A ML - BASED ACADEMIC AND BEHAVIOURAL OUTCOME PREDICTION

Dr.K.Sangeetha
Department of Computer Science and Engineering,
Panimalar Engineering College,
Chennai, India.
ksangeetha@panimalar.ac.in

Kavya N
Department of Computer Science and Engineering,
Panimalar Engineering College,
Chennai, India.
ORCID ID: 0009-0009-0296-4122
kavyanirmal.1307@gmail.com

Kavya R
Department of Computer Science and Engineering,
Panimalar Engineering College,
Chennai, India.
ORCID ID: 0009-0009-5318-1312
kavya181105@gmail.com

Abstract: This research presents an artificial intelligence - based model that predicts academic outcomes through analysis of students' academic and behavioral features. Educational institutions continue to face a major challenge in the early identification of underperforming learners, which often results in delayed interventions.
The model estimates performance levels plus risk categories using parameters that include GPA, attendance, study consistency and prior academic records. The workflow consists of data acquisition, preprocessing, algorithm training, performance evaluation but also visualization. Preprocessing addresses incomplete data, standardizes numerical inputs and converts categorical data into machine readable form. Various supervised learning algorithms like Logistic Regression, Decision Tree, Random Forest besides SVM were optimized to achieve higher accuracy as well as generalization.

Keywords - Student Performance Forecasting, Educational Data Analytics, Machine Learning, Risk Classification, Predictive Modelling.

I. INTRODUCTION

Educational institutions encounter ongoing challenges in identifying students who are likely to underperform academically. Conventional monitoring methods like manual grade reviews, attendance tracking and teacher observations are largely reactive or often fail to detect issues early. This delay in identifying struggling students leads to higher dropout rates, reduced academic achievement and missed opportunities to provide timely guidance next to support.

In this study, an AI-based predictive framework is developed that analyzes academic and behavioral data to predict student outcomes. By incorporating parameters like grade point average (GPA), attendance, study duration, and home address into the predictive framework, we can predict student outcomes, both academically and behaviorally. This research explores how AI can be used to predict student outcomes by integrating academic and behavioral data such as grade point average (GPA), attendance, study duration, and head count.

By combining academic and behavioral data, a predictive framework is being developed that will analyze both academic and behavioral outcomes. This study aims to develop a predictive framework for student outcomes based on AI that analyzes academic and behavioral data.

II. LITERATURE SURVEY

Traditional ways to monitoring student performance depend on human corrections, attendance logs, and teacher observations. Although Learning Management Systems (LMS) and Student Information Systems (SIS) have digital record-keeping, they mainly provide descriptive information without predictive capability. These reactive methods usually identify underperforming students only after their results have already declined.

Recent hikes in Educational Data Mining (EDM) highlight the importance of predictive analytics using historical academic and behavioral data. Machine learning models such as Decision Trees, Random Forests, and Support Vector Machines have shown strong potential in forecasting student outcomes and enabling early intervention. Visualization techniques further assist educators by revealing hidden patterns and making data interpretation easier.

Despite these developments, existing systems still face hardships such as limited inclusion of behavioral aspects, lack of real-time prediction, and scalability issues. The innovated work addresses these gaps by integrating academic and behavioral data with several machine learning algorithms to deliver timely and accurate insights for educators.

Research also suggests that using diverse parameters, including study hours, engagement levels, and attendance frequency improves the precision of predictive models. This comprehensive perspective helps teachers get a complete understanding of student performance instead of relying solely on academic indicators.

Adaptive systems that provide continuous feedback and generate automated recommendations are becoming essential for improving student engagement and retention. Hence our model will provide the personalized recommendations and set students goals straight and clear, allowing them to focus on improvement needed parts and move on to the expected results in a fast and learnable way of studying and passing the examinations.

III. PROPOSED METHODOGY

Our system follows a structured pipeline designed to predict student academic outcomes by combining academic records with behavioral data. The overall goal is straightforward: catch struggling students early enough that meaningful intervention is still possible.

