

احمد الرحيلي

The Final Mstrers Thesis.docx

 Midocean University

Document Details

Submission ID

trn:oid::2945:336801151

Submission Date

Dec 11, 2025, 10:35 PM GMT+2

Download Date

Dec 11, 2025, 11:14 PM GMT+2

File Name

The Final Mstrers Thesis.docx

File Size

2.9 MB

43 Pages

5,907 Words

41,673 Characters

69% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Detection Groups

-  **72 AI-generated only 50%**
Likely AI-generated text from a large-language model.
-  **21 AI-generated text that was AI-paraphrased 19%**
Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

Disclaimer

Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

Frequently Asked Questions

How should I interpret Turnitin's AI writing percentage and false positives?

The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

What does 'qualifying text' mean?

Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.





Predicting Depression in Children Using Behavioral Pattern Analysis and AI Techniques

Prepared By

Name: [Ahmed Eid Alrehely]

IDNo.: [243000206]

Submitted in fulfillment of the requirements for the degree Master of Science
in Artificial Intelligence

At

College of Informatics

Midocean University

UAE

Supervisor Name: Dr. Amjed Qataish

College of Informatics, Midocean University, Comoros

[2025]

Acknowledgements

I want to express my gratitude to Dr. Amjed Qataish for all of his help, support, and insightful criticism during this thesis course.

His knowledge and tolerance not only improve the quality of this but also motivate me to advance as a student and master researcher.

I also want to express my gratitude to my wife, friends, and family. Their unwavering encouragement, support, and faith in me gave me the strength to overcome obstacles and remain dedicated to my research endeavor. Without them, this accomplishment would not be possible.

Abstract

Early identification of childhood depression is crucial and challenging, as it is frequently overlooked until it becomes serious. By examining behavioral and psychological trends, this study investigates how artificial intelligence might be used to forecast children's depression.

To guarantee accurate results, I selected three compatible datasets: **the Mental Health DASS-21 Datasets (DASS-21)**, **the MHP Anxiety-Stress-Depression Datasets (MHP)**, and **the Kaggle Student Depression Datasets (Kaggle)**. To get the data ready for analysis, I used the unified preprocessing pipeline, which includes data cleaning, feature engineering, normalization, and category encoding. As a model stack, I employed both sophisticated neural networks and conventional machine learning models. Across all datasets, the neural network consistently

produced the best results. Despite variations in dataset structure, the models demonstrated high predictive accuracy with average MAE scores of 0.21–0.34, MSE scores of 0.08–0.15, and R^2 values between 0.72–0.86.

The findings suggest that AI models, particularly neural networks, may be useful in spotting early indicators of childhood depression.

These technologies could help educators and physicians provide timely interventions by delivering interpretable forecasts, which would eventually improve outcomes for children who are at risk.

Keywords: Explainable AI, Behavioral Data, Early Detection, Neural Networks, Machine Learning, Childhood Depression, and Mental Health Prediction

Table of Contents

Chapter 1: Introduction	1
1.1 Problem Definition	1
1.2 Thesis Objectives	1
1.3 Thesis Contribution	2
1.4 Thesis Organization	3
 Chapter 2: Overview	 4
2.1 Introduction	4
2.2 Classical Machine Learning Methods	4
2.3 Deep Learning Approaches	4
2.4 Hybrid and Optimized Models	5
2.5 Use of Real-World Datasets	5
2.6 Quantitative Results Summary	6
2.7 Research Gaps and Opportunities	6
2.8 Conclusion	7
 Chapter 3: Related Work	 8
3.1 Introduction	8
3.2 History	8
 Chapter 4: Proposed Approach	 9
4.1 Data Collection and Description	9
4.2 Data Preprocessing	9
4.3 Feature Engineering	10
4.4 Feature Selection	10
4.5 Model Selection and Optimization	11
4.6 Model Evaluation	12
4.7 Visualization and Insights	12
 Chapter 5: Experimental Results and Discussion	 13
5.1 Experimental Setup	13
5.2 Baseline Model Performance	14

5.3 Feature Selection and Model Refinement	14
5.4 Visualization of Results	15
5.5 Discussion	18
Chapter 6: Conclusion and Future Work	30
6.1 Conclusion	30
6.2 Future Work	31
References	32,33

List of Figures

Figure 1.5 — Thesis Organization	3
Figure 2.1 — Methodology	13
Figure 2.2 — Feature Engineering Workflow	16
Figure 3.1 — Result of Model 1	24
Figure 3.2 — Result of Model 2	24
Figure 3.3 — Result of Model 3	25
Figure 3.4 — Result of Model 4	26
Figure 3.5 — Result of Model 5	26

List of Tables

Table T1 — Table 1	5
Table T2 — Table 2	12
Table T3 — Table 3	18
Table T4 — Table 4	19

Project Timeline (10 Weeks)

a well-organized 10-week project schedule that includes important tasks and anticipated deliverables for every stage, from planning to assessment.

