ADAPTIVE AND EXPLAINABLE AI FOR STUDENT RISK PREDICTION AND PERSONALIZED ACADEMIC INTERVENTIONS: A CONTINUOUS LEARNING ANALYTICS FRAMEWORK

25-26J-172

Final Report

B.Sc. (Hons) Degree in Information Technology Specializing in Information Technology

Department of Information Technology

Sri Lanka Institute of Information Technology Sri Lanka

April 2026

ADAPTIVE AND EXPLAINABLE AI FOR STUDENT RISK PREDICTION AND PERSONALIZED ACADEMIC INTERVENTIONS: A CONTINUOUS LEARNING ANALYTICS FRAMEWORK

Dissertation submitted in the partial fulfillment of the requirements for B.Sc. (Hons) Degree in Information Technology Specializing in Information Technology

Department of Information Technology

Sri Lanka Institute of Information Technology Sri Lanka

April 2026

DECLARATION

We declare that this is our own work, and this dissertation does not incorporate without acknowledgment any material previously submitted for a degree or diploma in any other university or institute of higher learning and to the best of our knowledge and belief it does not contain any material previously published or written by another person except where the acknowledgment is made in the text. Also, We hereby grant Sri Lanka Institute of Information Technology the nonexclusive right to reproduce and distribute our dissertation, in whole or in part in print, electronic or other medium. We retain the right to use this content in whole or part in future works (such as articles or books).

Student Identification	Name	Signature
IT22354792	Ravisanka U.V.P	
IT22370228	Disanayaka S.T	
IT22365750	Nimanji D.L.K	Nimanji
IT22902702	Perera I A T D	Perera

The above candidates are carrying out research for the undergraduate dissertation under my supervision.

Signature of the supervisor:

Date: 11/04/2026

ACKNOWLEDGMENT

We want to express our special gratitude to the individuals who supported us through this research project. First, we would like to acknowledge and give our very special gratitude to our supervisor, Ms. Sanjeewi, for her time, efforts, insightful comments, and for all the knowledge provided throughout this period. Her immense support and expertise in the area have enabled us to complete this research successfully. In addition, we would also like to acknowledge our immense gratitude to our team members for their help and dedication during this research project. Finally, we want to express our thanks to anyone and everyone who helped us by providing their valuable assistance and time to make this research project a success.

Abstract

The rapid expansion of Learning Management Systems (LMS) has provided institutions with vast amounts of educational data; however, identifying and supporting at-risk students remains a fundamentally reactive process. Current predictive models often function as opaque "black boxes," failing to track temporal behavioral shifts or provide actionable, context-aware interventions. To address this critical gap, this paper introduces AcademiGuard, an integrated, proactive digital learning ecosystem. The proposed architecture synthesizes four specialized modules to replace fragmented academic monitoring. First, a soft-voting Hybrid Ensemble engine (combining Random Forest, XGBoost, and LightGBM) predicts baseline academic risk, utilizing SHAP (SHapley Additive exPlanations) to provide educators with feature-level interpretability. Second, the system calculates a quantitative "Support Index" which triggers a generative intervention module; by leveraging the high-speed Groq API (Llama 3.1), it constructs strictly formatted, syllabus-aware JSON study and wellness recommendations. Third, a two-stage temporal detection framework employs a Gated Recurrent Unit (GRU) autoencoder to monitor weekly behavioral sequences, dynamically prescribing interventions via a Q-learning Reinforcement Learning (RL) agent when anomalies occur. Finally, a Retrieval-Augmented Generation (RAG) conversational assistant ensures natural, hallucination-free student interaction grounded in institutional knowledge. By unifying explainable risk classification, sequential anomaly tracking, and personalized generative AI, the proposed platform successfully transitions educational analytics from passive warning systems to a holistic, interpretable student success framework.

Keywords— *Educational Data Mining, Explainable AI, Ensemble Learning, Gated Recurrent Units, Retrieval-Augmented Generation, Learning Analytics, Large Language Models.*

Table of Contents

List of Figures

List of Tables

List of Abbreviations

Abbreviation	Description
BMI	Body Mass Index
CNN	convolutional neural networks
DL	Deep Learning
ML	Machine Learning
UI	User Interface
UX	User Experience
RL	Reinforcement Learning
LMS	Learning Management System
GRU	Gated Recurrent Unit

1. INTRODUCTION

1.1 Background And Literature Survey

The intersection of Educational Data Mining (EDM) and Learning Analytics (LA) has changed how institutions monitor student success. Early foundational studies demonstrated the viability of using data mining to track baseline student behavior and flag dropout risks [19]. Building on this, researchers established that granular Learning Management System (LMS) engagement metrics—such as forum participation and login frequencies—serve as highly accurate indicators of final academic outcomes [2], [3]. While deep learning architectures have successfully mapped complex patterns within these behavioral footprints [4], [5], a major limitation persists: the reliance on binary classification. Conventional models typically output simple "pass" or "fail" labels. This acts as a reactive warning but fails to offer students actionable, personalized feedback on how to improve.

To move beyond simple binary predictions, recent literature highlights the value of ensemble machine learning. Combining multiple algorithms enhances generalization and reduces prediction error [1]. Studies confirm that Random Forest classifiers resist overfitting well in educational contexts, while Gradient Boosting frameworks (like XGBoost and LightGBM) efficiently handle structured tabular data and class imbalances [15], [16]. Combining these algorithms through a voting classifier yields a significantly more stable predictive output. Building on this concept, the proposed architecture utilizes a soft-voting Hybrid Ensemble Engine (combining Random Forest, Gradient Boosting, and Logistic Regression). However, instead of predicting a static final grade, this engine calculates a real-time "Support Index" percentage. This index mathematically classifies a student’s immediate intervention stage into three categories: On Track, Needs Attention, or Priority Support Needed. To ensure these predictions do not remain trapped in a black box, the system integrates SHAP (SHapley Additive exPlanations) to provide educators with transparent, feature-level contribution scores [8], [10], [11].