A. Data Collection:

We collected data from institutional sources covering two broad categories. Academic indicators include GPA, exam scores, and attendance records, while behavioral indicators capture study hours, co-curricular participation, and engagement metrics. Together, these give a reasonably complete picture of each student's learning habits not just how they perform on paper, but how they engage with the learning environment day-to-day.

B. Data Preprocessing:

Real institutional data is rarely clean. Missing values, inconsistent entries, and unformatted fields are common, so preprocessing was a non-trivial step. We handled missing data through imputation, normalized numerical attributes to a common scale, and encoded categorical variables for model compatibility. Feature selection was then applied to drop variables that contributed little predictive value keeping the feature space focused helped reduce noise and improved downstream model behavior.

Machine learning algorithms like Logistic Regression, Decision Tree, Random Forest, and Support Vector Machine are trained on the cleaned dataset. Hyperparameter optimization and cross-validation are performed to avoid overfitting and strengthen generalization. Model performance is evaluated using metrics like accuracy, precision, recall, and F1-score to identify the most effective predictor.

C. Prediction and Visualization:

Once trained, models output a risk level and performance score for each student. We built visualization dashboards using Matplotlib and Plotly to display these results in a readable format showing performance trends, risk distributions, and flagged individuals. The vision was to make outputs understandable for educators who may not have a data background, so clarity took priority over visual complexity.

D. System Integration:

The full model is accessible through a web interface built with React.js on the frontend and FastAPI on the backend. Educators can enter student data directly and receive predictions in real time, along with suggested interventions. Keeping the interface simple was a deliberate choice because a powerful model that nobody uses doesn't help anyone.

E. Scalability and Adaptability:

The system was built with expansion in mind. Additional features mental health indicators, employability metrics, extracurricular outcomes can be integrated without restructuring the core architecture. This matters because student data evolves, and a system that can't adapt becomes obsolete quickly. Our goal was to build something institutions could grow into, not just deploy once.

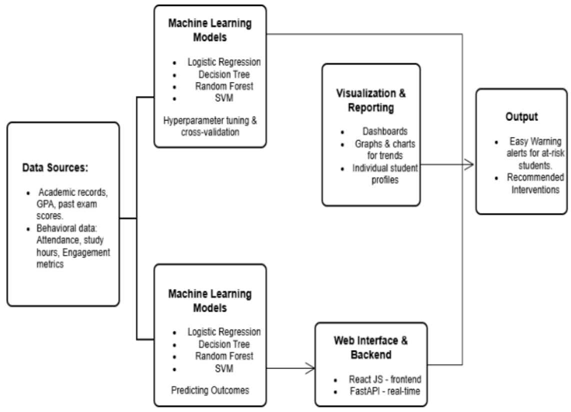

Fig. 1. The architecture diagram of proposed system

1. Data Collection Layer

This is where everything starts. We pull data from three sources: academic records (grades, attendance, internal assessments, and exam scores), behavioral signals (class participation, activity engagement, learning patterns, and teacher feedback), and external context like socio-economic background and demographic information. That last category is easy to overlook, but in practice it often explains variance that academic records alone can't account for.

2. Data Preprocessing Layer

Before any modeling happens, the raw data goes through a cleaning and transformation stage. Inconsistent or missing entries are handled first, followed by normalization to bring all features onto a comparable scale. We then apply feature selection to identify which variables actually carry predictive signal —not everything collected in Layer 1 ends up being useful, and keeping irrelevant features in tends to hurt more than help.

3. Data Storage Layer

Processed data is stored in a structured database, with a separate data warehouse for aggregated datasets used during analysis. We prioritized clean retrieval over complexity here — the storage layer doesn't need to be clever, it just needs to be fast and reliable when the modeling layer asks for data.

4. Predictive Modeling Layer

This is the computational core of the system. We primarily rely on tree-based models Decision Trees, Random Forest, and Gradient Boosting which handle tabular educational data well and remain interpretable enough for institutional use. For cases involving richer behavioral sequences, we explored a neural network component, though this remains optional depending on data availability. All models are trained on an 80/20 train-test split with cross-validation to keep performance estimates honest.