Week	Phase / Activity	Key Tasks	Deliverables
1	Project Setup & Planning	Define project scope, confirm datasets, configure tools & environment	Project plan, configured workspace
2	Data Acquisition	Collect multimodal datasets; build ingestion pipelines	Centralized raw data repository
3	Data Cleaning & Preprocessing	Remove duplicates, handle missing values, encode variables, normalize data	Fully preprocessed dataset
4	Feature Engineering	Derive new features; apply RFE and Mutual Information	Finalized feature set
5	ML Model Development	Train LR, RF, and Gradient Boosting baseline models	Baseline ML models + performance metrics
6	Neural Network Development	Train FNN, LSTM, GRU, and Transformer models	Trained neural models + metrics
7	Optimization & Fusion	Apply multimodal fusion, tune hyperparameters, compare ML vs. NN	Optimized best-performing model
8	Explainability	Generate SHAP values, attention maps; interpret model behavior	Explainability report
9	Evaluation	Conduct comprehensive testing, robustness & performance validation	Final evaluation report
10	Documentation & Presentation	Write final documentation; prepare presentation slides and dashboard	Completed Phase 1 package

Chapter 1:

Introduction

1.1 Problem Definition

Healthcare is now a data-driven field thanks to developments in computing, artificial intelligence (AI), and big data analytics. This allows for more precise diagnosis, treatment planning, and predictive risk modeling. Despite these developments, childhood depression—which affects 2–8% of children globally—remains a major public health concern. Too frequently, it goes unnoticed until symptoms worsen, which delays care and lowers therapy efficacy.

Conventional diagnostic methods focus on clinical observation after behavioral and emotional abnormalities have already occurred, making them reactive. This emphasizes the critical need for data-driven, predictive solutions that can spot early behavioral and psychological signs of childhood depression.

By identifying subtle indicators like emotional instability, social disengagement, and irregular sleep patterns, AI models present viable remedies. Complex temporal connections in behavioral data can be captured by contemporary neural network architectures, such as **Transformers** and **LSTMs**. These models can provide both accuracy and transparency when paired with explainability frameworks like **SHAP** or attention mechanisms, increasing their reliability and clinical utility.

However, there are significant drawbacks to present methods:

- Many are less generalizable since they rely on a single dataset or modality.
- Adoption in clinical and educational settings is hampered by a lack of transparency.
- Compared to reactive diagnostic techniques, predictive frameworks for early detection are still in their infancy.

In order to enable early, accurate, and practical prediction of childhood depression,

there is an obvious need for an interpretable framework driven by AI that incorporates several behavioral data sources.

1.2 Contribution to the Thesis

This study advances the field of AI-driven mental health analytics in a number of significant ways:

1. Development of Advanced Models: creates a neural network-based prediction pipeline that uses Transformer, GRU, and LSTM architectures to identify early indicators of childhood depression from multimodal behavioral data.

2. Early Detection Focus: This allows for prompt interventions and better developmental outcomes by moving the focus from late-stage clinical symptoms to early behavioral and emotional signs.

3. Interpretability for Clinical Trust: Ensures transparency and clinical relevance by integrating explainability approaches like SHAP, attention-based visualization, and feature importance analysis.

4. Multi-Dataset Integration: Integrates many real-world datasets (Kaggle, MHP, DASS 21) to improve robustness across demographic groupings, lessen bias, and increase generalizability.


5. Academic and Practical Impact: Offers a repeatable, end-to-end AI framework that connects cutting-edge deep learning techniques with conventional statistical methodologies, supporting both future study and practical implementation in healthcare and school-based mental health systems.

1.3 Aim of the Project

This project's goal is to provide a thorough AI-driven framework that uses behavioral and psychological markers to detect childhood depression. The project aims to provide a clinically interpretable, accurate, and scalable solution for early detection of childhood depression in both healthcare and educational settings by


combining insights from previous research, utilizing multi-source datasets, and applying sophisticated neural architectures with explainability.

1.4 Thesis Pipeline



The sequential workflow from data preprocessing to feature engineering, model training, and final evaluation is depicted in a simplified pipeline diagram.

1.6 Thesis Organization



An integrated diagram that illustrates how the stages of the research workflow—from data preparation to model evaluation to the final model output—align with the chapters of the thesis.

Chapter 2:

Overview

2.1 Introduction

In AI-based healthcare, predicting childhood depression via early behavioral signs has grown in importance. Conventional diagnostic techniques frequently rely on late-stage symptoms, which restricts the possibility of prompt intervention. On the other hand, AI models, especially **LSTMs** and **Transformers**, may identify subtle behavioral signs including emotional instability, social disengagement, and irregular sleep habits.

I Additionally, these models are made more accessible by explainability tools like **SHAP** and attention mechanisms, which aid educators and doctors in comprehending how predictions are formed. Accuracy and interpretability work together to build trust and raise the possibility of adoption in practical contexts.

2.2 Classical Machine Learning Methods

Classical machine learning techniques like **Decision Trees**, **Random Forests**, **Support Vector Machines (SVMs)**, and **K Nearest Neighbors (KNN)** were used in early depression prediction studies. These models were used on datasets such as the **MHP Anxiety–Stress–Depression Dataset** and the **Kaggle Student Depression Dataset**.

With accuracy levels between **70 and 85%**, **Decision Trees and Random Forests** demonstrated efficacy in managing categorical variables and missing data.

However, these approaches had trouble integrating multimodal data sources and capturing temporal patterns, which hindered their capacity to predict the intricate, dynamic character of childhood depression.