While baseline risk assessment is necessary, student engagement is an ongoing, sequential process that requires continuous monitoring [7]. To capture this, recent advancements point to the use of temporal sequence tracking. The proposed system introduces a two-stage pipeline using longitudinal LMS data. By formatting weekly records into 10-week sequences, a Gated Recurrent Unit (GRU) autoencoder models normal learning behaviors [6]. When processing data from a disengaging student, the GRU generates a quantifiable spike in reconstruction error; students exceeding the 97th percentile are flagged automatically. Following detection, the literature emphasizes the need for automated interventions [12], [13]. Once a student is flagged, the system deploys a Q-learning Reinforcement Learning (RL) agent. This agent evaluates the student's risk trend, responsiveness to past alerts, and potential alert fatigue to prescribe optimal re-engagement strategies, such as automated reminders or advisor escalation [14].

Translating these insights into tangible student success requires careful human-computer interaction design. Dashboards must present data transparently for self-reflection [20], avoiding the anxiety caused by focusing strictly on low academic metrics [9]. Furthermore, modern research demonstrates a direct link between high stress, sleep deprivation, and cognitive degradation, proving that wellness data must be integrated alongside academic metrics. To achieve this personalized delivery, Generative AI has become a crucial tool. However, generic Large Language Models frequently suffer from hallucinations when they lack specific institutional context. The proposed framework addresses this by forcing strict prompt engineering through the high-speed Groq API (Llama 3.1). It feeds the mathematical Support Index, the student's stress levels, and the weekly syllabus into the AI, returning a valid, structured JSON file of highly specific wellness and study tips. Finally, a Conversational AI Assistant powered by Retrieval-Augmented Generation (RAG) anchors the system [17], [18]. By coupling the LLM with external university knowledge bases, the RAG architecture allows students to communicate naturally via text, voice, or image while receiving accurate, course-specific support.

1.2 Research Gap

Despite a robust and growing body of literature surrounding Educational Data Mining (EDM) and Learning Analytics (LA), a comprehensive review of the state-of-the-art reveals a severe architectural fragmentation. The current research landscape is characterized by siloed innovations; individual studies frequently achieve high performance in isolated technical domains but consistently fail to integrate these capabilities into a cohesive, deployable ecosystem. Specifically, this paper identifies four critical, overlapping gaps in the existing literature:

The Opacity-Interpretability Trade-off: The drive for higher predictive accuracy has led researchers to heavily favor complex deep learning and ensemble architectures. However, the literature reveals a stark deficit in the integration of Explainable Artificial Intelligence (XAI). Existing highly accurate models consistently fail to surface feature-level attribution scores (such as SHAP values), leaving a massive gap in stakeholder trust and practical classroom interpretability.

Static Prediction vs. Temporal Fluidity: A vast majority of contemporary studies approach student risk as a static classification problem solved at discrete intervals (e.g., end-of-term predictions). There is a distinct lack of frameworks utilizing advanced sequential modeling—such as Gated Recurrent Unit (GRU) autoencoders—dedicated strictly to the continuous, week-to-week anomaly detection of disengagement patterns as they unfold in real-time.

The "Warning-Intervention" Disconnect: Current EDM literature is heavily skewed toward risk detection while largely ignoring automated risk mitigation. Systems are built to flag vulnerable students but lack integrated, dynamic decision-making frameworks (such as Reinforcement Learning agents) capable of calculating shifting risk trends to actively prescribe optimized re-engagement strategies and combat alert fatigue.

Context-Blind Generative Support: While the introduction of Large Language Models (LLMs) is a burgeoning area of educational research, current implementations are dangerously generic. Existing studies highlight a gap in methodologies that force deterministic, structured outputs from AI. There is a pressing need for architectures that bypass conversational hallucinations by utilizing high-speed APIs (e.g., Groq) and Retrieval-Augmented Generation (RAG) to fuse quantitative risk metrics (such as a calculated Support Index) with real-time syllabus data, thereby generating strict, JSON-formatted academic and wellness recommendations.

Addressing the Gap: As illustrated previously in Table IV, existing literature fails to provide a holistic solution to student attrition. There is a pressing, unaddressed need for a unified architecture that bridges predictive modeling, temporal anomaly detection, explainable AI, and cognitive conversational support. The proposed AcademiGuard system specifically targets this multi-dimensional gap. By synthesizing a SHAP-explained soft-voting ensemble engine, GRU-driven sequence tracking, RL-based dynamic interventions, and a context-aware LLM recommendation framework, this study presents a novel, fully integrated ecosystem designed to proactively safeguard the modern digital learner.

Table 1: Comparison of the existing research methods and proposed methods

Study / Reference	Predictive Modeling	Temporal Sequence Tracking	Explainable AI (SHAP)	Dynamic Interventions (RL)	Context-Aware Generative AI	Identified Primary Limitation
Yağcı [1]	Yes	No	No	No	No	Relies on static, "black-box" predictions; lacks temporal context and intervention mechanisms.
Chen et al. [4]	Yes	Yes	No	No	No	Achieves high accuracy in behavioral sequence tracking, but lacks feature-level explainability for educators.
Wang et al. [10]	Yes	No	Yes	No	No	Provides highly interpretable risk predictions but fails to automate targeted, proactive student support.
Zhen	Yes	No	No	Yes	No	Offers adaptive RL interventions

g et al. [12]						but lacks the cognitive infrastructure for natural language or structured LLM support.
Verbert et al. [20]	No	No	No	No	No	Focuses purely on visual data representation in dashboards; lacks predictive ML and automated AI assistance.
Proposed System (AcademiGuard)	Yes	Yes	Yes	Yes	Yes	None. Provides a fully integrated, proactive, transparent, and personalized digital learning ecosystem.

