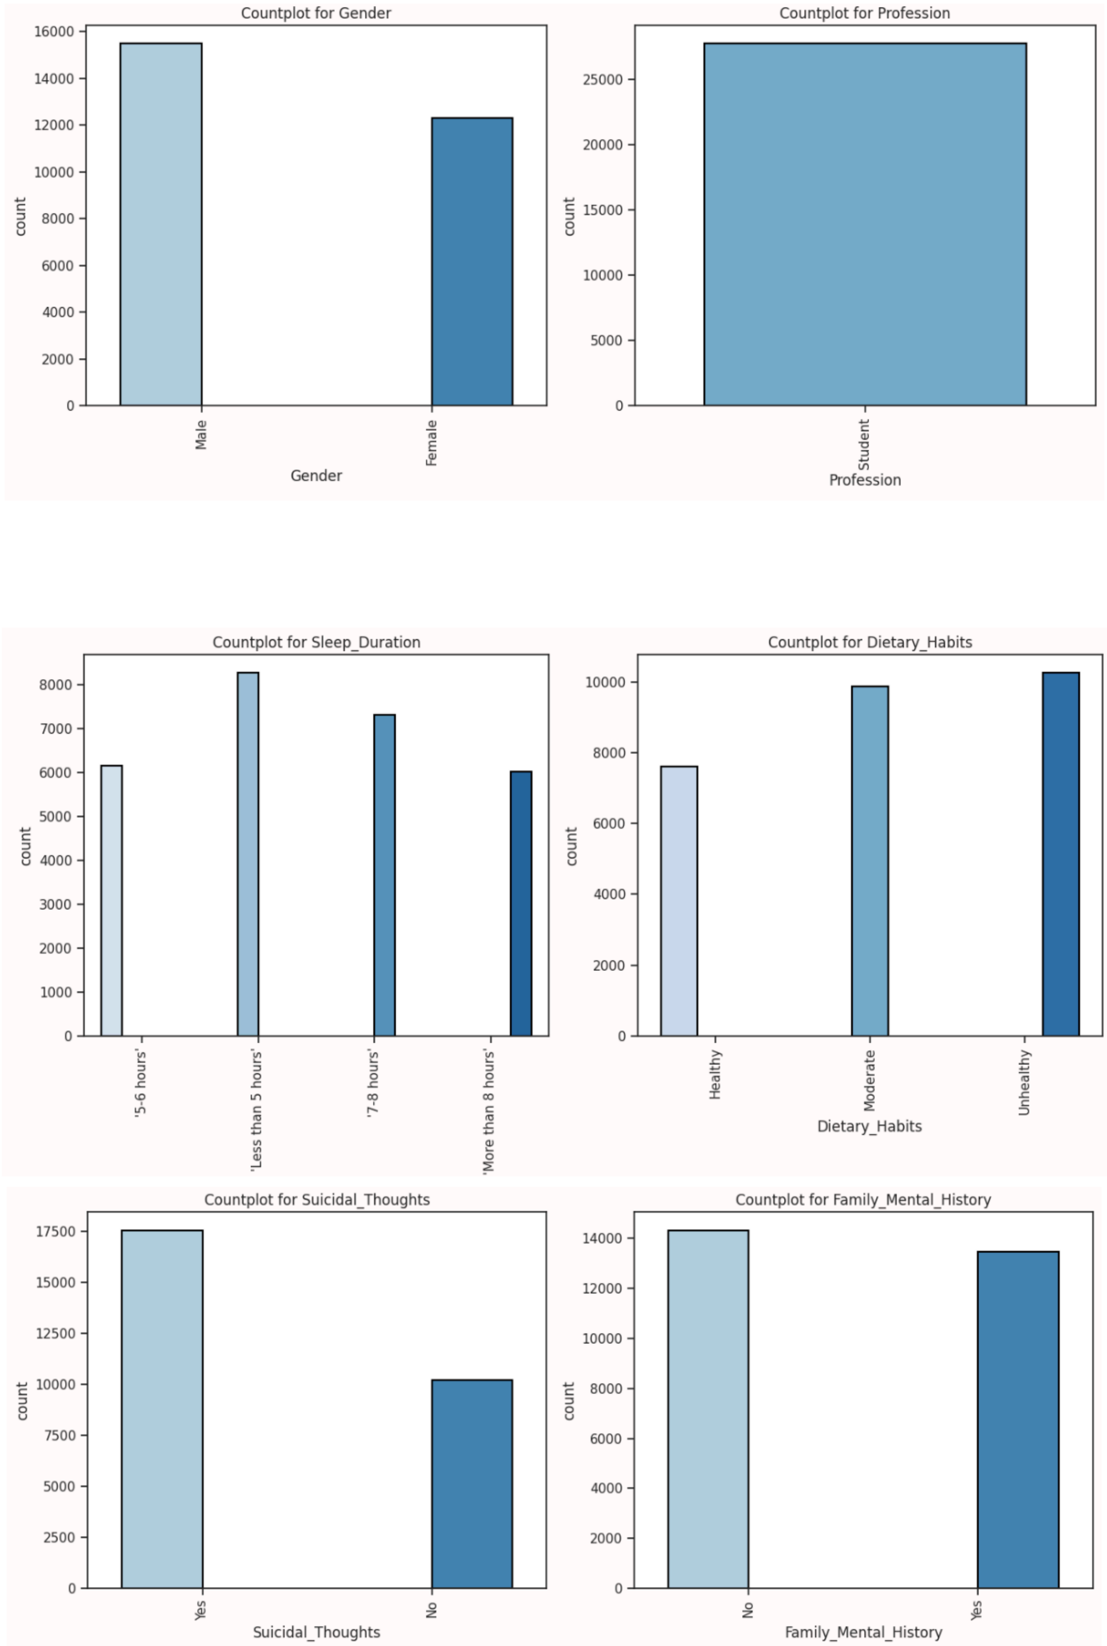

Fig 3.4: Bar Plots

3. **Histogram Plot (for numerical columns)**

Why plot it: A histogram helps to show the distribution of a variable's values. They highlight the number of points falling in each range (or interval) of the main variable.

What it signifies: The form of the histogram demonstrates whether the data is normally distributed, or if it is skewed or multi modal. It shows how most of the data are grouped, how spread out the data are, and if there are any numbers that are way different from the others.

How it's done:

- It pulls numerical columns by using the instruction data.select_dtypes(include=np.number).columns.
- It iterates through these columns.
- Seaborn.histplot is used for each column.
- bins=30 specifies the number of bins.
- Adding kde=True draws a Kernel Density smooth line, giving clearer insights into the shape of the distribution.
- Color, edgecolor, and alpha are used to determine the look of a Canvas.