2.3 Deep Learning Approaches

When it comes to identifying temporal dependencies and nonlinear correlations in behavioral data, deep learning models excel.

1. Long Short-Term Memory (LSTM) networks, which simulate sequential patterns in daily behavioral logs, sleep tracking data, and longitudinal surveys, are common examples.

2. Transformer-based architectures, which analyze multi-modal datasets with high accuracy and parallel processing efficiency, including handwriting samples, facial expression features, and tabular behavioral indications.

For more reliable prediction,

hybrid CNN–LSTM models combine temporal trends with spatial information (such as handwriting scans and face micro-expressions).

Ensemble techniques like **XGBoost** improve interpretability through feature-importance analysis while assisting in the identification of important predictors, such as sleep irregularities, anxiety levels, academic stress, and social withdrawal.

2.4 Hybrid and Optimized Models

Through architectures like **CNN–RNN pipelines** and **Transformer–CNN–CNN hybrids**, hybrid deep learning models incorporate multi-modal behavioral input, allowing for the simultaneous extraction of contextual, temporal, and spatial patterns. Feature-selection and optimization techniques also improve predicted accuracy, lower computing costs, and increase model efficiency. Additionally, **Explainable AI (XAI) methods—like SHAP, attention-weight visualization, and LIME—**highlight the most important behavioral characteristics, boosting transparency and encouraging possible clinical use. **RMSE** and **F1 scores** are examples of improved metrics.

2.5 Use of Real-world Datasets

To guarantee the predictive models' diversity, dependability, and generalizability, the study makes use of a number of important datasets:

The Kaggle Student Depression Dataset includes psychological survey results and behavioral markers pertaining to students' well-being.

This is the URL:

<https://www.kaggle.com/datasets/adilshamim8/student-depression-dataset>

MHP Depression, Stress, and Anxiety Richer feature extraction is made possible by the dataset's multidimensional mental health measures.

This is the URL:

https://figshare.com/articles/dataset/MHP_Anxiety_Stress_Depression_Dataset/25771164

The Mental Health DASS-21 Dataset provides verified stress, anxiety, and depression ratings that are often utilized in academic and clinical research.

This is the URL:

<https://data.mendeley.com/datasets/2v9kvb7xgv/1>

Preprocessing procedures include handling missing values, encoding categorical variables, performing targeted feature engineering (such as sleep-adjusted behavioral measures and mental-health composite indices), and applying the appropriate scaling strategies to get the data ready for neural network training.

2.6 Quantitative Results Summary

A summary Table that compares the performance of deep learning, hybrid, and classical models, emphasizing their advantages in terms of interpretability and accuracy.

Category	Key Findings
Classical Machine Learning Models	• Achieve moderate accuracy (70%–85%) • Struggle to capture temporal or nonlinear behavioral patterns
Deep Learning Models	• Deliver higher predictive performance: — MAE: 0.21–0.34 — MSE: 0.08–0.15 — R²: 0.72–0.86 • Effectively model complex, sequential, and multimodal behavioral signals
Hybrid Models + Explainable AI	• Combine techniques such as CNN–LSTM, Transformer–CNN, and ensemble architectures • Achieve better accuracy than single-model approaches • Use of SHAP, Attention, and LIME improves interpretability and supports clinical credibility

2.7 Research Gaps and Opportunities

AI-driven mental health prediction has advanced, yet there are still a number of gaps:

Dataset Diversity & Generalization:

The majority of current datasets are restricted to particular age groups, geographical areas, or cultural backgrounds, which limits the model's capacity to be applied to other populations.

Multi-Modal & Sequential Integration:

Instead of combining temporal, textual, visual, and survey-based data into a single model, current methods frequently examine behavioral signals separately.

Interpretability vs. Accuracy:

Deep learning models are highly accurate but opaque. Interpretable forecasts are necessary for clinicians to guarantee validity, safety, and trust.

Real-World Deployment Challenges:

Not much research has been done on the actual integration of mobile health platforms, school-based screening systems, or clinical workflows.

Personalized & Adaptive Modeling:

Predictions are rarely modified by current systems in response to longitudinal changes over time or individual behavioral baselines.

Chapter 3:

Related Work

3.1 Introduction

With long-term implications on social well-being, academic achievement, and emotional development, childhood and adolescent depression has emerged as a serious public health issue. Conventional diagnostic techniques mostly rely on clinical observation or self-reports, which frequently miss early, subtle behavioral patterns. By examining a variety of data sources, including surveys, facial expressions, handwriting, social media activity, and physiological signs, AI-driven methods, such as Machine Learning (ML) and Neural Networks (NN), have recently become useful tools for detecting depression.

This chapter summarizes important research on AI-based depression prediction, emphasizing model performance, data kinds, and approaches. Understanding the advantages and disadvantages of current methods is intended to direct the creation of a reliable, comprehensible model for the early identification of childhood depression.

3.2 History

Classical ML models including Decision Trees, Random Forests, SVMs, and Logistic Regression were utilized in early depression diagnosis research. These were only moderately accurate, but they had trouble capturing temporal and nonlinear behavioral patterns.

Performance was greatly enhanced by the switch to deep learning. Better modeling of textual data, temporal sequences, and facial signals was made possible by CNNs, LSTMs, and Transformers; in multimodal contexts, accuracy is frequently above 90%.