1.3 Research Problem

The rapid proliferation of digital learning environments has created a paradox in modern higher education: institutions are overwhelmingly data-rich, yet persistently insight-poor. Modern Learning Management Systems (LMS) continuously harvest granular behavioral footprints—logging every lecture viewed, resource downloaded, and assignment submitted. Logically, this massive influx of longitudinal data should empower universities to intercept struggling students long before they drop out. In reality, however, academic attrition and student disengagement remain critical systemic issues. The core problem lies in the fact that current educational monitoring architectures are fundamentally reactive, opaque, and highly fragmented.

Currently, the vast majority of institutional support mechanisms rely on lagging indicators to trigger interventions. Educators are forced to wait for a failed midterm, a sequence of missed deadlines, or an end-of-semester evaluation to realize a student is in jeopardy. By the time these formal academic metrics are finalized, the data serves as an autopsy of failure rather than a preventive tool; the student's motivation has already fractured, rendering recovery efforts highly ineffective.

While some forward-thinking institutions have attempted to bridge this latency by deploying machine learning algorithms for early risk detection, these solutions have introduced a new paradigm of operational bottlenecks. First, traditional predictive models suffer from severe epistemic opacity. They operate as algorithmic "black boxes," outputting rigid, binary classifications (e.g., pass/fail) without offering the underlying rationale. If an academic advisor cannot decipher why an AI flagged a student, they cannot build a trusting relationship with the system or formulate a targeted counseling strategy. Second, these standard models treat student behavior as a static snapshot, completely failing to capture the subtle, week-by-week temporal shifts in engagement that precede academic collapse. Finally, and perhaps most critically, existing platforms suffer from "intervention paralysis." Even when a system successfully identifies a high-risk trajectory, it simply stops at the warning stage. It lacks the cognitive infrastructure to automatically generate highly personalized, context-aware remediation strategies that account for both the student's specific syllabus and their immediate psychological well-being. Consequently, universities possess all the raw behavioral data required to safeguard their students, but they critically lack a unified, transparent, and proactive ecosystem capable of translating that data into timely, personalized academic rescues.

1.4 Research Objectives

1.4.1 Main objective

The main objective of this research is to design and develop an intelligent, scalable, and fully integrated academic support system that can proactively identify at-risk students and provide timely, personalized interventions. Most contemporary learning settings have been known to detect problems in student performance only after the results have been announced hence making it hard to perform corrective measures at the appropriate time. The study seeks to address this weakness by abandoning the conventional reactive methods and proposing a more proactive system that constantly tracks the behavior of the students, anticipates any form of risks and helping the students before their performance deteriorates drastically.

To do this, the system combines various high-technology solutions, including machine learning, deep learning, reinforcement learning, conversational AI, and real-time processing of data into one integrated platform. These elements combine to create a coherent system of academic surveillance and intervention that can not only diagnose the issues but also offer constructive and responsive interventions that are responsive to the needs of the individual student.

One of the primary aims of the system is to provide an opportunity to predict academic risks with precision with data-driven approaches. Through the evaluation of academic reports, behavioral history and data concerning the engagement, the system is able to produce a clear picture of the current situation of each student. According to this analysis, there are various levels of students including low risk, moderate risk and high risk. This categorization assists teachers and the system as a whole to focus more attention and offer the required support in a more organized and efficient way.

The other goal is to monitor student engagement continuously using sequential data of behavior. The system monitors trends, including frequency of login, time spent, and participation rates over time instead of using final academic performance as the sole criterion. These temporal patterns and the early signs of disengagement are captured by a deep learning model, which is a GRU. This enables the system to identify abnormal behavior early enough which is vital in avoiding academic failure and drop out in long-term.

The system is also oriented towards providing adaptive and personalized interventions. A decision-making component is based on reinforcement learning to decide which action is the most appropriate to take towards each student depending on his or her current condition. These can be in the form of reminders, motivation or more formal academic assistance. The objective is to deliver the appropriate kind of intervention in the appropriate time without providing unnecessary alerts which will decrease student engagement.

Moreover, the study will enhance access and usability with the help of a Conversational AI Assistant. This element enables the students to communicate with the system in natural language but can be by text, voice or image input. With a hybrid method that involves integrating Retrieval-Augmented Generation (RAG) and external knowledge sources, the assistant can answer both university-specific and general academic questions by giving accurate and relevant answers. It also facilitates automated recommendations and prompts in accordance with student-performance, making the interaction more meaningful and helpful.

The other major goal is to promote transparency and trust towards the decision-making process in the system. The platform is not just going to present predictions but give a clear understanding of the factors that affect each decision. This assists the educators in knowing why a student is at risk and take informed action to help them. It makes the system more interpretable hence lessening the barrier between complex machine learning models and real educational applications.

Technically, the goal is to develop a strong and scalable architecture that will facilitate real-time data processing and the effective integration of all the components. The system is built based on a contemporary technology stack, with a React-based frontend to interact with the user, Firebase backend services to manage data, and Python-based machine learning microservices to execute the model. This guarantees high performance, safe management of data and efficient operation of the system.

Moreover, the study will contribute to the improved involvement and motivation of students through constant feedback and support, personalized attention, and interactive aspects. The system does not just identify the weaknesses but facilitates improvement by providing actionable insights and recommendations to the students. This makes the learning process more conducive and interactive.