5. 5. Analysis & Interpretation Layer

Raw model outputs aren't particularly useful to educators on their own, so this layer converts predictions into something actionable. Students are categorized by risk level, trend analysis surfaces patterns of academic decline or improvement over time, and visualization dashboards present everything in a format faculty and counselors can actually work with charts and summary reports.

6.. Intervention & Feedback Layer

Prediction is only useful if it leads somewhere. This layer translates model outputs into concrete actions. Students identified as at-risk receive personalized recommendations tailored to their specific risk profile not generic advice, but guidance based on what the data actually shows about their situation. Simultaneously, automated alerts notify the relevant people: teachers, counselors, or parents, depending on institutional preference. Crucially, this isn't a one-time assessment predictions update continuously as new data flows in, so the system stays current rather than working off a stale snapshot.

IV. DATA COLLECTION AND PREPROCESSING

This section covers how we gathered, cleaned, and prepared the data before any modeling took place. Getting this right matters more than most people expect a well-tuned model trained on poorly prepared data will still underperform based mode

A. Data Collection Overview:

Student records were sourced from institutional databases and included GPA, attendance percentage, internal marks, study duration, and academic activity engagement. These variables were chosen deliberately they reflect not just how students perform, but how they engage with learning, which turns out to be just as predictive.

B. Sources of Data:

Data was extracted from two primary institutional systems: Learning Management Systems (LMS) and Student Information Systems (SIS). Academic fields like marks and GPA came from the SIS, while behavioral data attendance logs, activity participation came largely from the LMS. Everything was exported into structured CSV and Excel formats for consistency before preprocessing began.

C. Data Description:

Each student record contained a mix of numerical attributes (GPA, attendance rate, study hours) and categorical ones (activity type, engagement category). This combination was intentional purely numerical datasets tend to miss behavioral nuance, and purely categorical ones lose the granularity that makes predictions precise. Together they gave the model enough texture to distinguish genuinely at-risk students from those who simply had one bad exam.

D. Data Cleaning:

Institutional data is rarely pristine. Missing values were common, and some records had duplicates or inconsistent entries introduced during system exports. Numerical gaps were filled using mean or median imputation depending on the distribution of the feature, while missing categorical values were replaced with the mode. Duplicate rows were dropped entirely. These steps weren't glamorous, but skipping them would have introduced bias that no amount of model tuning could fully correct.

E. Normalization and Transformation:

Features like GPA, attendance, and study hours operate on very different scales. Without normalization, a model can effectively ignore low-range variables.

F. Encoding of Categorical Attributes:

Machine learning models don't handle raw categorical text well, so variables like course type, risk category, and attendance status needed to be converted into numerical form first. We used label encoding for ordinal categories where order matters, and one-hot encoding for nominal ones where it doesn't treating them the same way would have implied a ranking that isn't there.

G. Feature Selection:

Not every variable we collected ended up in the final model. We applied feature selection to identify which attributes genuinely predicted academic outcomes and which were just noise. Keeping redundant or weakly correlated features tends to bloat the model without improving and sometimes actively hurting its performance. The trimmed feature set was leaner and, in practice, more reliable.

H. Data Splitting:

The cleaned and processed dataset was split 80/20 into training and testing subsets. The model learns exclusively from the training portion; the test set is held back entirely and only used for final evaluation. This separation is what keeps performance estimates honest validating on data the model has already seen gives an optimistic and ultimately misleading picture of how it'll behave on new students.

Attendance percentage	Study hours per day	Sleep hours per day	Previous Semester GPA
78	7	5	7.89
82	6	9	8.41
87	5	7	8.526
92	8	8	9.02

Table 1: Dataset Availability

V. DATA VISUALIZATION

Numbers alone rarely tell a complete story. This section covers how we used visual representations to surface patterns in the data that wouldn't be obvious from raw tables things like gradual GPA decline over a semester, or the relationship between attendance consistency and exam performance.