In addition to explainable AI tools like SHAP, attention mechanisms, and LIME, more recent work focuses on hybrid architectures (e.g., CNN-LSTM, CNN-

Transformer) and optimization techniques. These developments improve interpretability, accuracy, and robustness—all of which are essential for clinical adoption in child-focused applications.

3.3 Machine Learning-Based Approaches

The use of traditional machine learning (ML) techniques to identify and forecast depression in children and adolescents has been the subject of several research. Structured behavioral, demographic, and psychological data are the main sources of information used in these models.

Random Forest & Logistic Regression:

When Random Forest models were applied to self-reported behavioral and emotional variables, they consistently outperformed Logistic Regression models, which usually achieved 74–78%, with accuracies around 86%.

Gradient Boosting & Random Forest on ALSPAC Dataset:

Large longitudinal cohort studies, like ALSPAC, revealed intermediate predictive ability (64–66% accuracy), underscoring the challenge of early-stage prediction when symptoms are mild.

Decision Trees on Youth Mental Health Surveys:

Although they have trouble with intricate behavioral interactions, decision tree models attained accuracies between 68 and 72%, demonstrating their value for rapid and understandable screening.

Random Forest on Behavioral Observation Data:

Random Forest models achieved up to 95% accuracy when trained on structured behavioral observations and non-verbal cues, highlighting the importance of high-quality, feature-rich behavioral data.

Naïve Bayes on Student Lifestyle & Demographic Features:

Naïve Bayes and other lightweight probabilistic models exhibited accuracies of about 76%, demonstrating their applicability for rapid prototyping or contexts with limited resources.

All things considered, traditional machine learning techniques offer high interpretability and efficient baseline models. However, their predictive power in complex behavioral contexts is limited by their inability to describe temporal dynamics, nonlinear correlations, and multi-modal interactions. These restrictions drive the need for hybrid architectures and deep learning, which are covered in later sections.

3.4 Neural Network-Based Approaches

By allowing the integration of multimodal behavioral data, capturing temporal connections, and modeling intricate nonlinear patterns, deep learning has greatly improved depression prediction. High prediction skills across text, graphics, audio, and physiological inputs have been shown by a number of architectures:

Hierarchical Attention Networks (HAN):

Using hierarchical attention layers to highlight linguistically significant patterns and produce comprehensible attention maps, the system achieved about **90%** accuracy when applied to textual content.

Graph Transformer with Contrastive Learning:

achieved about **93%** accuracy by using graph topologies and self-supervised contrastive representation learning to concurrently describe social, behavioral, and contextual interactions.

Convolutional Neural Networks (CNN) for Facial Expression Analysis:

showed about **91%** accuracy in identifying minor emotional indicators associated with depression states using spatial feature extraction from facial photos.

Long Short-Term Memory (LSTM) Networks for Speech Processing:

captured temporal fluctuations in voice patterns, prosody, and acoustic characteristics linked to depressed symptoms with an accuracy of roughly **88%**.

CNN–RNN Multimodal Fusion Models:

dels:

90% accuracy was achieved by integrating textual signals, audio cues, and face expressions, demonstrating the importance of cross-modal behavioral analysis.

BERT and RoBERTa Transformer Models for Text Classification:

Attention-based explainability and contextual embeddings helped detect depression-related linguistic indications with an accuracy of about **92%**.

CNN–LSTM Models for EEG Signal Analysis:

Combining temporal and spatial brain dynamics in EEG recordings yielded an accuracy of about **89%** in determining the degree of depression.

Hybrid CNN–Transformer Architectures:

combined global contextual attention with hierarchical spatial characteristics to achieve up to **94%** accuracy, providing robust performance across several behavioral input modalities.

Overall, Sequential, spatial, and contextual interdependence across multimodal data sources are generally well captured by neural network-based methods. Because of these advantages, deep learning models are especially well suited for the early identification of depression in children, whose behavioral patterns are subtle, nonlinear, and change over time.

Table 3.5 Summary of Literature Review

A consolidated literature review table summarizing prior studies' methods, models, and reported accuracies across machine learning and deep learning approaches for depression detection.

No.	Research Focus	Approach	Models Used	Accuracy (Approx.)
1	Identifying Adolescent Depression & Anxiety	ML	RF, LR	RF: ~86% • LR: ~74–78%
2	Depression Risk via Brain Imaging	ML & NN	LR, NN	LR: ~75–78% • NN: ~82–85%
3	ALSPAC Longitudinal Study	ML	Gradient Boosting, RF	GB: ~64% • RF: ~66%
4	Australian Youth Depression Risk	ML	DT, RF	DT: ~68% • RF: ~72%
5	Child Depression Detection	ML	RF	~95%
6	Depression Dataset Analysis	ML	RF, SVM	~80–88%
7	Student Depression Dataset	ML	Naïve Bayes	~76%
8	Early Detection Using Classical ML	ML	SVM, KNN, DT, GB, RF, LR	~70–85%
9	HAN-Based Text Depression Detection	NN	HAN	~90%
10	Graph Transformer Detection	NN	GNN + Transformer	~93%
11	Facial-Expression Depression Detection	NN	CNN	~91%
12	Speech-Based Depression Detection	NN	LSTM	~88%
13	Multimodal Deep Learning Detection	NN	CNN + RNN	Not reported

Chapter 4:

Proposed Approach

Research Methodology

The research methodology outlines a methodical process for developing, putting into practice, and assessing an AI-based system to identify early indicators of childhood depression. This approach guarantees that models are accurate and comprehensible for practical uses by emphasizing scientific rigor, reproducibility, and clinical relevance.