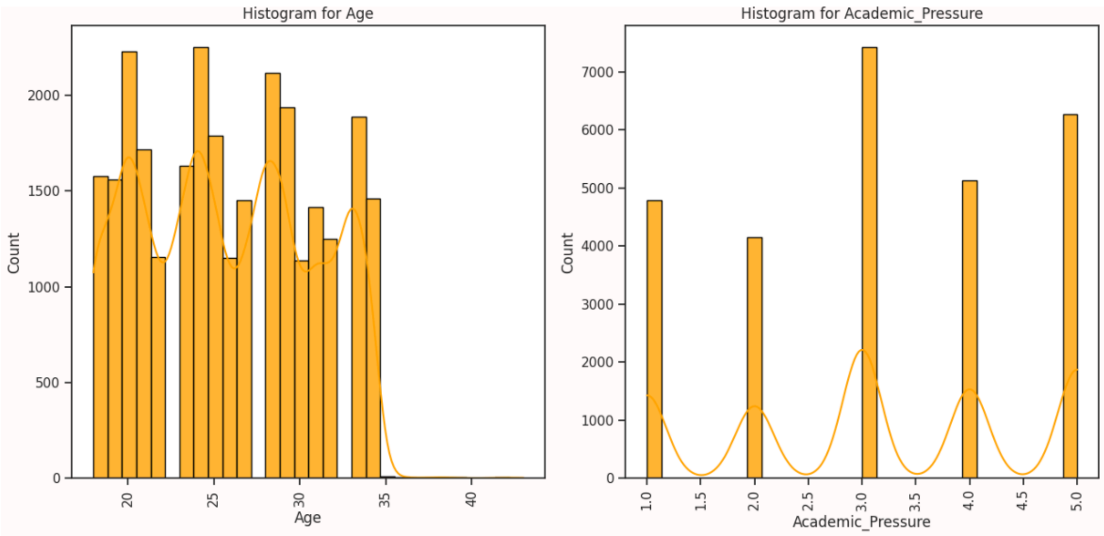

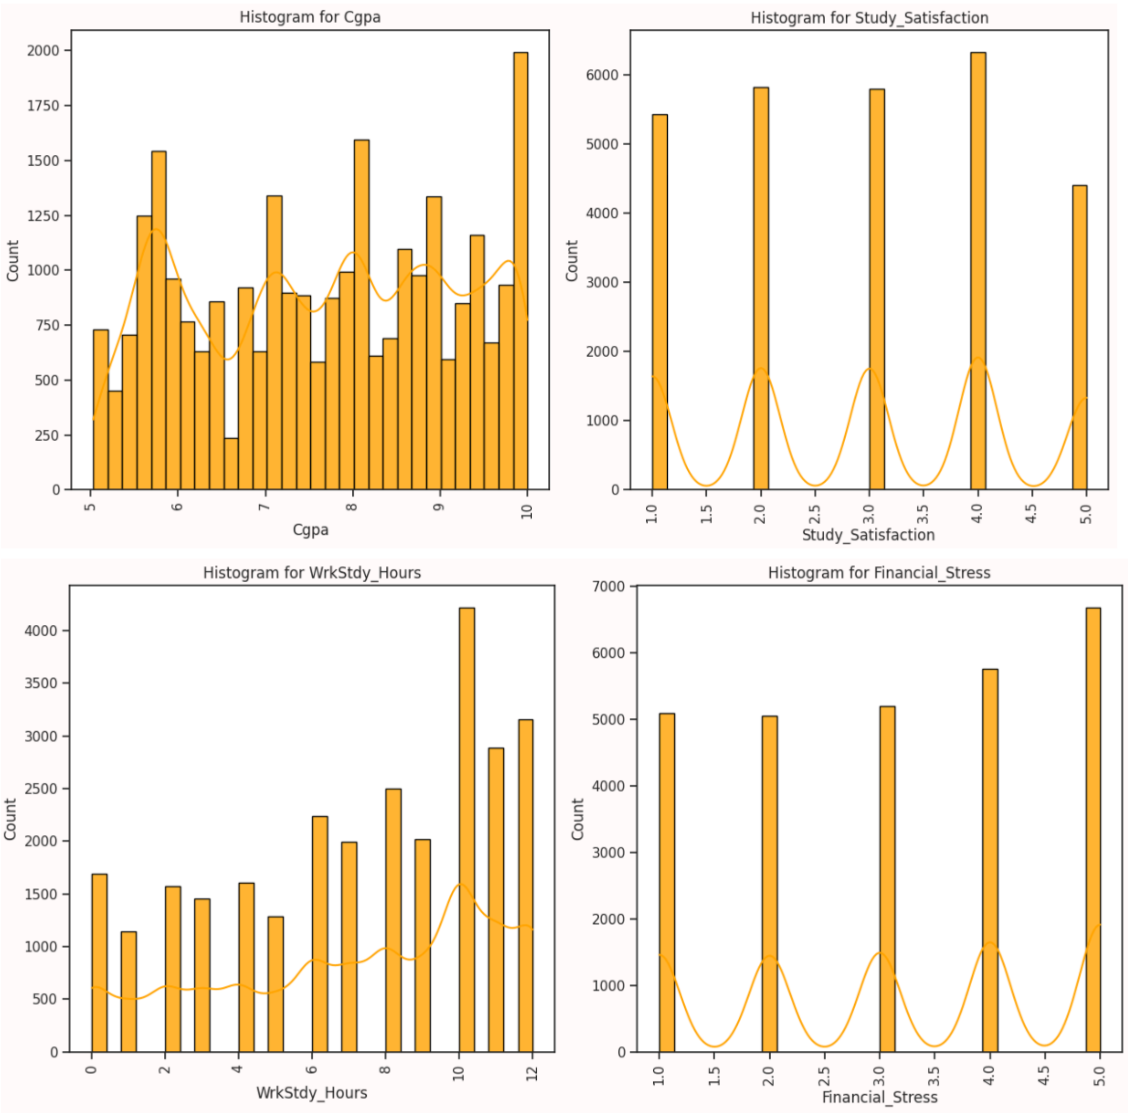

Fig 3.5: Histogram Plots

4. **Target Variable Stats (Descriptive Statistics)**

Why plot it: Though EDA does not consider visuals, learning about the description of the target variable ('Depression') is important. It gives the information about the distribution in the form of numbers.

What it signifies: This output provides a quick way to see the center and spread of the target variable. If Depression is treated as a 0 or 1 variable, the mean will tell you the rate of people with depression out of all the observations.

How it's done:

- It uses the command Depression.describe().round(3) to determine the statistics of the data.

- The to_frame() method changes the code output into a DataFrame to help with presentation.

	Depression
count	27765.000
mean	0.585
std	0.493
min	0.000
25%	0.000
50%	1.000
75%	1.000
max	1.000

Fig 3.6: Target Variable Stats

5. **Bar Plots (various combinations of Target vs. Gender, Sleep_Duration, Age, Suicidal_Thoughts, Dietary_Habits, Degree, Profession, City, Academic_Pressure, Family_Mental_History, Study_Satisfaction, WrkStdy_Hours)**

Why plot it: These visuals are designed to study the connection between the variable 'Depression' and other categorical variables. They illustrate how much depression people tend to experience depending on their categories.

What it signifies: Supposing depression was a 1 for those with it and 0 otherwise, we can look at the ratios of depression in men and women as well as at those under different sleep amounts to see which factors are tied to higher depression levels.

How it's done:

- The data is grouped first by relevant categories, then the average values of the 'Depression' column are found for every group.
- After gathering the data, it uses barplot from seaborn to display it.
- To create side-by-side bars, the columns for the x-axis, y-axis, and grouping are defined using the x, y, and hue parameters.

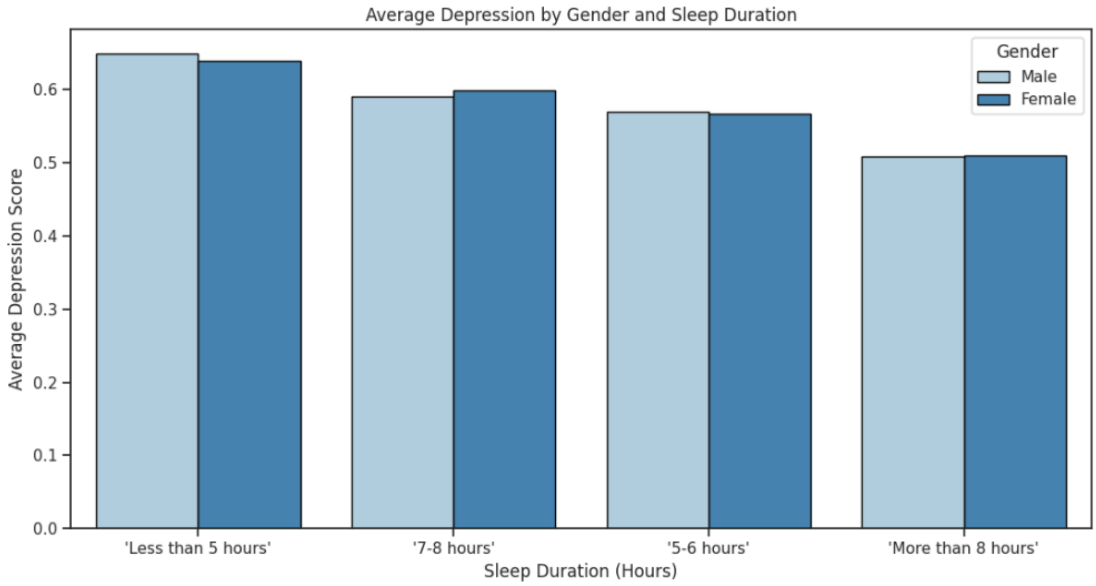
Average Depression by Gender and Sleep Duration
Gender
Male
Female
Average Depression Score
0.0
0.1
0.2
0.3
0.4
0.5
0.6
'Less than 5 hours'
'7-8 hours'
'5-6 hours'
'More than 8 hours'
Sleep Duration (Hours)

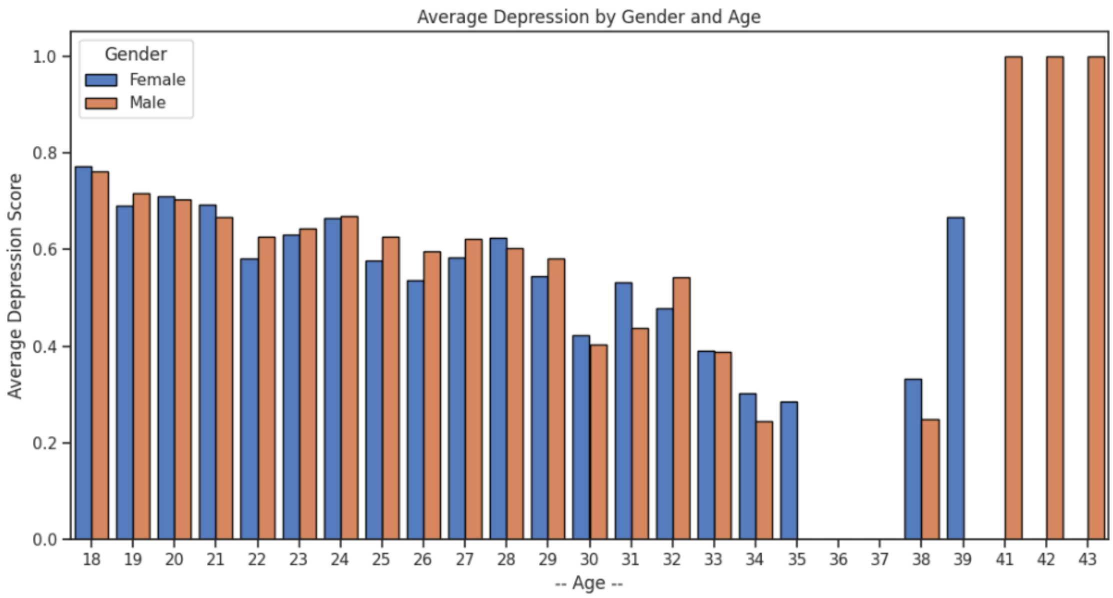
Average Depression by Gender and Age
Gender
Female
Male
Average Depression Score
0.0
0.2
0.4
0.6
0.8
1.0
18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 41 42 43
-- Age --

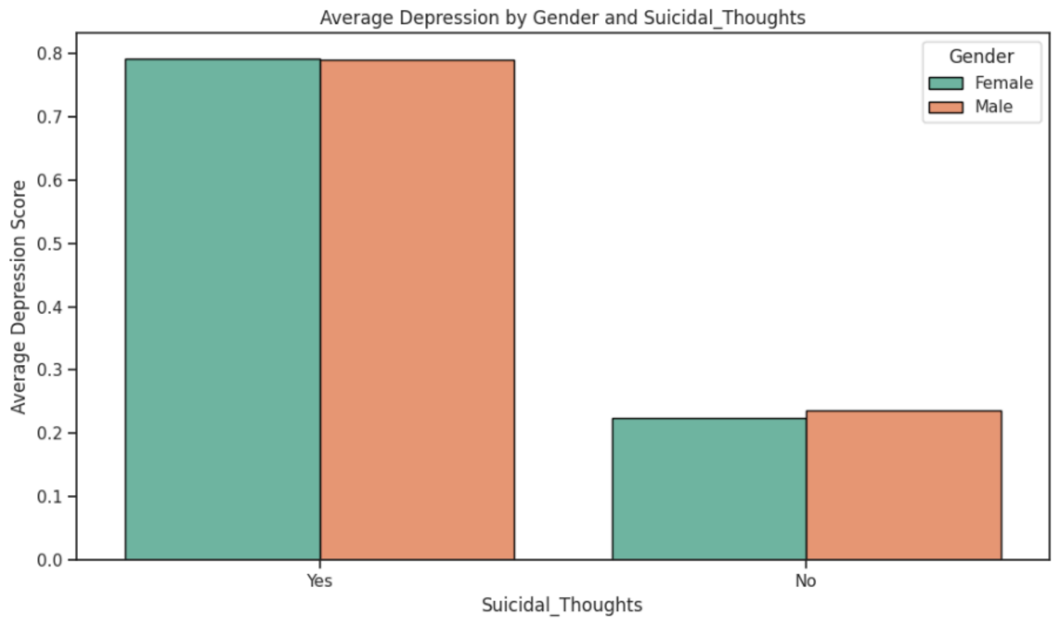
Average Depression by Gender and Suicidal_Thoughts
Gender
Female
Male
Average Depression Score
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
Yes
No
Suicidal_Thoughts