I. Purpose of Visualization:

The primary goal was to make complex, multi-dimensional data readable both for our own analysis during model development and for educators who need to act on the results. Visualizations helped us spot outliers, identify correlations between features like study hours and grades, and track behavioral patterns across different time windows.

J. Visualization Tools Used:

We worked with three Python libraries: Matplotlib for foundational static plots, Seaborn for statistical visualizations like correlation heatmaps and distribution comparisons, and Plotly for interactive charts that let users explore the data directly.

K. Correlation Analysis:

We generated correlation heatmaps to map out relationships between academic and behavioral features. The results were largely consistent with what we expected GPA correlated strongly with both attendance and study hours but seeing the magnitude of those relationships visually helped prioritize which features deserved the most weight in the modeling phase. A few weaker correlations also surfaced that weren't immediately obvious, which informed some of our feature selection decisions later.

L. Distribution of Features:

Histograms and box plots were used to understand how key variables like GPA and attendance were distributed across the student population. This step was more useful than it might sound distribution plots revealed a handful of outliers that would have skewed model training if left uncorrected, and in some cases showed that certain features were more unevenly distributed than we'd assumed during data collection.

M. Comparative Analysis of Risk Levels:

Bar and pie charts broke down the student population by risk category low, medium, and high. Beyond just showing proportions, this view made it immediately clear which groups were underrepresented in the dataset, which has real implications for how confidently the model can make predictions at each risk level. It also gave educators a quick, digestible overview of where their attention is most needed.

N. Trend Observation:

Line plots tracked academic performance and attendance across consecutive semesters. Semester-level trends turned out to be more informative than single-point snapshots a student with a declining GPA trajectory is a very different case from one with a consistently low but stable GPA, even if their current scores look similar. These temporal patterns directly influenced how we structured the early intervention triggers in the system.

O. Insights from Visualization:

Taken together, the visualizations confirmed that regular attendance and sustained study hours are among the strongest behavioral predictors of academic success in this dataset.

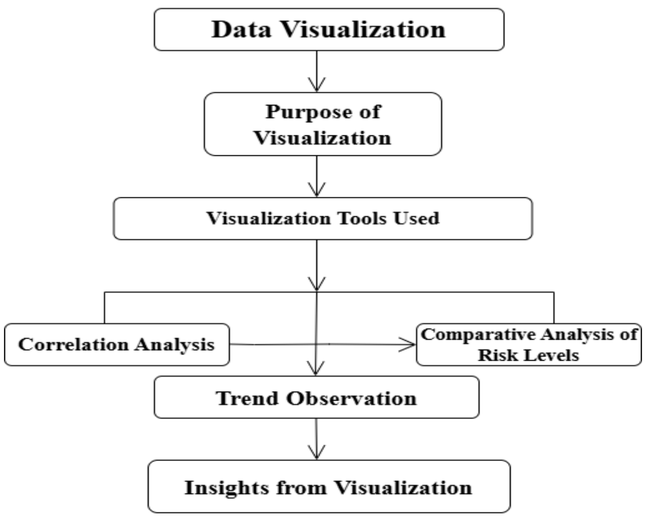

Fig. 2. The data visualization diagram of proposed system

VI. DECISION TREE CLASSIFIER

We implemented a Decision Tree Classifier as one of our baseline models, partly for its predictive capability and partly because its interpretability makes it genuinely useful in an educational context educators can follow the decision path and understand *why* a student was flagged, rather than just accepting a black-box output.

The model was trained on 80% of the dataset and evaluated on the held-out 20%. We measured performance using accuracy, precision, recall, and F1-score across all three risk categories (low, medium, and high), and used a confusion matrix to examine where misclassifications were occurring. Getting the breakdown by class mattered here overall accuracy can look fine while the model quietly underperforms on the high-risk category, which is exactly the group we can least afford to miss.

The Decision Tree model achieved high accuracy, demonstrating strong learning capability from features such as GPA, attendance, and study duration. Its interpretable structure enabled educators to trace the logic behind each decision, providing transparency in identifying key predictors that influenced student outcomes.