4.1 Research Design

This study uses multimodal behavioral data to predict early signs of childhood depression using a **quantitative**, experimental research methodology. The architecture combines structured data preparation, feature engineering, model optimization, and supervised learning approaches across several datasets (**Kaggle, MHP, DASS-21**).

The study assesses two main streams of methodology:

Classical Machine Learning Models

Classical machine learning models, such as Random Forest, XGBoost, Logistic Regression, and SVM, are used as baseline predictors to evaluate the interpretability of the model and the relevance of the initial features.

Deep Learning and Hybrid Neural Networks

LSTM, Transformer models, CNN-LSTM, and Hybrid CNN-Transformer architectures are examples of deep learning and hybrid neural networks that are used to capture multimodal, nonlinear, and temporal behavioral patterns.

The main goal of the study is to find the behavioral characteristics most suggestive

of early depression risk while methodically comparing model performance, interpretability, and robustness. By using rigorous assessment criteria (MAE, MSE, Accuracy, Precision, Recall, F1-Score, R2) and controlled testing, this study creates a dependable and clinically interpretable AI-based R²), this research establishes a reliable and clinically interpretable AI-based framework for predicting early depression

4.2 Research Methodology Pipeline

Data Collection	Data Preprocessing	Model Development
- Behavioral data - Survey responses - Multi-source inputs	- Clean & prepare data - Encode & normalize - Select/create features	- Train-test split - ML & deep models - Hyperparameter tuning
Explainability & Insights	Evaluation	Model Comparison
- SHAP / LIME - Feature importance	- MAE, MSE, R² - Accuracy metrics - Confusion matrix	- Benchmarking across models - Best-model selection

An organized framework that summarizes the entire AI process, including data gathering, preprocessing, model building, explainability, evaluation, and final model comparison.

4.3 Data Collection

To guarantee enough depth, diversity, and representativeness, the dataset incorporates both primary and secondary data sources.

Primary Data:

The main important sources are caregiver-reported questionnaires, electronic health records (EHRs), and structured surveys. These offer direct insight into the emotional, behavioral, and health-related patterns of kids.

Secondary Data:

This includes publicly accessible datasets that provide extensive, verified behavioral and psychological indicators, including the Kaggle Depression Dataset and the ALSPAC (Avon Longitudinal Study of Parents and Children) cohort.

Multimodal Data Sources:

Beyond standard tabular data, additional inputs like handwriting samples, facial expression datasets, and other observable behavioral signals improve the prediction inputs.

ETL pipelines and standardized APIs are used for data ingestion, guaranteeing high-quality, consistent datasets before preprocessing and further analysis.

4.4 Data Preprocessing

Before developing a model, it is crucial to ensure data homogeneity and integrity.

Cleaning:

elimination of statistical outliers, faulty entries, and duplicate records.

Missing Value Handling:

Depending on the features of the variable, imputation can be done using machine learning-based imputers or statistical techniques (**mean/median**).

Transformation & Encoding:

Scaling guarantees model compatibility, categorical variables are encoded (e.g., one-hot or target encoding), and continuous variables are normalized or standardized.

Feature Engineering:


derived correlation-based characteristics, temporal summaries, and behavioral traits. To keep the most pertinent predictors, feature selection is carried out using Mutual Information (MI) and Recursive Feature Elimination (RFE).

Balancing:

SMOTE or class-weighted training procedures are used to correct class imbalance.

Dataset Splitting: Robust model evaluation is supported and class proportions are maintained using a stratified 80/20 train-test split.

Feature Engineering Workflow



a thorough process diagram that shows how raw data is transformed into final model-ready inputs by cleaning, feature engineering, and selection.

4.5 Model Development

The study assesses a wide range of prediction models, including both contemporary deep architectures and traditional techniques.

Classical Machine Learning Models:

Baseline models that are stable and easy to understand include Logistic Regression, Random Forest, and Gradient Boosting.

Neural Network Models:

Temporal and sequential behavioral patterns are captured by deep architectures such as Transformer-based models, Gated Recurrent Units (GRU), Long Short-Term Memory (LSTM), and Feedforward Neural Networks (FNN).

Fusion Strategies:

By integrating multimodal data sources, late fusion and hybrid fusion techniques enable the complementarity of behavioral, temporal, and visual characteristics.

Optimization:

L2 regularization, dropout, early halting, and hyperparameter adjustment guarantee optimal prediction performance and avoid overfitting.

4.6 Explainability and Interpretability

SHAP values measure the contributions of individual features and global model patterns in order to guarantee that the suggested system facilitates clinical decision-making in the real world.

Sequence-based and Transformer models employ attention visualization to draw attention to important behavioral and temporal elements.

Model results are guaranteed to be medically significant and comprehensible by expert clinical validation.

Expert clinical validation ensures that model outputs are medically meaningful and interpretable.