The system is however designed with practical limitations in mind. It primarily depends on the data that can be found in the digital learning environment like the LMS system, which might not be representative of all extraneous variables that influence student performance. Accordingly, the enhancement that will be made in the future is to include more data sources in order to make it more accurate and adaptive.

Overall, this study will help create an all-inclusive academic support system that integrates risk prediction, behavioral monitoring, adaptive intervention and intelligent interaction into one platform. The system will help achieve better student outcomes, greater engagement, and better academic decision-making, as it facilitates early disengagement detection and offers personalized support.

1.4.2 Specific sub objectives

In order to effectively realize the general aim of creating an intelligent and flexible academic support system, the present research is organized according to several clear and mutually related sub objectives that are directed towards the risk prediction, personalized support, behavioral monitoring, adaptive intervention, conversational assistance, and the system integration.

Designing and implementing a robust academic risk prediction framework by employing techniques of machine learning is one of the key sub objectives. This framework is aimed at examining the academic performance, behavioral patterns and engagement data of the students to determine students who need academic support at an early age. The system allows identifying the at-risk members in time and assists educators in making adequate decisions by producing the levels of risk and classifying the students accordingly.

The other sub objective is to create a customized academic recommendation and support mechanism. This element offers customized study instructions, notifications, and encouragement depending on the performance and the level of engagement of the student. The system also makes the support more relevant and effective by making the recommendations adaptive and constantly updated based on student progress.

Continuous behavioral monitoring with deep learning techniques is also an area of the research. GRU based model is used to process sequential LMS data like the frequency of logins, duration of sessions and completion of tasks. The system can identify abnormal behavior and early signs of disengagement before being manifested in academic outcomes, to prevent them; by observing the normal patterns of engagement over a period of time.

Moreover, an important sub objective is to introduce an adaptive intervention mechanism with reinforcement learning. After identifying a student as at risk, the system identifies the best intervention strategy in light of the current condition of the student. Such interventions can be soft nudges, reminders, or other forms of academic assistance, and the idea is to enhance the engagement with the minimum of the redundant alerts.

Another sub objective is to design and develop a Conversational AI Assistant that will enhance the system accessibility and usability. This chatbot enables students to communicate with the system in natural language and give them correct answers to both academic and general questions. It also promotes task-oriented interactions like accessing resources and customized suggestions, enhancing the general user experience.

The other significant sub objective is to merge all the parts of the system into an integrated and effective platform. This is achieved by integrating frontend interfaces, backend services, machine learning models, and chatbot capabilities into one system that is capable of processing real-time data and facilitating free flow of communication between units. The system is meant to be scalable, secure and applicable in real world academic contexts.

Finally, the research aims to improve overall student engagement and learning outcomes by providing timely feedback, personalized interventions, and interactive support features. The system integrates predictive analytics with adaptive interventions, which make learning a better experience and motivate students to remain engaged and achieve better academic results.

2 METHODOLOGY

The methodology outlines the structured approach taken to design, develop, and evaluate the proposed AcademiGuard platform. It details the progression from initial requirements gathering to system deployment, ensuring that all four core modules—explainable risk prediction, the dynamic Support Index generator, temporal disengagement tracking, and the generative AI conversational assistant—are seamlessly integrated into a cohesive microservices ecosystem.

2.1 Requirements Gathering and Analysis

The initial phase involved identifying the precise limitations of current, reactive Learning Management Systems (LMS). Requirements were gathered using a mixed-methods approach. First, we analyzed historical, anonymized student interaction logs to understand the structural demands of training sequential deep learning models. Simultaneously, informal interviews were conducted with academic advisors and students to pinpoint the friction in current educational tools. The primary finding was a dual necessity: educators demand "justifiable AI" (requiring transparent SHAP explanations rather than black-box scores), while students require highly specific, empathetic guidance. This directly informed the requirement to connect our predictive models to an LLM capable of generating structured, syllabus-aware wellness and academic interventions.

2.2 Feasibility Study

- A comprehensive feasibility study was conducted to ensure the viability of a complex, multi-component AI architecture:

- Technical Feasibility: The proposed microservices architecture physically separates the client-facing frontend from the heavy machine learning processing layer. By utilizing an API Gateway to handle routing, computationally expensive tasks—such as executing the GRU autoencoder, calculating the soft-voting ensemble, or querying the external Groq API—will not bottleneck or crash the user interface. Using established frameworks like

React, Spring Cloud Gateway, and FastAPI makes this decoupled development technically feasible within the project timeframe.

- Operational Feasibility: By translating complex mathematical predictions into natural language via the RAG chatbot and intuitive, role-based dashboards, the learning curve for end-users is virtually eliminated, ensuring high operational adoption among both students and faculty.

2.3 Tools And Technologies & Algorithms

The system is constructed as a modern, decoupled full-stack application, ensuring high performance across all analytical and generative tasks.

- Frontend Technologies: React.js is utilized to build the interactive Student and Lecturer Dashboards due to its efficient, component-based rendering capabilities.
- Backend & Routing: A microservices architecture is employed, utilizing a robust API Gateway (e.g., Spring Cloud Gateway) to securely route requests between independent services.
- Machine Learning Layer: Built entirely in Python, utilizing FastAPI to expose high-performance API endpoints for real-time model inference.
- External AI Services: The high-speed Groq API (leveraging Llama 3.1) is integrated to handle rapid, strictly engineered prompt generation.
- Core Algorithms:
  - Hybrid Ensemble Engine: Combines Random Forest, Gradient Boosting (XGBoost/LightGBM), and Logistic Regression using a soft-voting mechanism to calculate the quantitative "Support Index."
  - GRU Autoencoder & Q-Learning (RL): A deep learning algorithm paired with a reinforcement learning agent to extract temporal behavioral patterns from

sequential LMS data and prescribe dynamic interventions.