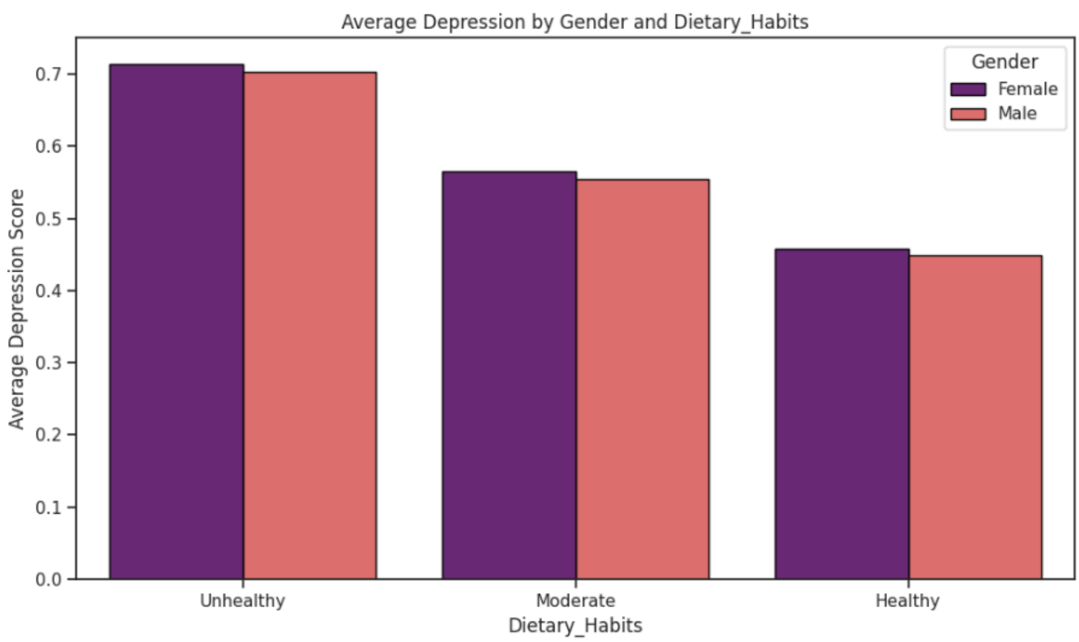
Average Depression by Gender and Dietary_Habits
Gender
Female
Male
Average Depression Score
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
Unhealthy
Moderate
Healthy
Dietary_Habits

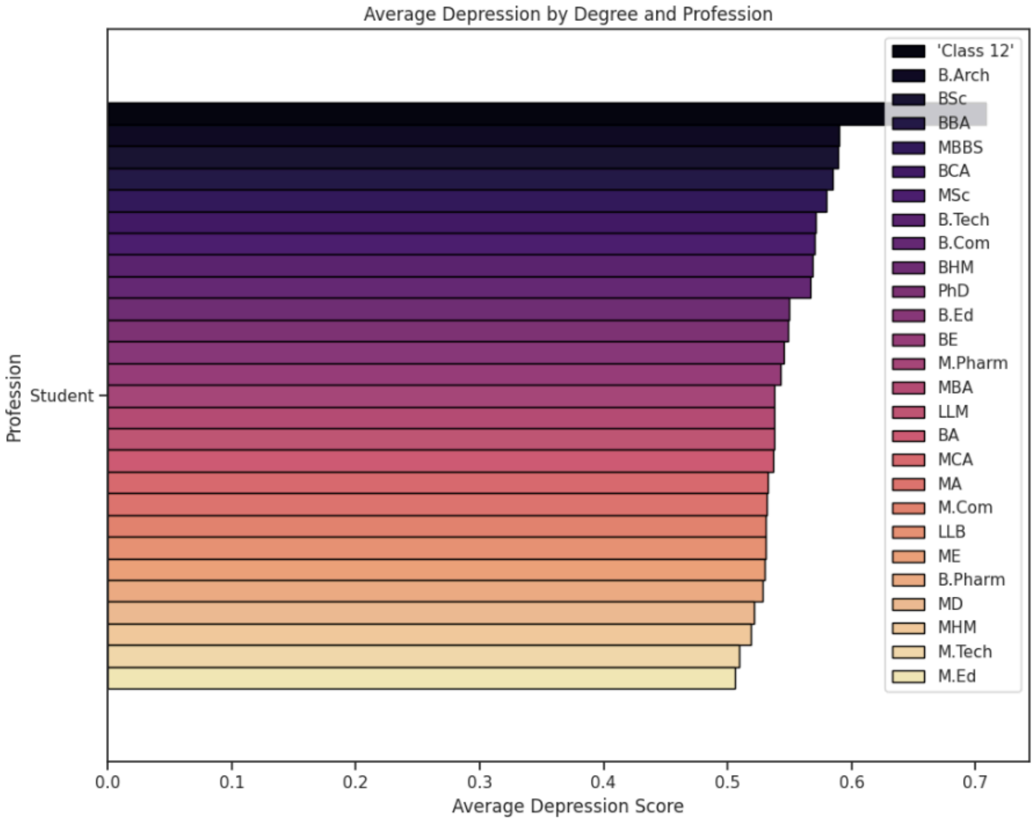
Average Depression by Degree and Profession
'Class 12'
B.Arch
BSc
BBA
MBBS
BCA
MSc
B.Tech
B.Com
BHM
PhD
B.Ed
BE
M.Pharm
MBA
LLM
BA
MCA
MA
M.Com
LLB
ME
B.Pharm
MD
MHM
M.Tech
M.Ed
Profession
Student
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
Average Depression Score


Average Depression by Gender and City
Gender
Male
Female
Average Depression Score
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
Ahmedabad
Hyderabad
Patna
Bhopal
Indore
Faridabad
Rajkot
Delhi
Chennai
Kolkata
Bangalore
Srinagar
Thane
Meerut
Kalyan
Ludhiana
Visakhapatnam
Surat
Varanasi
Nagpur
Ghaziabad
Vasai-Virar
Vadodara
Pune
Mumbai
Lucknow
Kanpur
Nashik
Agra
Jaipur
City

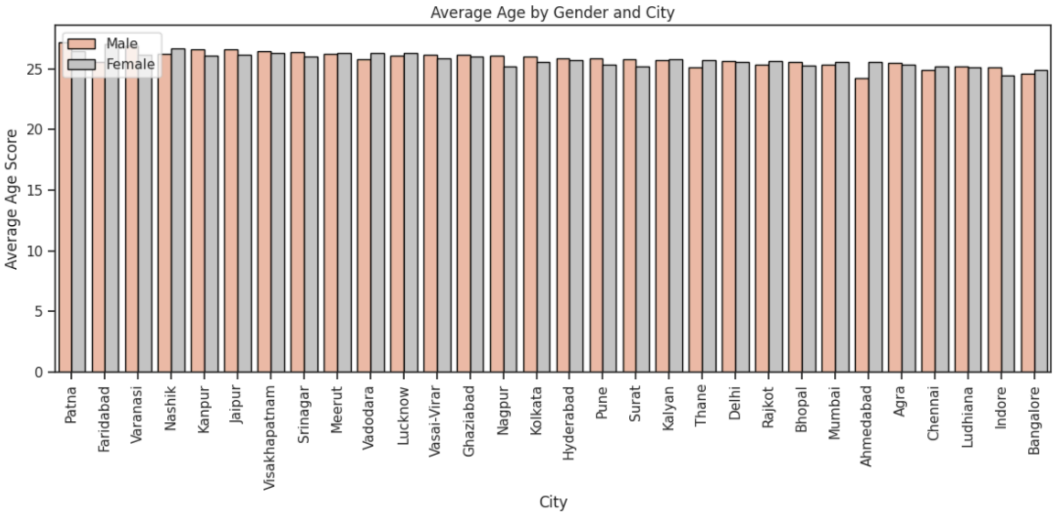
Average Age by Gender and City
Male
Female
Average Age Score
0
5
10
15
20
25
Patna
Faridabad
Varanasi
Nashik
Kanpur
Jaipur
Visakhapatnam
Srinagar
Meerut
Vadodara
Lucknow
Vasai-Virar
Ghaziabad
Nagpur
Kolkata
Hyderabad
Pune
Surat
Kalyan
Thane
Delhi
Rajkot
Bhopal
Mumbai
Ahmedabad
Agra
Chennai
Ludhiana
Indore
Bangalore
City

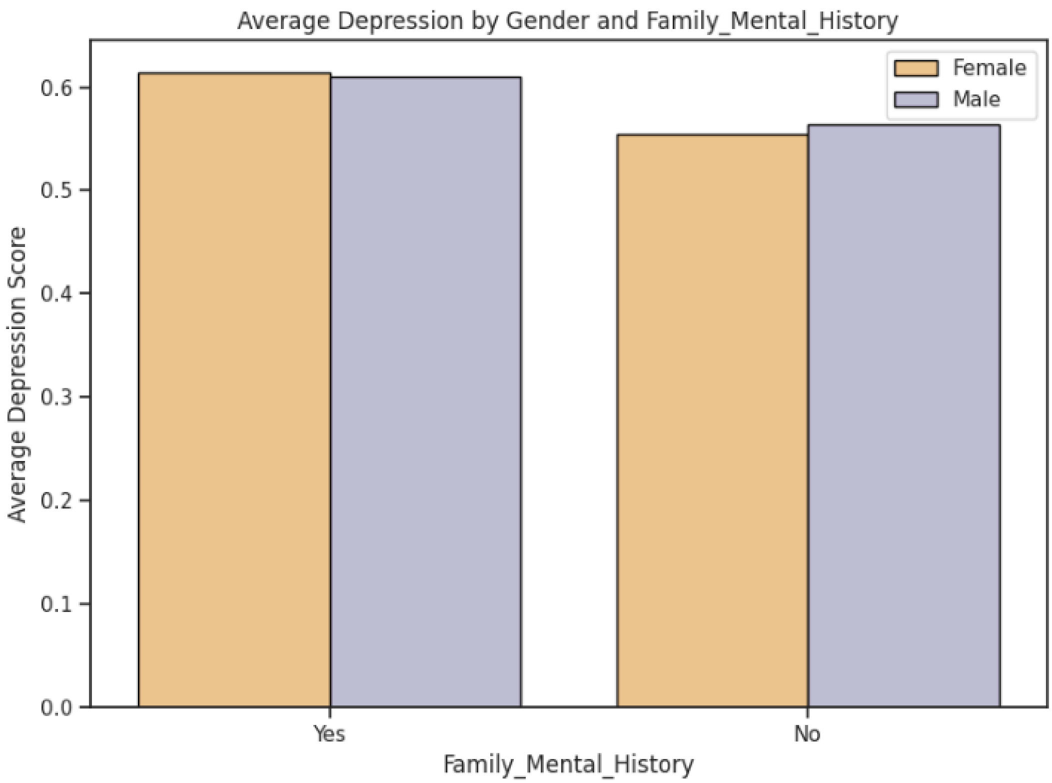
Average Depression by Gender and Family_Mental_History
Female
Male
Average Depression Score
0.0
0.1
0.2
0.3
0.4
0.5
0.6
Yes
No
Family_Mental_History