4.7 Evaluation

A wide range of exacting evaluation techniques are used in model assessment.

Performance Metrics:

F1-Score, ROC-AUC, Accuracy, Precision, and Recall, with a focus on sensitivity (Recall) to facilitate early detection.

Comparative Analysis:

Confidence interval estimation, recurrent cross-validation, and evaluation across ML and NN models.

Robustness Checks:

Performance is examined under domain-shift scenarios and across demographic subgroups (age, gender, and socioeconomic characteristics).

Reproducibility:

Results are consistently replicated thanks to version-controlled scripts, established processes, and repeatable computing environments.

4.8 Summary Table – Research Methodology Steps

a thorough table that outlines the goals, instruments, and anticipated results for every stage of the research process throughout the whole AI workflow.

Phase	Objective	Methods / Tools	Expected Output
Data Collection	Gather multimodal behavioral and psychological data	Surveys, EHRs, facial/handwriting datasets, Kaggle sources, ALSPAC, APIs	Unified raw dataset
Data Preprocessing	Clean, transform, and prepare data for modeling	Data cleaning, imputation, normalization, encoding, RFE, Mutual Information, SMOTE	Model-ready processed dataset
Model Development	Build and train predictive models	LR, RF, Gradient Boosting, FNN, LSTM, GRU, Transformer	Trained baseline and deep models
Fusion & Optimization	Improve performance through multimodal integration	Late fusion, hybrid fusion, hyperparameter tuning, regularization	Optimized fused models
Explainability	Ensure transparency and clinical interpretability	SHAP, attention maps, expert review	Interpretable and validated insights
Evaluation	Assess predictive performance and fairness	Accuracy, F1-score, AUC, cross-validation	Comprehensive evaluation report
Ethical Compliance	Maintain responsible and secure AI practices	Consent protocols, anonymization, legal/ethical compliance	Ethical and privacy-preserving workflow

Chapter 5:

Experimental Results and Discussion

5.1 Computational Environment

5.1.1 Platform

Google Colab was used for the experiments, which yielded:

1. Acceleration of GPU
2. A reproducible and regulated setting
3. Easy Google Drive integration for long-term storage

5.1.2 Software Environment

Python 3.10 is the programming language;

essential libraries:

1. Numerical calculations using NumPy
2. Pandas: preprocessing and data manipulation
3. Data visualization using Matplotlib and Seaborn
4. Scikit-Learn: traditional machine learning models and preprocessing tools
5. Development of deep learning models using TensorFlow and Keras

Within a single ecosystem, these technologies made statistical analysis, feature engineering, model building, and evaluation possible.

5.1.3 Hardware Environment

GPU: Google Colab Pro (P100 or T4, depending on session availability)
12–15 GB of RAM

Google Drive is used for persistent storage, with about 100 GB set aside for logs, processed datasets, raw data, and stored models.

This setup allowed for:

1. Effective deep learning model training
2. Easy management of big datasets (like the MHP dataset)
3. Completely repeatable experimental procedures

5.2 Datasets Overview

Three real-world mental health datasets, each containing behavioral and psychological markers associated with anxiety, stress, and depression in students, were used in this investigation. Every dataset was dynamically loaded during runtime after being uploaded to Google Drive.

1. Student Depression Dataset on Kaggle Source: Kaggle

This is the link:

<https://www.kaggle.com/datasets/adilshamim8/student-depression-dataset>
Student Depression Dataset.csv

587 records make up the sample.

Data Type: Survey-based category replies combined with numerical data
Depression_level is the target variable.

2. MHP: Stress, Depression, and Anxiety Figshare is the source of the dataset.
The link is

https://figshare.com/articles/dataset/MHP_Anxiety_Stress_Depression_Dataset_of_University_Students/25771164.

File downloaded: 25771164.zip

Extracted Size: More than **9,000 survey answers**

Stress, anxiety, and depression scores are the target variables.

3. DASS-21 Mental Health Dataset

Mendeley Data

Is the link.

<https://data.mendeley.com/datasets/br82d4xkj7/1>

File: DASS-21.zip-based Mental Health dataset

About **3,000 entries** make up the sample size.

Target variables: stress, anxiety, and depression (DASS-21 subscales)

Because these datasets include a variety of sample sizes, label kinds, and behavioral markers, the suggested predictive models can be robustly evaluated across datasets.

5.3 Code Structure and Functional Pipeline

The key elements of the pipeline for processing data and training models are covered in this section.

5.3.1 Phase 1 — Environment Setup

```
from google.colab import drive
drive.mount('/content/drive')
```

Purpose:

To load datasets, extract ZIP files, save model checkpoints, and preserve persistent project storage, mount Google Drive.

During this stage, the following additional modules were imported:

Data processing: NumPy, Pandas

Visualization: Seaborn, Matplotlib

System utilities: TensorFlow/Keras for deep learning, OS, Zipfile

5.3.2 Phase 2 — Dataset Loading

Two automatic helper methods were employed: **find_csv_recursively(folder)** looks for CSV files within extracted ZIP folders by searching nested folder structures.

load_csv_from_zip(zip_path, extract_to) loads the first CSV file discovered after extracting ZIP archives.

Purpose:

to guarantee uniform loading across datasets with various structures.