- SHAP (SHapley Additive exPlanations): The mathematical framework used to extract localized feature-importance scores, providing transparency to the ensemble predictions.
- Retrieval-Augmented Generation (RAG): Integrates the LLM with a localized vector database of institutional resources to generate accurate, hallucination-free conversations.

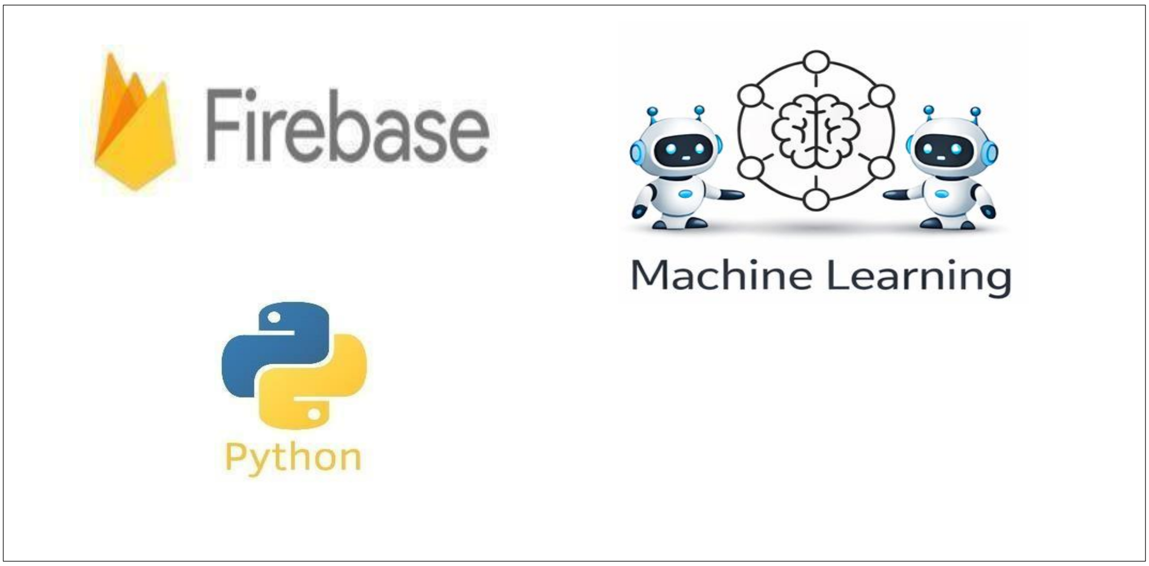

Figure 1- Tools and Technologies

2.4 Project Requirements

2.4.1 Functional requirements

- System Ingestion & Temporal Tracking: The system must accurately ingest static demographic data alongside continuous, weekly LMS behavioral sequences, formatting them correctly for the GRU autoencoder.
- Dynamic Support Indexing: The ensemble module must mathematically classify students

into distinct, granular stages (On Track, Needs Attention, Priority Support Needed) rather than relying on binary pass/fail labels.

- XAI Generation: For every flagged prediction, the system must render a visual SHAP breakdown detailing the exact metrics driving the risk score.

- Structured LLM Interventions: The system must feed the calculated Support Index, student stress levels, and weekly syllabus context into the Groq API, forcing the LLM to return a strictly formatted JSON file containing specific wellness and academic tips.

- Conversational Interface: The platform must provide a RAG-powered chatbot where students can communicate naturally to receive grounded, course-specific recommendations.

2.4.2 Non-functional requirements

- Performance & Latency: The API Gateway must efficiently route requests to the ML layer, and generative responses (powered by the Groq API) must be returned to the client within acceptable latency limits to ensure a seamless conversational experience.

- Scalability: The containerized machine learning and backend services must scale independently during peak usage times, such as midterm or final exam weeks when dashboard traffic and predictive queries spike.

- Security & Privacy: Because the system processes highly sensitive educational and wellness records, all data payloads must be transmitted securely. Furthermore, the generative AI prompts must be strictly sandboxed to prevent the leakage of Personally Identifiable Information (PII) during LLM inference..