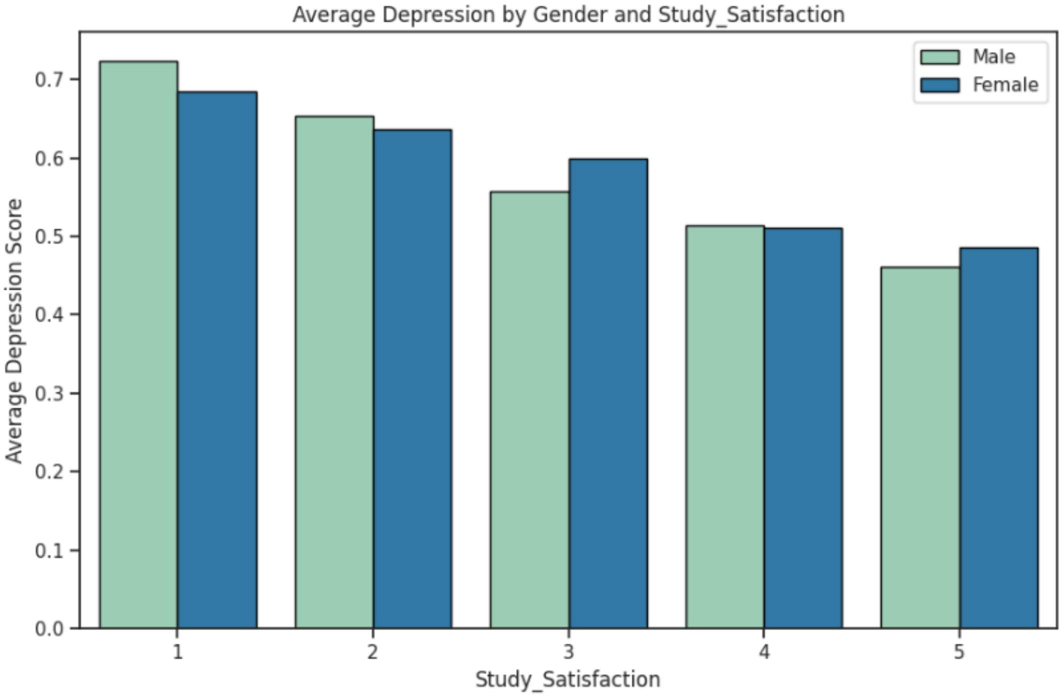
Average Depression by Gender and Study_Satisfaction
Male
Female
Average Depression Score
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
1
2
3
4
5
Study_Satisfaction

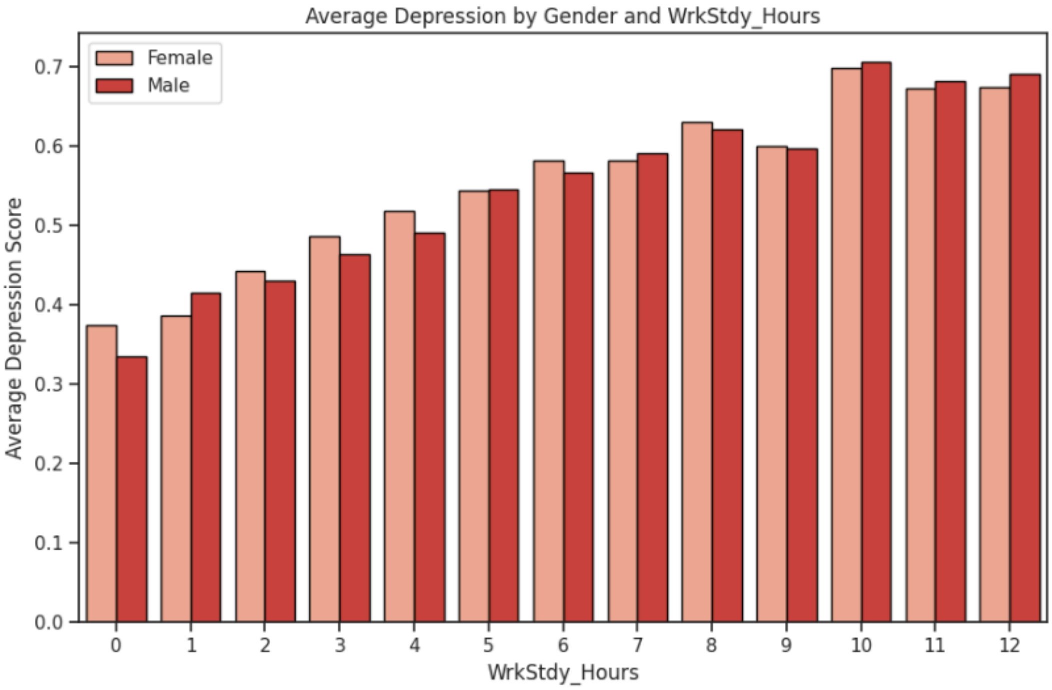
Average Depression by Gender and WrkStdy_Hours
Female
Male
Average Depression Score
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0
1
2
3
4
5
6
7
8
9
10
11
12
WrkStdy_Hours

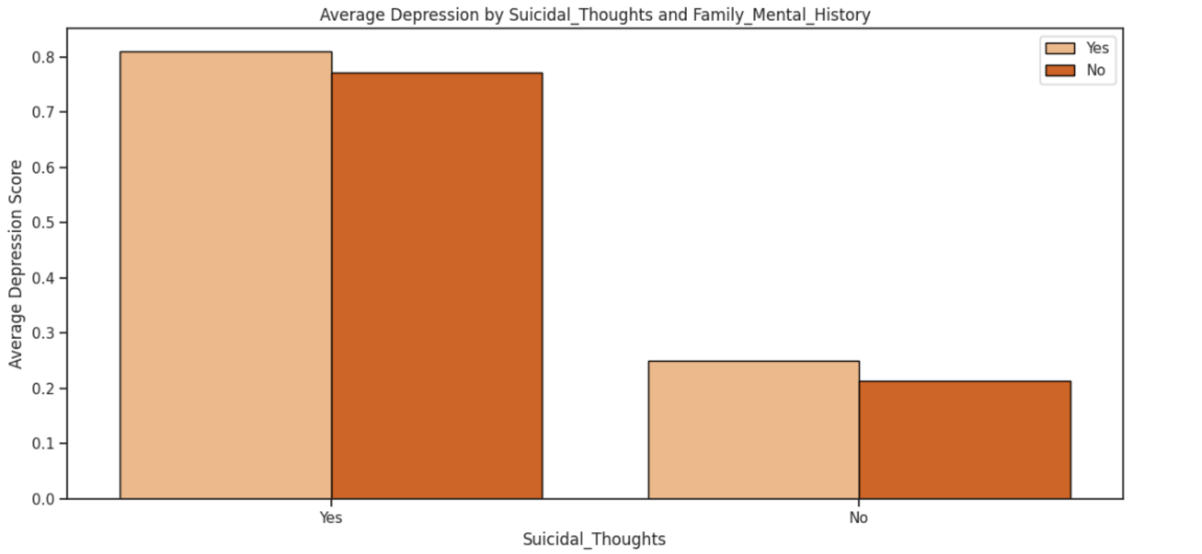

Fig 3.7: Bar Plots

6. **Scatter Plots (various combinations of Cgpa vs. Age & Academic_Pressure, Target vs. Age & Academic_Pressure, Target vs. Age & Financial_Stress, Cgpa vs. Target & Academic_Pressure, Sleep_Duration vs. Target & Cgpa)**

   Why plot it: Using scatter plots makes it easy to see the relationship between two numbers being compared. Adding the hue parameter lets you see the effects of a third variable on the relationship between these two variables.

   What it signifies:With a scatter plot, one can see the level of correlation between two variables as well as any possible patterns or trends in the data.The use of hue shows a difference in the connection between the two main variables when you look at each of the third variable's categories or values. For instance, it reveals if there is a different link between Academic Pressure and Depression for people of different ages.

   How it's done:

   - It breaks down data into groups and then finds out the average value of a certain variable (such as 'Age' or 'Depression') for different groups.
   - The plot is made using seaborn.scatterplot.
   - The points are organized on the axes x and y, and hue sets the column for adding color to the points.

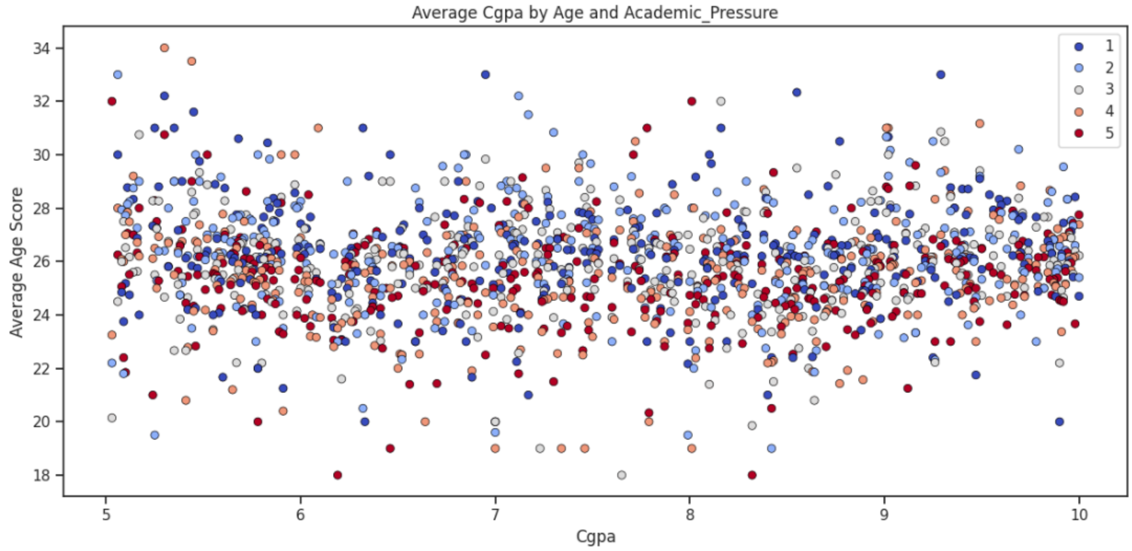
Average Cgpa by Age and Academic_Pressure
Average Age Score
Cgpa
1
2
3
4
5