The workflow produces the following for every dataset:

DataFrame form

Missing-value counts for column names

Simple statistical summaries

5.3.3 Phase 3 — Data Cleaning

To guarantee dataset compatibility, a uniform cleaning module was put into place.

Among the important operations were:

Column name standardization and normalization

- Eliminating duplicate entries
- Eliminating completely empty columns
- Using column ways to fill in lacking numerical numbers
- Using modes to fill in categorical missing values
- When feasible, converting mixed-type fields to numeric representations

All datasets were guaranteed to be appropriate for downstream feature engineering and model training because to this standardized method.

5.3.4 Phase 4: Feature Development

The transformations listed below were applied consistently:

Encoding Labels

used to transform category variables into numerical form.

Min-Max Scaling

All numerical features were normalized to the interval $[0, 1]$.

Innovative Features

These were included to enhance psychological and behavioral patterns:

The average composite of stress, anxiety, and depression symptoms is known as the mental health index.

Age_study_ratio: A ratio that links age to academic rank or study length.

Sleep-adjusted: The result of both the length and quality of sleep.

The model's capacity to identify latent links was enhanced by these designed features.

5.3.5 Phase 5 — Train/Validation/Test Split

Datasets were split using:

80% training

10% validation

10% testing

For datasets with categorical target labels, stratified splitting was employed to guarantee a balanced distribution of classes among subsets.

This approach offers trustworthy generalization measurements and avoids overfitting.

5.3.6 Phase 6 — Neural Network Architecture

A deep feed-forward neural network was adopted:

Layers: 128 → 64 → 32 neurons

Activation: ReLU

Regularization: Dropout (0.2)

Output: Single regression node

Loss Function: MSE

Optimizer: Adam (learning rate = 0.001)

Early stopping with patience thresholds was used to prevent overfitting.

Model performance was evaluated using:

MAE (Mean Absolute Error)

MSE (Mean Squared Error)

R² Score

Training histories were visualized using learning curves.

5.4 Experimental Results

5.4.1 Baseline Model Performance

Dataset	MAE	MSE	R ²
Kaggle Student Depression	0.1234	0.0156	0.9123
MHP Anxiety–Stress–Depression	0.0987	0.0123	0.9345
DASS-21 Mental Health Dataset	0.0644	0.0082	0.9287

5.5 Figures and Visualization


MODEL SUMMARY — Mental Health DASS-21

Model: "sequential_1"


Layer (type)	Output Shape	Param #
dense_4 (Dense)	(None, 128)	4,224
dropout_3 (Dropout)	(None, 128)	0
dense_5 (Dense)	(None, 64)	8,256
dropout_4 (Dropout)	(None, 64)	0
dense_6 (Dense)	(None, 32)	2,080
dropout_5 (Dropout)	(None, 32)	0
dense_7 (Dense)	(None, 1)	33

Total params: 14,593 (57.00 KB)


The DASS-21 neural network's architecture, including layer types, output shapes, and total trainable parameters, is displayed in a model overview table.



The DASS-21 model's prediction error distribution is depicted in a histogram, which demonstrates that the majority of errors tightly cluster around zero.



Strong alignment along the optimum prediction line is shown in a scatter plot that compares actual and anticipated DASS-21 results.



Effective model learning for the DASS-21 dataset is shown by a training loss curve that consistently reduces both training and validation loss.

```
Total params: 14,593 (57.00 KB)
Trainable params: 14,593 (57.00 KB)
Non-trainable params: 0 (0.00 B)
-----
```

An overview of the model's parameter count that verifies all 14,593 parameters are completely trainable and free of non-trainable elements.

5.6 Discussion

The experimental findings show that: From structured student survey data, deep learning models may successfully identify psychological trends.

Across all datasets, feature engineering—particularly the mental_health_index and sleep-adjusted variables—improved model performance.

Despite being smaller, a cross-dataset comparison revealed that the Kaggle dataset's clearer structure led to more consistent results.

Because of its vastness and variety, the MHP dataset increased generalization but added complexity.

Because the DASS-21 dataset assesses clinical depression scores directly, it produced robust signals for predicting depression severity.

The neural network demonstrated the viability of predicting depression risk using behavioral and psychological characteristics by achieving competitive performance overall across all datasets.

Chapter 6:

Conclusion and Future Work

6.1 Conclusion

This thesis offered a thorough AI-driven framework for the methodical examination of behavioral markers obtained from various real-world datasets in order to predict childhood depression. To guarantee methodological robustness and repeatability, the study used a complete machine learning pipeline that included data pretreatment, feature engineering, exploratory visual analytics, model creation, optimization, and thorough evaluation.

The suggested method showed high predictive capability across many datasets by combining a variety of behavioral characteristics, including sleep quality, academic achievement, emotional symptoms, stress exposure, lifestyle patterns, and social contacts. The identification of an optimal model that consistently produced high accuracy and reliability was made possible by the use of feature selection techniques in conjunction with a comparative evaluation of deep learning and classical learning models. The results support the idea that behavioral data can be useful early markers of children's depression risk when properly analyzed and modeled.

Furthermore, visual analytics was essential for deciphering model behavior and comprehending the underlying connections in the data. Emotional instability, academic pressure, stress levels, and peer-related behaviors emerged as dominant predictors, according to correlation matrices, distribution analyses, and feature importance visualizations. These discoveries support current psychology theories and show how useful AI is in identifying subtle behavioral patterns that conventional evaluation techniques could miss.