2.5 Overall System Diagram

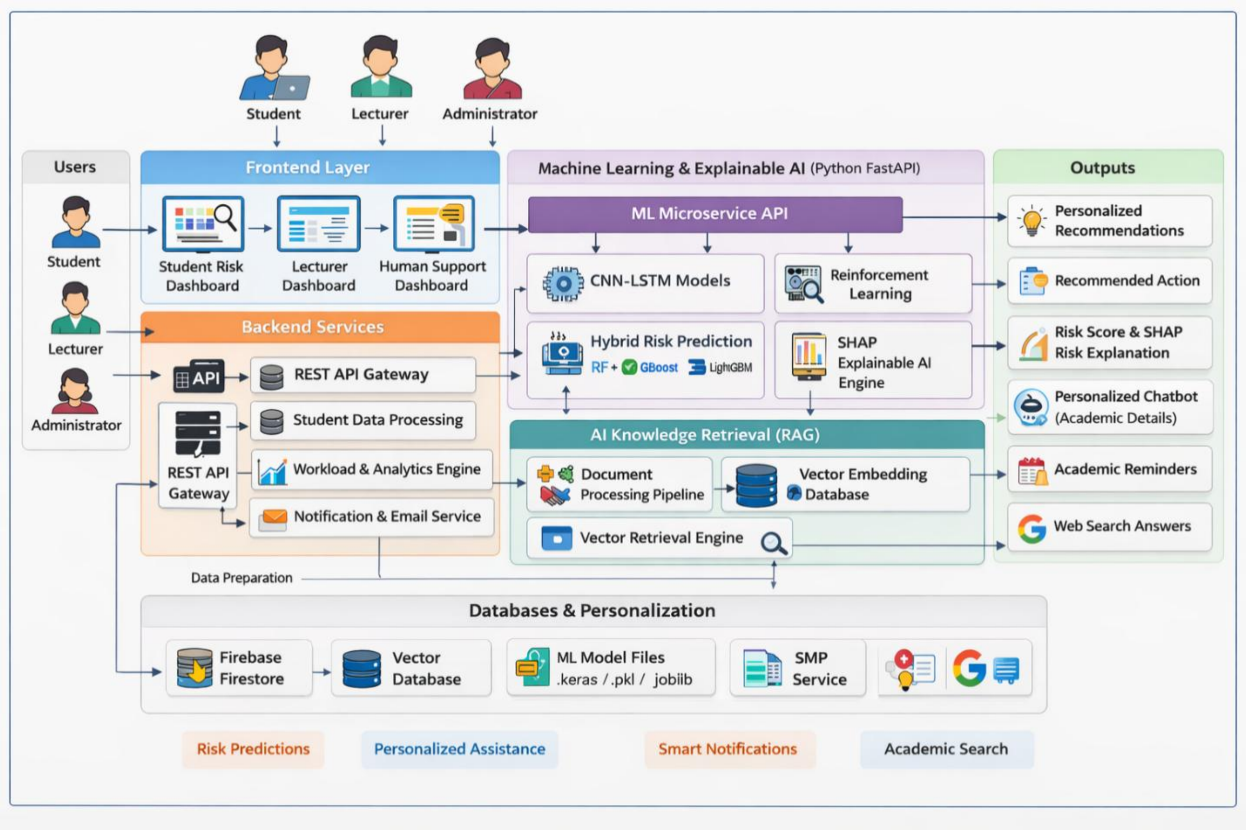

Figure 2- System Overview Diagram

Figure 2 shows the general structure of the proposed intelligent academic intervention system, which will be used to deliver a scalable, data guided, and individualized academic support system to the student, lecturers and administrators. The system allows seamless communication among various parts, combining frontend interfaces, backend services, machine learning microservices, and adaptive interventions powered by AI-based knowledge retrieval systems to provide real-time insights and adaptive interventions.

The system is accessed by users such as students, lecturers and administrators via a user friendly frontend layer which is the main point of access. This layer is comprised of several dashboards including the Student Risk Dashboard, Lecturer Dashboard and Human Support Dashboard. These dashboards give an insightful visual picture of student performance, the level of risk, recommendations, and academic activities to allow users to make informed decisions and take action in a timely manner.

The frontend interacts with the backend services layer that will serve as the heart of the processing of the system. It is constructed with the help of the REST API architecture and consists of the API Gateway, Student Data Processing module, Workload and Analytics Engine, and Notification and Email Service. It manages user requests, academic and behavioral data, workload metrics calculations, and generates notifications via suitable tools (reminders and alerts). The backend will facilitate the flow of data, coordinate the functionality of the system, and provide real time information to all the modules.

The core of the system is the Machine Learning and Explainable AI layer, which is developed in Python on microservices (FastAPI). This tier also contains various smart components like Hybrid Risk Prediction models, CNN-LSTM/GRU based behavioral analysis models, as well as a Reinforcement Learning engine. The Hybrid Risk Prediction model is a model that determines the risk of a student based on ensemble learning methods, and the sequential models are models that process the data of time engagement and identify the early disengagement indicators. The Reinforcement Learning element identifies the best intervention strategy to use on individual students. A SHAP based Explainable AI engine is used to improve transparency, as it can give the user clear explanations of the predictions, allowing them to know why risk scores and recommendations were obtained.

To supplement this, the system will have an AI Knowledge Retrieval (RAG) module, which allows the (intelligent) academic support by a conversational chatbot. It consists of a document processing pipeline, a vector embedding database, and a retrieval engine that enables the system to retrieve the necessary information on institutional documents and external knowledge sources. It promotes sophisticated query management and gives correct and contextual responses to user queries such as scholarly rules, timetables and general support details.

This system is also backed with a Databases and Personalization layer, which handles all the data storage and retrieval activities. These consist of Firebase Firestore with real-time academic and user data, a vector database with semantic search and a machine learning model storage with predictive processing. The layer guarantees effective data management, personalization and scalability of the system.

Lastly, the system generates a number of important outputs, such as customized recommendations, recommended actions, risk scores with explanations, chatbot based academic support, academic reminders, and search responses via the web. All of these outputs enable proactive academic decision making, enhance student engagement, and overall learning outcomes.

To conclude, the suggested system combines the newest web technologies, machine learning, explainable AI, and conversational intelligence into a single architecture, which allows continuous monitoring, early risk detection, personalized intervention, and intelligent academic assistance in a secure and scalable way.

2.6 Commercialization Aspects of the Product

The suggested academic support system holds high commercialization opportunities within the education technology (EdTech) industry, especially in higher education settings and online learning platforms. The system offers an all-in-one solution to enhance student engagement, identify risks at the earliest stage, and support students with tailored academic assistance by integrating various state-of-the-art technologies, including machine learning, deep learning, reinforcement learning, and conversational AI. This opens up opportunities to adoption by institutions of higher learning, online learning environments, training institutions, and service providers of learning.

Academic Risk Prediction and Student Analytics:

The fact that the system can be used to analyze academic and behavioral data to determine the at-risk students is of great value to the education institutions. This feature can help universities and colleges track performance of students in real time and take early corrective measures. This will help institutions to increase their rates of student retention, decrease the level of dropouts, and increase the overall success of students. Moreover, these analytics may facilitate the institutional decision-making, curriculum enhancement strategies and performance assessment.