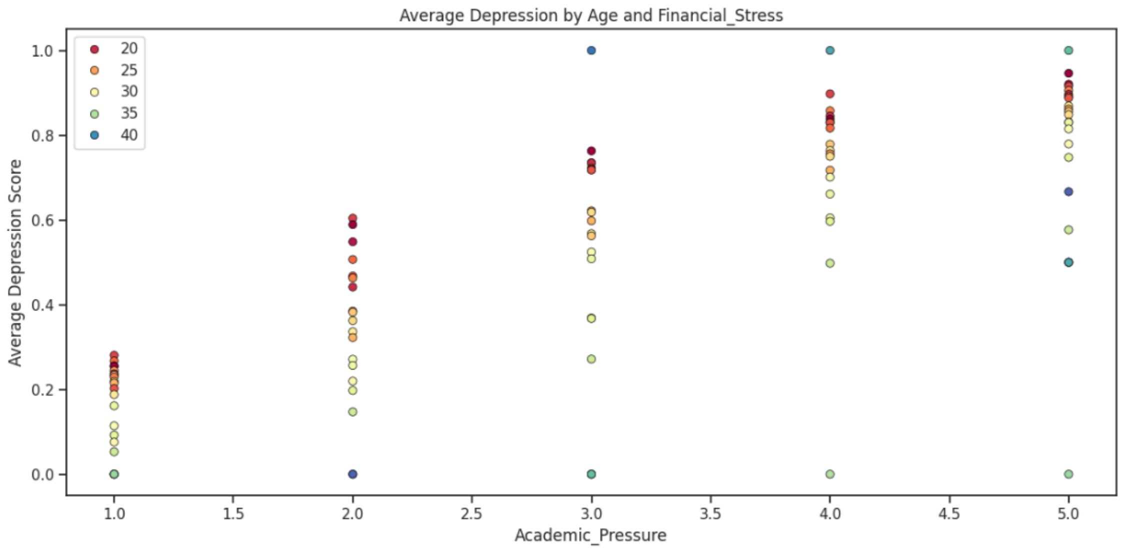
Average Depression by Age and Financial_Stress
Average Depression Score
Academic_Pressure
20
25
30
35
40

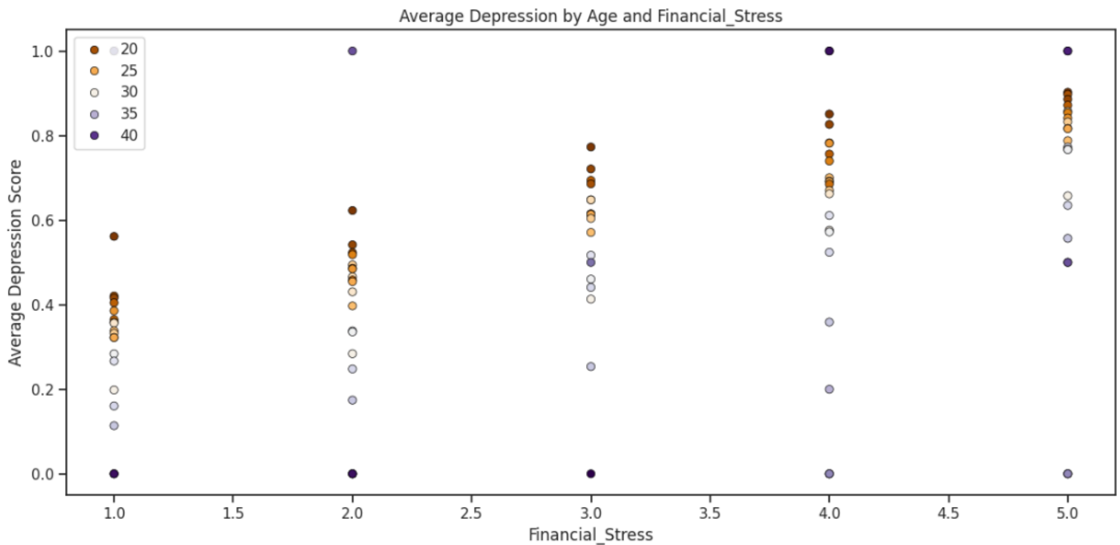
Average Depression by Age and Financial_Stress
Average Depression Score
Financial_Stress
20
25
30
35
40

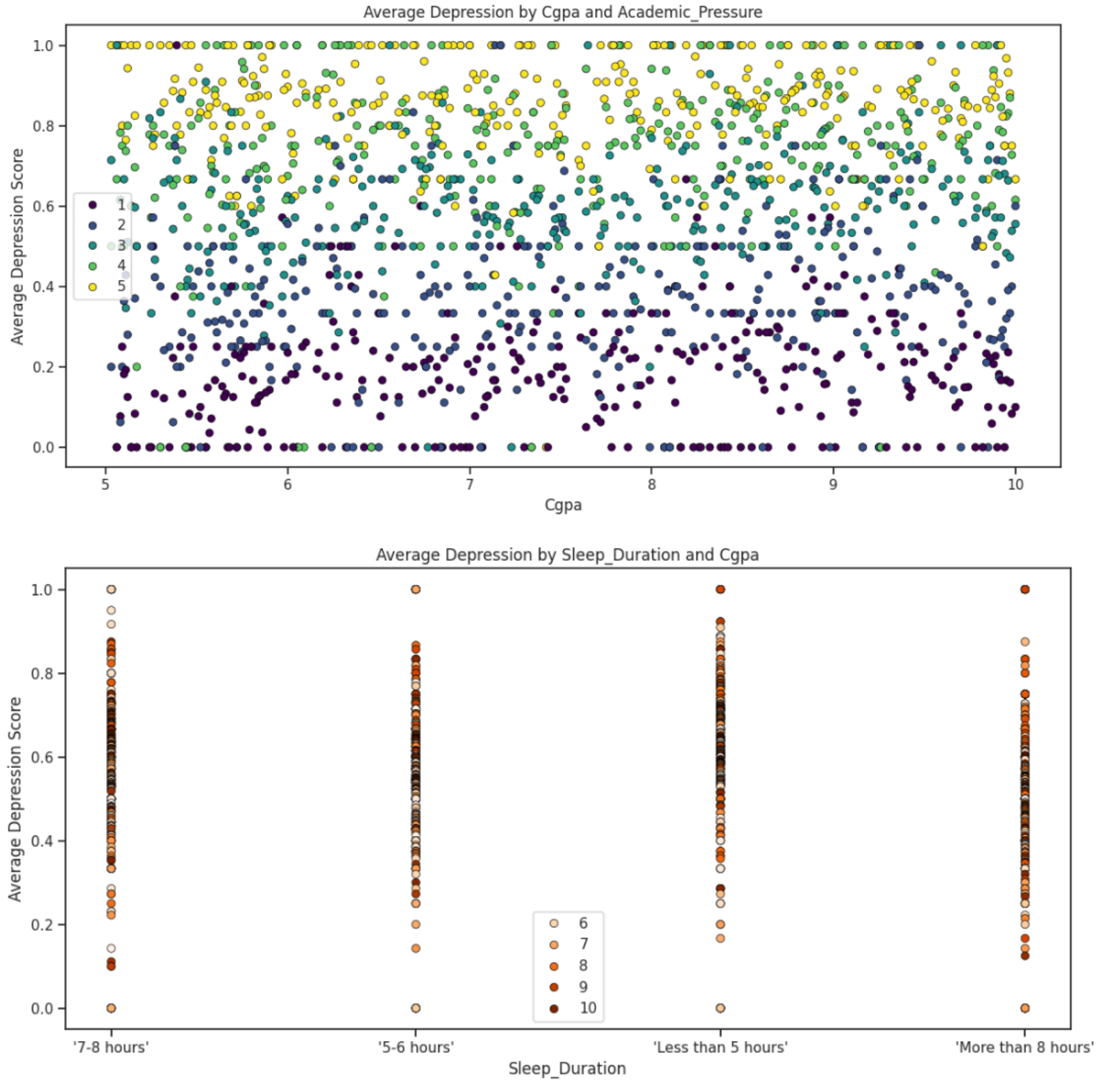

Fig 3.8: Scatter Plots

7. **Correlation Matrix (Heatmap)**

   Why plot it: It displays each numerical variable's correlation with every other numerical variable in a table.

   What it signifies: The color in the heatmap helps you see how strong a correlation is and in which direction it is going. It finds out which variables can predict the target variable best. If the scatterplot shows a strong link between variables, then there may be a multicollinearity issue to watch out for while modeling.

   How it's done:

   - The code selects numerical columns.
   - The function .corr() is used to find the correlation matrix.

- ○ The seaborn.heatmap function is used to generate the visualization.
- ○ Using annot=True will graph the correlation values on the heatmap.
- ○ cmap='RdYlBu' sets the color scheme.