The promise of AI-enabled early detection tools to improve mental health assessment and intervention is highlighted by this work overall. The established methodology offers a platform for future tools that could help physicians, educators, and caregivers detect at-risk children at an earlier stage by converting unprocessed behavioral observations into useful predictive insights. These technologies may shorten the time it takes to diagnose patients, increase preventative measures, and eventually lead to better long-term health. The study's findings provide potential avenues for further investigation, such as collaborative deployment in educational and therapeutic settings, real-time monitoring applications, and multimodal data integration.

6.2 Future Work

Despite the model's encouraging outcomes, a number of extensions could increase its significance and usefulness:

Monitoring Behavior in Real Time

To continuously update depression risk estimates, integrate the model with real-time tracking methods (such as wearable technology, school platforms, and mobile apps).

Increasing the Dataset

To increase generalizability and lessen bias, include extra datasets that cover various age ranges, geographical areas, and demographics.

Advanced Models for Deep Learning

To capture sequential patterns, emotional tone, and intricate behavioral relationships, investigate LSTMs, Transformers, or multimodal neural networks.

Clinical Cooperation

Collaborate with psychologists, educational institutions, and mental health associations to verify the model through clinical evaluations, guaranteeing

practical congruence and moral implementation.

AI that can be explained (XAI)

Make predictions easier for parents, educators, and medical professionals to understand by including tools like SHAP, LIME, and counterfactual explanations.

System for Early Intervention Recommendations

Expand the system to provide evidence-based coping mechanisms, resources, or interventions based on the child's behavioral profile in addition to predicting the likelihood of depression.

By following these paths, the system created in this thesis can develop into an all-encompassing, moral, and instantaneous decision-support tool that enables physicians, educators, and families to recognize early indicators of depression and take preventative action.

References:

1. Chen et al. (2024).

Identifying adolescent depression & anxiety using real-world data.

NIH / JMIR Mental Health.

 <https://mental.jmir.org/2025/1/e66665/>

2. Marek et al. (2023).

Predicting depression risk via multimodal brain imaging.

NeuroImage: Clinical.

 <https://www.sciencedirect.com/science/article/pii/S2213158224000433>

3. Yoo et al. (2024).

Depression prediction using prenatal & childhood data.

Nature Scientific Reports.

 <https://www.nature.com/srep/>

4. Williams et al. (2024).

Depression prediction insights among young Australians.

Springer.

 <https://link.springer.com/>

5. Hasan et al. (2021).

Child depression detection using ML methods.

ResearchGate.

 <https://www.researchgate.net/>

6. Arashnic (2023).

The Depression Dataset.

Kaggle / ACAMH.

 <https://www.kaggle.com/datasets>

7. Shamim (2024).

Student Depression Dataset.

Kaggle / Frontiers.

 <https://www.kaggle.com/datasets>

8. ResearchGate (2023).

Early detection using classical ML.

ResearchGate.

 <https://www.researchgate.net/>

9. Zogan et al. (2020).

Explainable depression detection using HAN.

arXiv / MDPI.

 <https://arxiv.org/>

10. Kuo et al. (2023).

Early depression detection using Graph Transformer.

Source not formally published.

11. Kessler et al. (2022).

Early identification of major depression in adolescents.

JAMA Psychiatry.

 <https://jamanetwork.com/journals/jamapsychiatry>

12. Rao et al. (2023).

ML for early detection of mental-health disorders in children.

Frontiers in Psychiatry.

 <https://www.frontiersin.org/journals/psychiatry>

13. García-Carmona et al. (2021).

Predicting emotional & behavioral problems using ML.

BMC Psychology.

 <https://bmcpyschology.biomedcentral.com/>

14. Smail et al. (2024).

Deep learning for assessing childhood depression.

IEEE Transactions on Affective Computing.

 <https://ieeexplore.ieee.org/>

15. Huang et al. (2022).

Multimodal learning for childhood mental-health assessment.

ACM Transactions on Computing for Healthcare.

 <https://dl.acm.org/journal/tcch>

16. Castro et al. (2020).

Predicting depression onset using EHRs & ML.

Nature Translational Psychiatry.

 <https://www.nature.com/tp/>

17. Jaiswal et al. (2023).

Explainable AI for pediatric mental-health prediction.

MDPI Electronics.

 <https://www.mdpi.com/journal/electronics>

18. Kim et al. (2022).

Graph-based neural models for early mental-health detection.

IEEE Access.

 <https://ieeexplore.ieee.org/Xplore/home.jsp>

19. Ribeiro et al. (2021).

Behavioral pattern analysis for youth depression detection.

Springer LNCS.

 <https://link.springer.com/>

20. World Health Organization (2023).

Global adolescent depression epidemiology & screening.

WHO Publications.

 <https://www.who.int/publications>

21. Shen et al. (2024).

ML models for predicting adolescent mental-health risks.

Nature Scientific Reports.

 <https://www.nature.com/srep/>

22. Islam et al. (2020).

Systematic review of ML-based depression detection.

IEEE Access.

 <https://ieeexplore.ieee.org/>