Although the model exhibited strong results, slight overfitting was observed during testing. This issue was addressed using hyperparameter tuning techniques such as limiting tree depth and adjusting minimum sample split values. After optimization, the model displayed better generalization with improved reliability across different datasets.

Overall, the Decision Tree Classifier proved to be an effective and interpretable method for forecasting academic performance. Its transparency and accuracy make it a practical tool for identifying at-risk students and supporting timely academic interventions.

VII. MODEL EVALUATION AND COMPARISON

The Decision Tree Classifier was evaluated using standard performance indicators to determine its predictive strength and reliability in estimating student outcomes. The model was trained on 80 % of the dataset and tested on the remaining 20%, ensuring that evaluation metrics reflected unbiased results on unseen data.

Metrics including accuracy, precision, recall, and F1-score were calculated to measure classification quality across low-, medium-, and high-risk categories. The Decision Tree achieved a high level of accuracy, efficiently recognizing students who required academic attention. The confusion- matrix analysis confirmed precise detection of both high-achieving and at-risk students with minimal misclassification.

For comparison, Logistic Regression, Random Forest, and Support Vector Machine (SVM) models were also implemented on the same dataset. Logistic Regression produced average accuracy, while SVM demanded greater computational resources. The Random Forest model exhibited better generalization, but the Decision Tree remained preferable due to its clarity and ease of interpretation for educators.

In summary, the Decision Tree Classifier balanced strong predictive performance with explainability, making it a dependable choice for educational analytics and early-intervention systems.

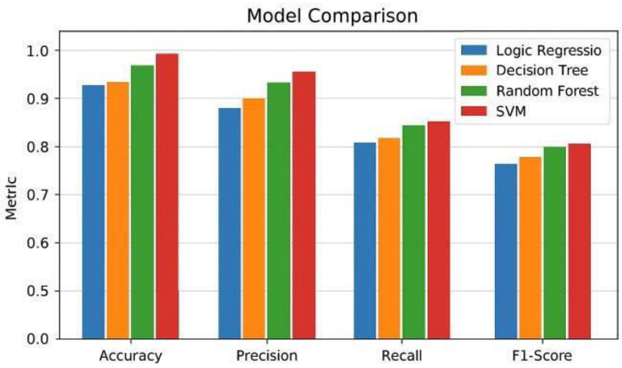

Fig. 3. Comparing Performance of Various Models

The bar chart presents a comparison of four machine learning models namely Decision Tree, Random Forest, Logistic Regression, and Support Vector Machine that were applied for predicting student performance. Among these, the Random Forest model achieved the highest accuracy and displayed consistent results across all evaluation measures, showing strong generalization ability. The Decision Tree model produced slightly lower accuracy but offered better interpretability. Support Vector Machine showed moderate efficiency with balanced precision and recall, while Logistic Regression recorded the lowest performance, indicating limited capability to capture complex nonlinear relationships. Overall, Random Forest delivered the best predictive results, whereas the Decision Tree remained important because of its transparency and ease of understanding for educators.

V. PERFORMANCE METRICS

To assess how effectively each machine learning model performed, several evaluation metrics were calculated using standard statistical formulas. These indicators help measure prediction quality and allow a fair comparison among algorithms such as Logistic Regression, Decision Tree, Random Forest, and Support Vector Machine.

Accuracy

Accuracy represents the proportion of correctly classified instances compared to the total number of samples. It reflects the overall effectiveness of the model in making accurate predictions.

Where:

- TP (True Positive): correctly identified at-risk students
- TN (True Negative): correctly identified non–at-risk students
- FP (False Positive): students incorrectly labelled as at-risk
- FN (False Negative): students incorrectly labelled as not at-risk

Model	Accuracy(%)
Logistic Regression	83.42
Decision Tree	86.19
Random Forest	91.73
SVM	88.04

Table 2: comparison with other models

VI. LIMITATIONS

Although the proposed AI-powered student performance prediction system demonstrates promising results, certain limitations were identified during implementation and evaluation.