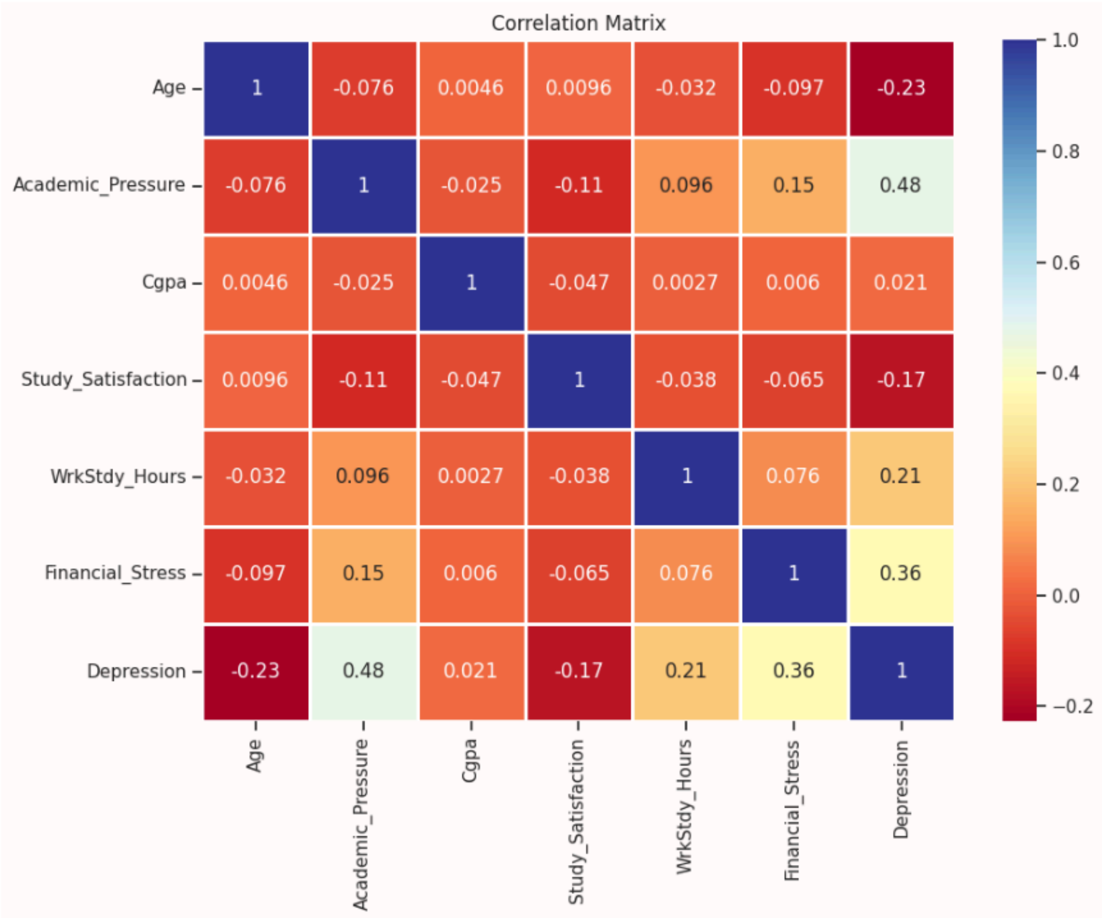

Fig 3.9 : Heatmap

8. **Pair Plots (Scatter and Regression types)**

Why plot it: Pair plots are also called scatterplot matrices and they place scatter plots side by side for every pair of numerical variables in the dataset. The diagonal describes the way each variable is spread out (using KDE plots).

What it signifies:

- ○ The lines in the off-diagonal cells represent the association between pairs of variables.
- ○ The diagonal plots highlight the distribution of all the variables.

- Analyzing the entire grid allows you to easily find related numbers, groupings, and how features are arranged.
- By setting hue='Depression', you can compare the ways the relationships and distributions differ in each category of the target class.
- The regression line is added on every scatter plot in the kind='reg' version, highlighting the line of the trend.

How it's done:

- The code selects numerical columns.
- It uses seaborn.pairplot.
- The height of each plot depends on the overall number of plots that will fit inside.
- plot_kws customizes the scatter points.
- kind is chosen as either 'scatter' or 'reg'.
- Changing diag_kind to ‘kde’ makes a kernel density plot on the main diagonal.
- The plot colors points according to the outcome of the target variable.

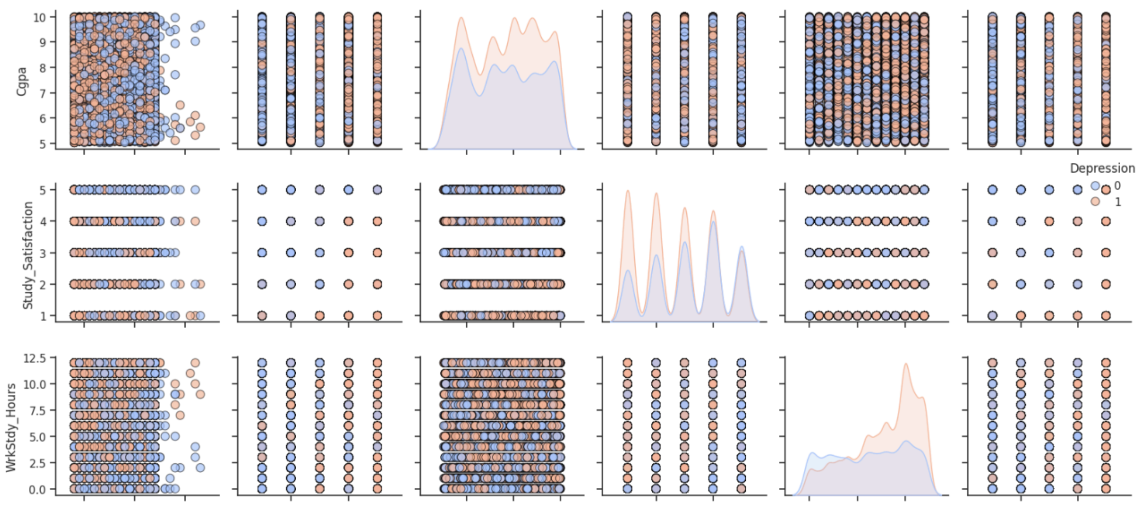

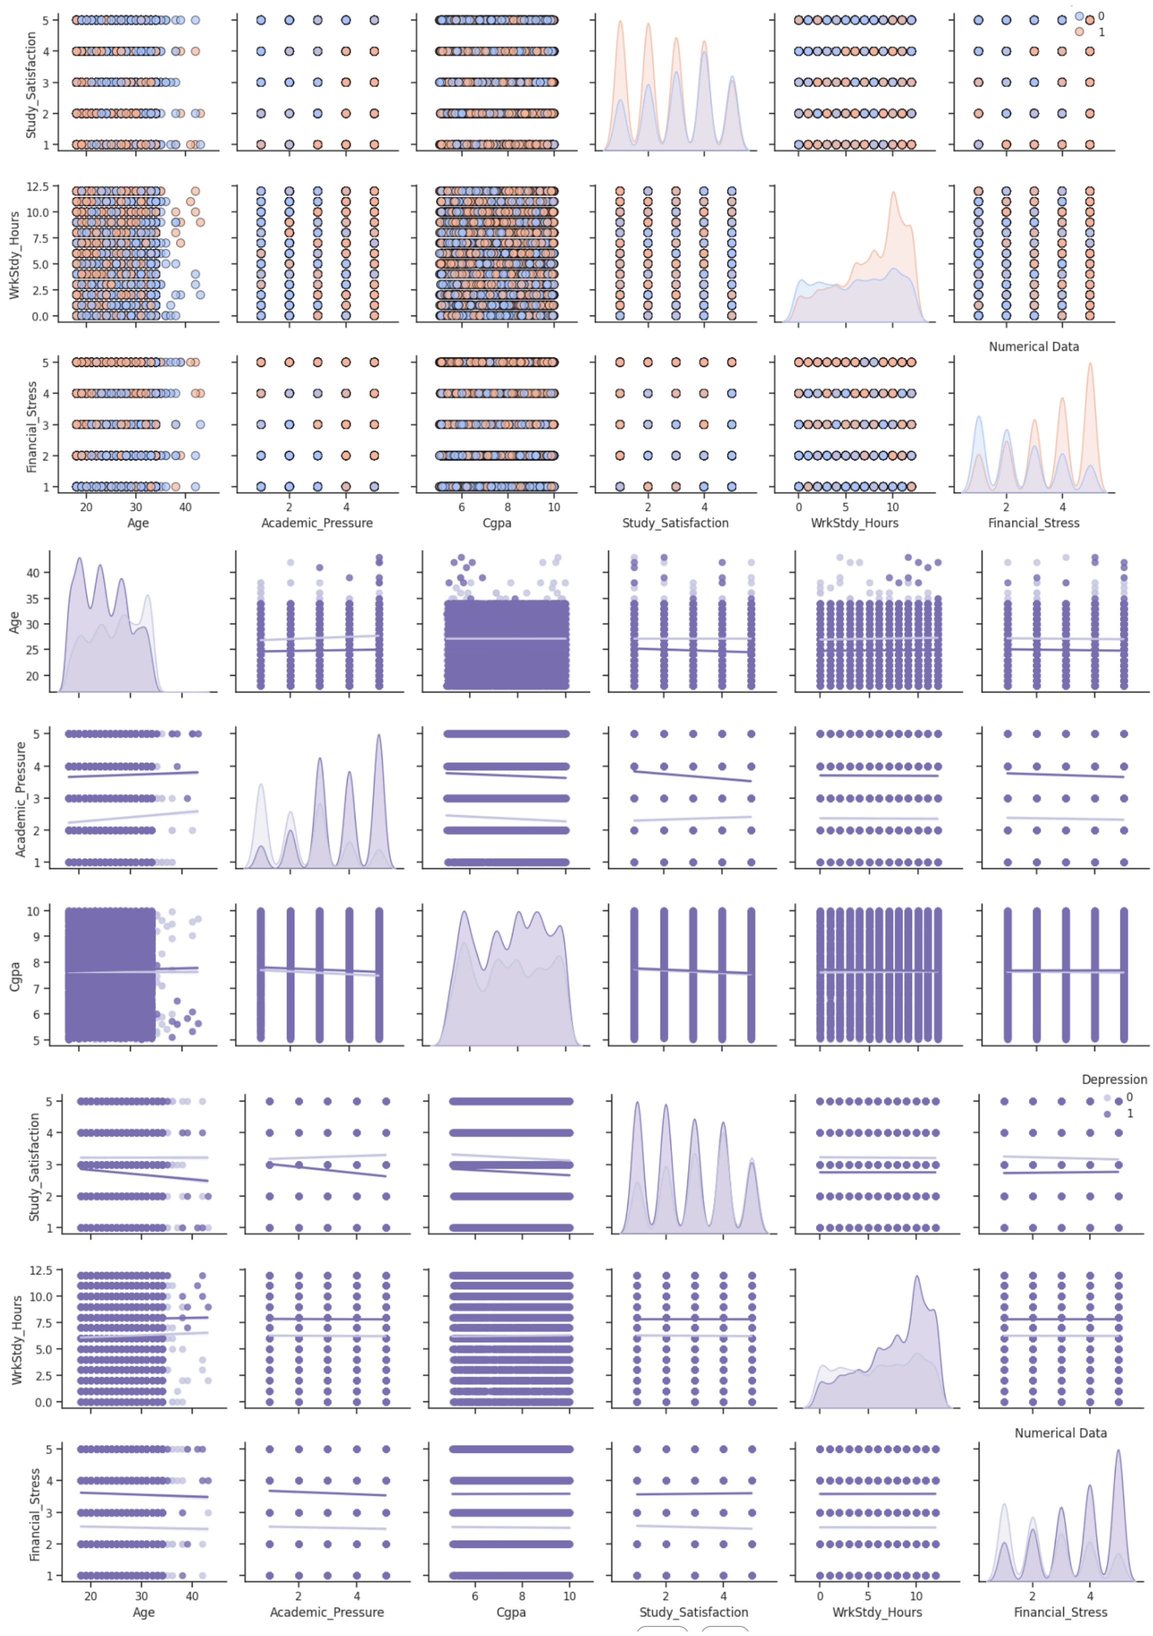

Fig 3.10: Pair Plots

9. ANOVA Test on Categorical Columns

Why plot it: ANOVA (Analysis of Variance) is a statistical test that has been used for a very long time to know if there are any statistically significant differences between the means of three or more independent (unrelated) groups. In context to our thesis, it's used to see if the mean 'Depression' score and other scores is significantly different across the different categories of each categorical variable of our dataset.