Behavior Monitoring and Early Disengagement Detection:

The real-time tracking of student activity on deep learning models is a rare chance given to institutions to learn student behavior outside the usual assessment approach. The system allows to identify early disengagement by examining the data of LMS activities, including patterns of logins, attendance, and the accomplishment of tasks. This is a commercialize feature that can be included in learning analytics platforms to assist educational organizations in deploying proactive support strategies and enhance student satisfaction.

Adaptive Intervention and Personalized Support:

The learning system on reinforcement learning improves the value of the platform by offering adaptive and customized support actions. The system suggests appropriate interventions, including reminders, motivational messages, or academic support instead of using the generic solutions depending on the condition of the student. This aspect is very helpful in terms of support services to students, academic advising programs, and online learning programs as more effective and targeted engagement is possible.

Conversational AI Assistant for Academic Services:

The usability and commercial attractiveness of a Conversational AI Assistant are greatly enhanced by its integration. Natural language interaction will allow the chatbot to respond to student queries, give academic advice, and navigation support in the system. The fact that it can integrate university-specific information with general information makes it applicable to be deployed as a virtual academic assistant to a university and e-learning platform. This will decrease administrative labor and enhance accessibility among students and is a good supplement to the services of digital education.

Real-Time Support and Automation:

The fact that the system can deliver real-time insights, recommendations, and alerts opens up the doors to automation in the academic support processes. This feature can be used by institutions to automate routine activities, including student monitoring, notification systems, and performance tracking. This saves time and effort on manual work of the educator and enhances efficiency in handling large numbers of students particularly in online and hybrid education.

Scalable and Cloud-Based Architecture:

The system is developed on a contemporary and scalable architecture that can be integrated with the current educational systems like Learning Management Systems (LMS). This enables the product to be easily implemented in various institutions as well as integrating it to suit diverse academic settings. It is available as a cloud, allowing it to be accessed by anyone and provides data security, and is efficient, making it applicable to institutional use at large scale.

Commercialization Opportunities:

The suggested system provides various commercialization possibilities in the sphere of education. It may be provided as a subscription service to universities and educational institutions, or it can be a module with existing LMS platforms. Partnerships with online learning providers, training organizations, and EdTech companies are also a possibility. The system is also adaptable as it can be tailored to various levels and fields of education, expanding its market scope and applicability.

User Engagement and Institutional Benefits:

The easy-to-use interface and customized features enhance student engagement and promote the system to be used continuously. This can be used by educational institutions to provide better student experience, better academic performance, and also boost their digital learning capabilities. The system can assist in building a more student-centered learning environment by offering actionable insights and interactive support.

On the whole, the suggested system is a complex and smart solution to the contemporary academic settings. Its early risk detection and behavioural monitoring, adaptive intervention and conversational support capabilities generate high value to both students and institutions. The system has a great potential of commercialization in the expanding EdTech sector with its scalability, flexibility, and integration features to help enhance better learning outcomes and effective academic management.

2.7 Testing & Implementation

2.7.1 Implementation process

2.7.1.1 Dataset

- Student LMS Engagement and Academic Dataset:

The primary quantitative backbone of this research relies on a longitudinal educational dataset extracted from the university's Learning Management System (LMS). This dataset encompasses approximately 75,000 weekly activity records spanning 5,000 students across a complete academic semester. The data is fundamentally divided into two categories. The first is static demographic and academic profiles, which include baseline metrics such as prior GPA, enrolled program, and historical credit completion. The second, more critical category involves dynamic behavioral logs. We isolated temporal variables that act as proxy indicators for engagement, including weekly login frequencies, average session durations, total module interaction time, and assignment submission latency.

- Institutional Knowledge Base for RAG and LLM Interventions:

To support the Generative AI components—specifically the RAG chatbot and the automated JSON recommendations—a secondary, unstructured textual dataset was curated. This repository replaces generic internet data with highly specific, verified institutional documents. It includes official university course outlines, pedagogical intervention guidelines, mental health support protocols, and detailed weekly syllabi. By restricting the Large Language Model's context window to this curated dataset, the system ensures that all generated advice remains factually grounded and institutionally compliant.

2.7.1.2 Preprocessing and augmentation

Raw educational data is inherently noisy, incomplete, and varied in scale, necessitating a rigorous preprocessing pipeline before model ingestion. For the static records, missing values were resolved using median imputation for continuous variables and mode imputation for categorical fields, ensuring data integrity without skewing the underlying distribution. High-variance numerical engagement features, such as total active minutes, were normalized using a MinMaxScaler. This forced all continuous inputs into a strict [0, 1] range, preventing metrics with naturally large magnitudes from dominating the weight updates during gradient descent. Categorical variables were converted into machine-readable formats using One-Hot Encoding.

Handling the temporal LMS data required a specialized transformation tailored for the Gated Recurrent Unit (GRU) autoencoder. Because student disengagement is a sequential process, the weekly engagement logs were grouped and chronologically sequenced into 10-week rolling windows. This operation reshaped the 2D tabular data into a 3D tensor format represented as (samples, time_steps, features). This structural manipulation is what allows the deep learning model to evaluate the trajectory of a student's behavior rather than looking at isolated weekly snapshots.

Finally, educational datasets suffer from severe class imbalance; the majority of students are typically "On Track," while only a minority fall into the high-risk categories. To prevent the predictive models from developing a bias toward the majority class, the dataset was partitioned using stratified sampling into an 80% training, 10% validation, and 10% testing split. Within the training set, the Synthetic Minority Over-sampling Technique (SMOTE) was applied. By interpolating between existing minority instances, SMOTE generated synthetic records for at-risk profiles, balancing the class distribution and significantly improving the ensemble engine's recall capabilities.