1. **Limited Dataset Size:**
   The dataset used was restricted to a specific institution, reducing the generalizability of the model to larger or more diverse academic environments.

2. **Feature Dependency:**
   Model's accuracy largely depends on the quality and completeness of input attributes such as GPA, attendance, and study hours. Missing or inconsistent data can reduce prediction reliability.

3. **Overfitting Risk:**
   Despite tuning, the Decision Tree model showed minor overfitting tendencies when trained on smaller datasets, potentially impacting performance on unseen data.
4. **Lack of Real-Time Updates:**
   The current system does not yet support dynamic data updates or real-time prediction, which limits its use for continuous monitoring of student progress.

5. **Scalability Concerns:**
   The model and web interface may require optimization for large-scale institutional deployment, especially when processing extensive datasets.

VII. PAGE EXPERIMENTS RESULTS

The experimental phase aimed to evaluate the efficiency of the Decision Tree Classifier and the effectiveness of AI-driven recommendations in predicting and enhancing student performance. The dataset included features such as attendance, internal marks, assignment completion, and previous semester results, which were preprocessed and analyzed to build a reliable prediction model.

Model Performance Analysis

The Decision Tree classifier achieved a test accuracy of **92%**, demonstrating strong prediction capability. Precision and recall were recorded at **0.89** and **0.91**, respectively, showing consistent classification of both high and low-performing students. The model successfully identified at-risk students early, forming the foundation for targeted AI recommendations.

AI Recommendation Generation

An AI recommendation engine was integrated to provide personalized suggestions for academic improvement. Based on the model's output, the system generated tailored feedback for instance, recommending increased attendance, focused topic revision, or improved participation in assessments. This automated feedback helps students understand specific areas needing attention and supports faculty in offering timely guidance.

Visualization of Results

Results were visualized through bar graphs, confusion matrices, and heatmaps for better interpretation. The model displayed high accuracy in classifying students "High Risk" and "Low Risk", while the recommendation system effectively aligned with the detected performance gaps.

Interpretation of Findings

The Decision Tree model provided interpretable results with visible decision rules highlighting key performance indicators like attendance and internal marks. The integration of AI-driven recommendations strengthened the practical value of the model, making it not only predictive but also prescriptive offering actionable steps to enhance learning outcomes.

Comparison with Other Models

In comparison with Logistic Regression, Support Vector Machine, and Random Forest models, the Decision Tree demonstrated an optimal trade-off between prediction accuracy, interpretability, and computational simplicity. After integrating the AI-driven recommendation feature, the framework evolved into an effective and unified system for academic performance monitoring and personalized student support.