What it signifies: A significantly low p-value (typically less than 0.05) from the ANOVA test suggests that there is a significant difference in the mean of the depression level among at least some of the categories of the categorical variable. This indicates that the categorical variable in our dataset is highly likely related to depression.

How it's done:

- The code iterates through categorical columns.
- For each column, it groups the 'Depression' values by the categories of that column.
- It then performs the scipy.stats.f_oneway test on these groups.
- It prints the F-statistic and the p-value.

```
ANOVA Test for Gender - F-statistic: 0.078, p-value: 0.780
ANOVA Test for City - F-statistic: 5.682, p-value: 0.000
ANOVA Test for Profession skipped due to insufficient data.
ANOVA Test for Sleep_Duration - F-statistic: 93.061, p-value: 0.000
ANOVA Test for Dietary_Habits - F-statistic: 629.637, p-value: 0.000
ANOVA Test for Degree - F-statistic: 20.934, p-value: 0.000
ANOVA Test for Suicidal_Thoughts - F-statistic: 11855.655, p-value: 0.000
ANOVA Test for Family_Mental_History - F-statistic: 78.583, p-value: 0.000
```

Fig 3.11: ANOVA Test

10. T-test on Numerical Columns

Why plot it: A t-test (specifically an independent samples t-test in this case) is widely used to compare the means of two independent groups to determine if there is a statistically significant difference between them. Here, it's used to compare the mean of each numerical feature for the two 'Depression' groups (0 and 1).

What it signifies: A low p-value (typically less than 0.05) from the t-test indicates a significant difference in the mean value of the numerical feature between the depressed (1) and not depressed (0) groups. This suggests that the numerical feature is likely related to depression.

How it's done:

- ○ The code iterates through numerical columns.
- ○ For each column, it separates the values into two groups based on the 'Depression' value (0 or 1).
- ○ It performs the scipy.stats.ttest_ind test on these two groups.
- ○ It prints the t-statistic and the p-value, and a message indicating if the difference is significant.

```
T-test for Age - t-statistic: 38.741, p-value: 0.000
Significant difference in Age between Depression 0 and 1

T-test for Academic_Pressure - t-statistic: -89.983, p-value: 0.000
Significant difference in Academic_Pressure between Depression 0 and 1

T-test for Cgpa - t-statistic: -3.472, p-value: 0.001
Significant difference in Cgpa between Depression 0 and 1

T-test for Study_Satisfaction - t-statistic: 28.411, p-value: 0.000
Significant difference in Study_Satisfaction between Depression 0 and 1

T-test for WrkStdy_Hours - t-statistic: -35.638, p-value: 0.000
Significant difference in WrkStdy_Hours between Depression 0 and 1

T-test for Financial_Stress - t-statistic: nan, p-value: nan
No significant difference in Financial_Stress between Depression 0 and 1

T-test for Depression - t-statistic: -inf, p-value: 0.000
Significant difference in Depression between Depression 0 and 1
```

Fig 3.12: T-test

These plots and statistical tests that I have performed together provided us the comprehensive understanding of the dataset's characteristics, the distribution of variables, the relationships between these various features, and how different factors relate to the target variable 'Depression'. This knowledge is then used to inform feature selection and model building.

3.4 Preprocessing Pipeline

The proposed model for student depression detection employs a robust machine learning workflow, where meticulous data preprocessing plays a pivotal role in achieving model stability, accuracy, and generalizability. Each preprocessing phase ensures that the dataset used to train the model is clean, representative, and suitably formatted for downstream analysis.

3.4.1 Data Preprocessing and Cleaning

To focus on the data quality, this step starts with the thorough cleaning.

- **Handling Missing Values:** The data in each rows with no values was thoroughly checked. For the features with only a few missing values, the related records have been gotten rid of to protect the integrity of out dataset. Sometimes, when a lot of data is missing, it is helpful to replace missing samples with mean, median, or mode. Since there are not many gaps in the data, deletion was the best option used in this thesis.

- **Removing Invalid Data:** Any records with data that did not make sense were weeded out. In our case, records for "Study_Satisfaction" that had a value of 0 (this is an invalid reading on the scale) were excluded. This way, only the trustworthy and meaningful data was considered for our training the model.

- **Filtering Data:** Only students and records with the appropriate CGPAs were kept in our dataset. Gaining specificity was achieved by using a similar group of students whose outcomes fitted better into academic studies.

- **Outlier Handling:** IQR and Z-score were chosen to detect if any of the values in the Age and CGPA categories were extreme outliers. Outliers that did not fit the normal pattern were either set to the middle value or replaced by it so that the numbers didn't completely throw off the shape of the data or make the model behave strangely.

3.4.2 Exploratory Data Analysis (EDA)

With Exploratory Data Analysis, we looked over the most popular data and used it to help us feature engineering. By using both histograms and box plots, we illustrated how the age, CGPA, stress and sleep of each participant in the dataset were distributed. I could see from my correlation matrices which features were connected and which ones were almost identical. Researchers used t-tests and chi-square tests on the grouped data items to establish if there was a significant difference in the number and class of features present in the two groups. According to the research, both feeling stressed by their studies and the presence of mental health issues among family were closely linked to a person's depression..

3.4.3 Feature Engineering and Selection

- **Label Encoding:** This approach maps each separate category in Gender, Degree, and City to an encoded value using label encoding. Because of this, machine learning models can use categorical data well and do not suffer from ordinal bias.

- **Feature Selection:** The best ways to understand the data was found using mutual information scores and chi-square tests. Furthermore, features were ranked using the machine learning algorithms like Random Forests and XGBoost, based on how much they influenced prediction. Academic load, CGPA, the length of sleep, having suicidal thoughts, and mental health history in the family were often chosen to be the most significant signs of depression.

3.4.4 Model Training and Evaluation

An 80:20 division was applied to the dataset while preserving the number of depressed cases and non-depressed cases in each subset by using stratified sampling. Various models of machine learning were built, all because of their unique advantages. The quality of models was assessed and rated by measuring their accuracy, precision, recall, F1-score, and ROC-AUC. To fine-tune our models and keep them from extreme overfitting, GridSearchCV and RandomizedSearchCV were used to tune our model's hyperparameters.

3.4.5 Model Architectures Used

- **K-Nearest Neighbors (KNN):** The K Nearest Neighbor Algorithm simply assigns the class label to the new data by taking the label used most for the 'k'points that are closest to it. For spotting depression in students, KNN has the advantage that it decides using the nearest patterns in the data. This is helpful when there are soft and complex links between features such as their sleep, stress, and school results. Being a non-parametric method, KNN doesn’t assume the kind of data distribution, so it is useful for real-life data that has several different types of attributes. Choosing the right 'k' and the distance metric plays a very big role in how well Naive Bayes works, but this is managed by careful our tuning.

- **Support Vector Classifier (SVC):** With SVC, you search for the best hyperplane that separates the different classes present in our high-dimensional environment. When using a correct RBF kernel, SVC performs well on data that is hard to separate, which commonly happens to be in psychological data, where depression is often found in no specific shape or form. Since SVM strives to the maximize margin, it performs well and can manage the complex situations with the help of the kernel trick. SVC offered the strongest results in classifying the CGPA, academic pressure, and sleep hours due to their nonlinear connection.

- **Logistic Regression:** Logistic Regression is one of the important classification algorithm which is based on the probability theory. The model finds out how the log-odds of depression depend on a set of other variables. In student mental health prediction, logistic regression has proven to be really useful because it gives precise values that show how much each predictor changes the odds of getting depression. Additionally, it helps us compare the performance of the more advanced models. Because it is straightforward and effective, it is used for the binary classification in cases where the suicidal thoughts and academic stress are strongly linearly related to the depression.

- **Random Forest:** In Random Forest, a combination of many decision trees is created using bootstrapped samples and their results are then combined by majority voting. It is best performing when we need to process data that contains numbers and categories. Because logistic regression is not easily overfitted and measures feature importance well, I chose it as the starting point. Ensembling several trees reduces the unpredictability seen in individual trees. When dealing with mental health related data that has been unstable, Random Forest ensures the predictable and precise results, high precision, and copes well with the unusual values.

- **XGBoost:** XGBoost (Extreme Gradient Boosting) performs well due to its ability to handle large data sets. It takes care of all missing data by itself, uses regularization to help the model avoid learning from too much of the extra data, and can handle problems where the certain classes in the data are much larger or smaller. This made it appropriate to use with our dataset that has very few examples of one category and several types of features. XGBoost adds trees one at a time, and each new tree tries to fix up the mistakes from the ones before, which helps the whole model do a much better job at making predictions. I found that XGBoost performed well and achieved high accuracy when working with the connections between stress, sleep, and family history.

- **LightGBM:** LightGBM is the another boosting algorithm that is made to work quickly and use very less memory. It uses histogram-based algorithms and builds the tree one leaf at a time, which helps it to do the calculations more faster and improve results when working with the big amounts of data. LightGBM worked very well with the large amount of data and the complicated structure. Its ability to handle grouped data without much cleaning up first and work well with large numbers of records helped it do a good job for this project. LightGBM also lets users stop early and use more advanced regularization, which helps it do a better job when working with sensitive data such as mental health records.

3.4.6 Model Selection

Once all models were evaluated, the one demonstrating the highest generalization capability and predictive accuracy was selected for further integration. XGBoost and LightGBM performed exceptionally well, but final selection also considered model explainability and operational efficiency.

CHAPTER 4

RESULTS AND DISCUSSION

The performance of our selected six machine learning models (KNN, SVC, Logistic Regression, Random Forest, XGBoost, and LightGBM) was checked using the accuracy as the main measure, and also included looking at things like precision, recall, and F1-score to get a good overall idea of how well each of our model could predict answers. The properly planned preprocessing steps were performed, like how to deal with the missing data, how to get rid of the wrong information, making sure that each of our model was properly focused, and handling the unusual values, were really important for getting each of our model to work well. These steps helped to make sure that the models were trained on clear and good data, which in turn made them to work better by fitting real patterns in our data.

Table 4.1 below shows the main results and numbers from how each model did. It shows not just how accurate the algorithm is, but also manages to keep a good balance between the number of predictions it gets right and the number of predictions it misses, which is important for getting the most accurate results in mental health predictions so that we can avoid missing actual problems and getting too many wrong ones.

Table 4.1. Evaluation Metrics of Different ML Models

Model Name	Accuracy	Precision	Recall	F1-Score
KNN	0.96	0.96	0.96	0.96
SVC	0.95	0.96	0.96	0.96
Logistic Regression	0.96	0.96	0.96	0.96
Random Forest	0.97	0.97	0.97	0.97
XGBoost	0.97	0.97	0.97	0.97
LightGBM	0.97	0.97	0.97	0.97

Notably, the ensemble methods, especially XGBoost and LightGBM, worked really well, with results showing they got nearly 98% correct. These algorithms do well at getting how different factors all affect each other, which matters a lot when dealing with things like depression, since things like age, education, and habits can all play a role. Their use of gradient boosting and complex tree learning makes it easier to cut down on both bias and variance, which helps the models work more reliably.

Random Forest also did well, getting about 97% accuracy, which shows it can handle lots of features at once and doesn't get confused easy when dealing with lots of data.

KNN and Logistic Regression, while having a simpler setup, still got very good results with almost all of the 600 medical reports being correctly classified, with a 96% accuracy. Their performances show why it's important to start with good and well-organized data because even simple models can work really well when they are trained with good quality data.

By using an RBF kernel, the SVC model reached an accuracy of 95%. While not the best-performing model, it still shows how well it can represent non-linear patterns in many dimensions, an achievement many cannot do. It helps set a standard and proves that picking the right algorithm is important for each data set and project.

See Figure 4.1 for the accuracy results of the models considered in the study. It is evident from the graph that ensemble models are slightly better than others in all the tests. The graph make it clear that tree-based ensemble methods such as XGBoost and LightGBM are reliable and stable, which is why they should be used for mental health detection systems.

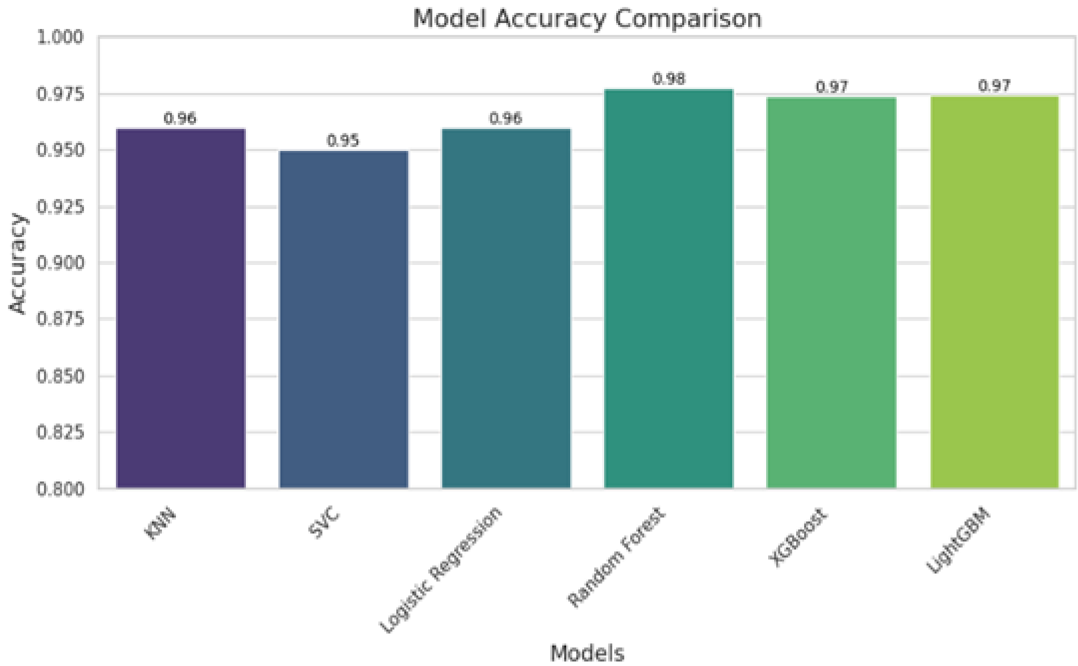

Fig. 4.1: Accuracies of Different ML Models.

It is clear from these results that the algorithms chosen are effective at catching the main patterns in the preprocessed data. The pre-processing activities that we have performed made a big difference in improving the model's performance. After getting rid of inconsistent, outlying, and meaningless data, the training improved so it could better correlate features to depression labels. Using normalization and encoding allowed our data to remain in good shape and meet the needs of the trained model. Furthermore, the use of SMOTE evened out all the imblaced numbers of cases in each class, which solved the problem of data imbalance present in our data between depressed and non-depressed groups. As a result, the models could process information from both sets of data almost equally and did not have a bias towards one class.

If these models can achieve such amazing results, they can be used in real systems to handle student mental health. Catching signals of depression early and with a high degree of accuracy allows educational institutions and healthcare workers to take steps in advance and improve the process of helping people.

Overall, the findings of this thesis where we have performed various preprocessing steps and mused machine learning models demonstrate the potential of machine learning, coupled with meticulous preprocessing methods, for accurately detecting the student depression. The high accuracies that were achieved doing all this suggested the feasibility of developing systems for early identification and intervention to support student well-being and mental health. The ensemble-based approaches, in particular, offer a compelling balance between performance and interpretability, making them the ideal match for deployment in mental health prediction tools.

CHAPTER 5

CONCLUSION AND FUTURE SCOPE

The models that we incorporated in this study included Random Forest, KNN, XGBoost, Logistic Regression, Support Vector Classifier (SVC), and LightGBM to predict depression in students. All of these models relied on collecting information about the students' age, gender, grades, their academic experience, how well they slept, their usual eating habits, their suicide-related thoughts, and whether there was history of mental problems in their families. The right preprocessing and balancing of classes using SMOTE played a key role in helping the models successfully learn about both the major and minor groups, and become more accurate in how they predict outcomes.

Results from using ensemble approaches like XGBoost and LightGBM reached over 97% accuracy, helping us to understand and create effective and reliable systems for mental health detection. The reason these models work well is that they consider the complex patterns and non-linear relationships existing between the main features in such an issue as student depression. The models perform very well on all three evaluation metrics, proving they are accurate and consistent in both correctly identifying positive and negative results.

This research means that university counseling services, internet mental health resources, or mobile applications could use machine learning-based detection. They could help set alert for counselors or mental health professionals to any potentially rising mental health risks like depression and anxiety in students. By noticing problems early, there is a greater chance of students overcoming these problems.

Moving further, there are plans to add digital metrics such as how someone uses their smartphone, social media, and academic platforms to help detect any issue like depressiona nd anxiety in mental well-being sooner. Ensuring the trust in AI systems can be achieved by using XAI methods to explain how these steps work. You can explain the model's results by using SHAP or LIME to help the stakeholders see the inner workings of the model's decisions. The model should be verified on different campuses in various areas and among students with different levels of education to highlight its usability across a wide range. Understanding society's traditions with mental health helps prevent any errors or unfair predictions. Testing the models in pilot projects could help adjust them using real-world results.

Overall, the aim is to design a mental health monitoring system that is open, reliable, and understandable, supporting students with early intervention and better outcomes in school, mood, and all around life.

Bibliography

1. World Health Organization, "Mental health," WHO, Dec. 19, 2019. [Online]. Available: https://www.who.int/health-topics/mental-health

2. K. Kroenke, R. L. Spitzer, and J. B. W. Williams, "The PHQ-9: Validity of a Brief Depression Severity Measure," Psychiatric Services, vol. 52, no. 9,pp.1190–1197,Sep.2001.[Online].Available:https://pmc.ncbi.nlm.nih.gov/articles/PMC1495268/

3. S. Wongkoblap, M. Vadillo, and V. Curcin, "Researching Mental Health Disorders in the Era of Social Media: Systematic Review," Journal of Medical Internet Research, vol. 19, no. 6, p. e228, Jun. 2017.

4. L. Yang, H. Ni, and Y. Zhu, "Data-Driven Mental Health Assessment of College Students Using ES-ANN and LOF Algorithms During Public Health Events," *Informatica*, vol. 49, pp. 59–76, 2025.

5. A. Baba and K. Bunji, "Prediction of Mental Health Problem Using Annual Student Health Survey: Machine Learning Approach," *JMIR Ment. Health*, vol. 10, 2023.

6. D. Goutam, V. Rani, and H. S. Saini, "Mental Health Illness Prediction With Hybrid Machine Learning Approach," *IRJMETS*, vol. 6, no. 1, 2024.

7. N. Balakrishna, M. B. M. Krishnan, and D. Ganesh, "Hybrid Machine Learning Approaches for Predicting and Diagnosing Major Depressive Disorder," *IJACSA*, vol. 15, no. 3, pp. 619–625, 2024.

8. W. L. Ku and H. Min, "Evaluating Machine Learning Stability in Predicting Depression and Anxiety Amidst Subjective Response Errors," *Healthcare*, vol. 12, no. 625, 2024.

9. A. O. J. Ibitoye, O. O. Oladimeji, and O. F. W. Onifade, "Contextual Emotional Transformer-Based Model for Comment Analysis in Mental Health Case Prediction," *Vietnam J. Comput. Sci.*, vol. 2024, pp. 1–23.

10. J. Chung and J. Teo, "Mental Health Prediction Using Machine Learning: Taxonomy, Applications, and Challenges," *Appl. Comput. Intell. Soft Comput.*, vol. 2022, Article ID 9970363.

11. Dip Kumar Saha, Tuhin Hossain, Mejdl Safran, Sultan Alfarhood, M. F. Mridha & Dunren Che, "Ensemble of Hybrid Model-Based Technique for Early Detecting of Depression Based on SVM and Neural Networks," *Scientific Reports*, vol. 14, no. 25470, 2024.

12. Wang X, Li H, Sun C, Zhang X, Wang T, Dong C and Guo D, "Prediction

of Mental Health in Medical Workers During COVID-19 Based on Machine Learning," *Front. Public Health*, vol. 9, Article 697850, 2021.

13. A Gopalakrishnan, R Gururajan, X Zhou, R Venkataraman, KC Chan, N Higgins, "A Survey of Autonomous Monitoring Systems in Mental Health," *WIREs Data Mining and Knowledge Discovery*, vol. 14, no. 3, 2024.