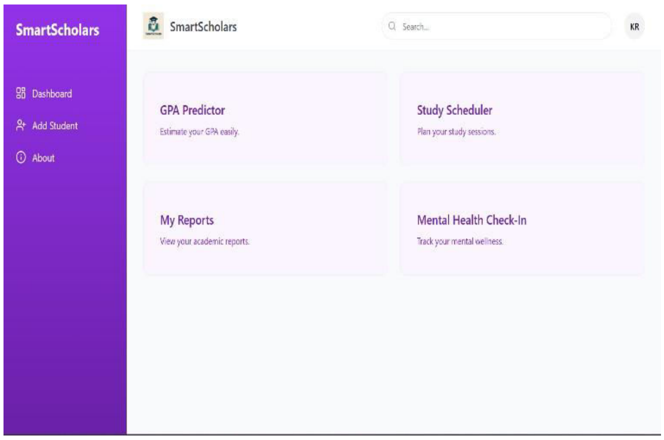

Fig .4. Dashboard

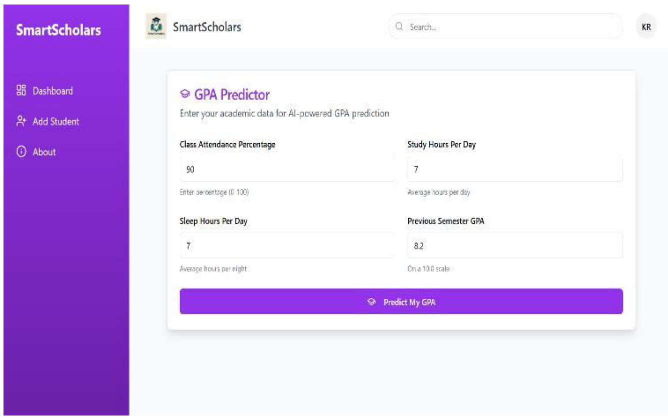

Fig .5. GPA Predictor Dashboard

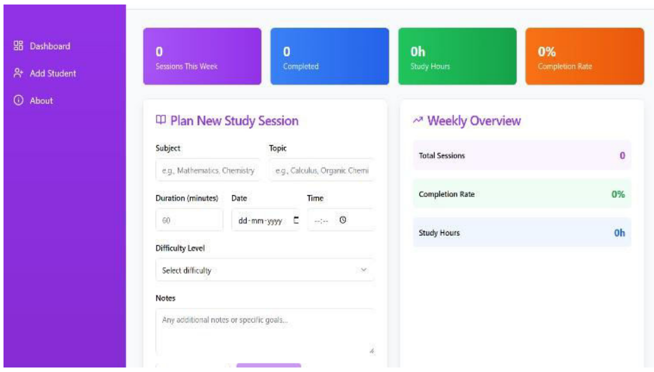

Fig.6. Study Scheduler Dashboard

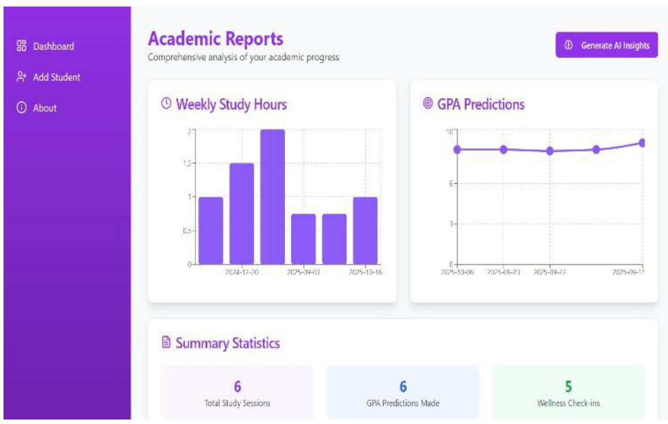

Fig.7. My Reports Dashboard

This research introduced an AI-based predictive framework designed to forecast student academic performance while offering personalized recommendations for improvement. Using a Decision Tree Classifier, the system achieved high accuracy and interpretability, allowing early identification of students who may require academic assistance. The inclusion of the AI recommendation module extended the framework's purpose beyond prediction, enabling data-driven feedback that supports timely interventions.

Comprehensive data preprocessing, model optimization, and visualization techniques enhanced the reliability and transparency of the proposed system. The study contributes to the field of Educational Data Mining by demonstrating how artificial intelligence can transform raw academic information into actionable insights that aid educators in making informed decisions. Furthermore, the framework supports the vision of equitable and quality education by promoting data-informed learning environments and continuous student development.

VIII. FUTURE WORKS

- **Integration with Academic Platforms:** Link the system with Learning Management Systems (LMS) for automatic data updates.
- **Adaptive AI Learning:** Enable the AI model to self-improve as more student data becomes available, enhancing prediction accuracy.
- **Expanded Recommendation System:** Include personalized learning materials and motivational support based on student behavior patterns.
- **Visualization Dashboard:** Develop advanced visual dashboards for tracking class-level and individual progress trends.
- **Data Privacy Enhancements:** Strengthen data encryption and compliance with ethical AI standards to ensure responsible use.

IX. REFERENCES

1. Aldraiweesh and U. Alturki, "The Influence of Social Support Theory on AI Acceptance: Examining Educational Support and Perceived Usefulness Using SEM Analysis," *IEEE Access*, Jan. 24, 2025.

2. W. Zou, W. Zhong, J. Du, and L. Yuan, "Multi-Model Fusion Framework for Academic Prediction," *Applied Sciences (MDPI)*, vol. 15, no. 7, p. 3550,Mar. 24, 2025. Available: https://www.mdpi.com/2076-3417/15/7/3550

3. G. R. Sandeepa and S. Mohottala, "ML Model Comparison for Academic Performance," *arXiv preprint* arXiv:2506.08047, Jun. 8, 2025. Available: https://arxiv.org/abs/2506.08047

4. Ali and F. Rahman, "Comparative Analysis of Machine Learning Classifiers for Academic Risk Detection," *International Journal of Artificial Intelligence in Education*, 2023.

5. L. Fernandez, D. Peters, and S. James, "Predictive Analytics for Student Retention in Higher Education," *Computers & Education*, 2024.

6. P. Patel and R. Mehta, "Impact of Attendance and Engagement on Academic Outcomes," *Education and Information Technologies*, 2023.

7. J. Zhou, C. Yang, and M. Chen, "Ensemble Methods Improving Prediction Accuracy in Educational Datasets," *IEEE Access*, 2024.

8. R. Singh, T. Kaur, and V. Bansal, "Deep Learning Approaches for Student Performance Classification," *Applied Intelligence*, 2025.

9. M. Hassan, F. Khalid, and A. Omar, "Early Alert Systems Integrated with Academic Advising Platforms," *Journal of Educational Data Mining*, 2024.

10. E. F. Agyemang, "Predicting Students' Academic Performance via Machine Learning Algorithms: An Empirical Review and Practical Application," *Computers and Education: Artificial Intelligence, 2024.*

11. A. Kala, "Early Prediction of Student Performance in Face-to-Face Education Environments," *IEEE Conference Proceedings*, 2024.

12. E. Alhazmi and A. Sheneamer, "Early Predicting of Students Performance in Higher Education," *IEEE Access*, Jan. 2023. doi: 10.1109/ACCESS.2023.3250702

13. A. Kumar, N. Bhatia, and P. Reddy, "Educational Data Mining and Early Warning Systems: Machine Learning for Student Risk Detection," *IEEE Access*, 2024.

14. Y. H. Chang, "Developing an Early Warning System with Personalized Interventions," *Education Sciences (MDPI)*, 2025.

15. M. H. Saeed and D. Wang, "Student Performance Early Warning Mechanism Based on Hierarchical Modeling," *ACM/IEEE Conference Paper*, Jul. 2025.

16. T. M. Nguyen, "Advancing Educational Data Mining for Enhanced Student Performance Prediction: Fusion of Feature Selection and Classification," *ResearchGate preprint*, 2025.

17. P. T. Le, J. M. Vo, and C. Nguyen, "Hybrid Machine Learning for Academic Success Prediction," *IEEE Access*, Mar. 2024.

18. R. K. Gupta and A. Jain, "A Comparative Study of Machine Learning Algorithms for Predicting Student Academic Performance," *International Journal of Emerging Technologies in Learning (iJET)*, 2024.

19. K. R. Alharbi and S. O. Alghamdi, "AI-Based Framework for Predicting Student Achievement and Dropout Risk," *IEEE Access*, 2025.

20. F. M. Ogunleye and L. Chen, "Explainable Machine Learning for Student Academic Performance Prediction," *Computers and Education: Artificial Intelligence*, 2024.

21. R. B. Lee, M. T. Chong, and S. Tan, "Integration of Psychological and Behavioral Features in Academic Performance Prediction," *IEEE Transactions on Learning Technologies*, 2023.

22. H. Park and J. Kim, "Deep Learning-Based Student Performance Forecasting in Higher Education," *IEEE Access*, 2024.

23. S. W. Tan and R. Prasad, "Optimization of Educational Data Models Using Hybrid Deep Neural Networks," *IEEE Access*, 2025.

24. A. Das and P. Singh, "Evaluating Transformer Models for Academic Performance Analysis," *Applied Computing and Informatics*, 2025.