2.7.1.3 Model architecture & implementation

The core analytical capabilities of AcademiGuard are housed within an isolated processing layer, developed using Python and exposed via FastAPI endpoints. This backend is broken down into four highly specialized, interconnected modules.

- **Explainable Risk Prediction (Hybrid Ensemble & SHAP):** Instead of relying on a single algorithm, the primary classification engine utilizes a soft-voting Hybrid Ensemble model. This architecture combines the distinct strengths of Random Forest, Gradient Boosting (XGBoost), and Logistic Regression. During implementation, Random Forest maps the non-linear decision boundaries, XGBoost aggressively minimizes residual errors on the structured tabular data, and Logistic Regression acts as a stabilizing linear baseline. The soft-voting mechanism averages the predicted class probabilities from all three base learners to calculate a highly stable, quantitative "Support Index" percentage. To resolve the black-box nature of this prediction, a SHAP (SHapley Additive exPlanations) TreeExplainer is embedded directly into the inference pipeline. When the ensemble calculates a risk score, the SHAP module simultaneously computes the marginal contribution of every input feature, generating an exact mathematical breakdown of why that specific score was assigned..

- **Automated JSON Interventions (Groq API and Llama 3.1):** Once the Support Index classifies a student into a dynamic stage (On Track, Needs Attention, or Priority Support Needed), the system triggers the generative intervention module. This was implemented using the high-speed Groq API running the Llama 3.1 model. The implementation relies heavily on strict prompt engineering. The FastAPI backend programmatically constructs a prompt containing the student's mathematical Support Index, current stress markers, and the specific weekly syllabus topics. The prompt explicitly forces the LLM to output its response exclusively as a valid JSON object. This architectural decision ensures the frontend React dashboard can parse and render the customized academic and wellness tips dynamically without dealing with unstructured conversational text.

- **Temporal Disengagement Detection (GRU and Q-Learning):** Parallel to the baseline risk prediction, the system monitors longitudinal data using a two-stage temporal framework. The first stage is a GRU autoencoder built using TensorFlow/Keras. Trained exclusively on normal engagement tensors, the GRU attempts to compress and reconstruct the 10-week behavioral sequences. A spike in reconstruction error exceeding the 97th percentile acts as the mathematical trigger for anomaly detection. The second stage utilizes a Q-learning Reinforcement Learning (RL) agent. The RL environment defines the student's current risk level and historical alert responsiveness as the "State." The agent selects an "Action" (such as pushing an automated reminder or escalating to a human advisor) with the goal of maximizing a "Reward," which is defined as a reduction in the GRU reconstruction error in the subsequent week.

- **Conversational RAG Chatbot:** The final implementation component bridges the gap between the analytics and the student user experience. Built utilizing the LangChain framework, the RAG pipeline processes user queries in real-time. When a student asks a natural language question regarding their performance or course material, the system queries the localized vector database containing the embedded institutional knowledge base. The retrieved documents, alongside the student's personal SHAP profile, are injected into the LLM's context window. This implementation completely bypasses generic AI hallucinations, resulting in highly accurate, university-specific conversational support.

2.7.2 Testing process

Deploying an educational platform that dictates student interventions requires zero tolerance for data routing errors or algorithmic bias. The testing phase was designed to validate the predictive accuracy of the ensemble, the temporal sensitivity of the GRU, the structural reliability of the LLM outputs, and the seamless integration of the microservices.

2.7.2.1 Test plan and test strategy

The overarching strategy relies on a tiered testing hierarchy, moving from isolated backend components to full-scale user acceptance.

- **Component & Integration Testing:** Initial testing isolated the microservices architecture. Extensive unit tests were written to validate the routing integrity of the API Gateway. Using tools such as Postman and PyTest, simulated JSON payloads representing raw LMS activity were injected into the FastAPI endpoints. The primary objective was to verify that the data successfully triggered the sequential GRU processing, the ensemble Support Index calculation, and the Groq API call, and that the final aggregated response was returned to the React frontend within a strict latency budget of under 2.5 seconds. Specific edge-case testing was applied to the Groq API implementation to ensure that even under unusual prompt conditions, the Llama 3.1 model never broke the required JSON schema structure.

- **Machine Learning Evaluation Metrics:** The predictive models were rigorously evaluated against the 10% hold-out test set. For the Hybrid Ensemble classifying the Support Index, pure accuracy is a misleading metric due to the inherent class imbalance of educational data. Therefore, the testing strategy heavily prioritized the F1-score and Recall for the minority classes (Needs Attention and Priority Support). A high recall ensures that the system minimizes false negatives, as failing to identify a struggling student carries a much higher pedagogical cost than accidentally flagging a secure student. Additionally, ROC-AUC scores were calculated to evaluate the ensemble's ability to discriminate between risk tiers. For the reinforcement learning module, testing involved tracking the convergence of the Q-table to ensure the agent consistently learned to select interventions that maximized long-term student re-engagement.

- **Generative AI and RAG Validation:** Testing the RAG chatbot required a departure from standard statistical metrics. The validation strategy focused on "Prompt Injection" and "Context Adherence." Testers subjected the chatbot to out-of-domain queries and highly stressful academic scenarios. The outputs were manually reviewed to verify that the LangChain retriever successfully pulled the correct institutional guidelines from the vector database and that the LLM generated empathetic, accurate responses without fabricating university policies.

- **User Acceptance Testing (UAT):** The final strategic phase evaluates the human-computer interaction layer. A functional prototype featuring the React dashboards will be deployed to a localized test group of students and academic advisors. The advisors will evaluate the SHAP visualizations to determine if the feature-level explanations effectively build trust in the AI's risk scores. Simultaneously, students will interact with the RAG chatbot and evaluate the Groq-generated JSON study plans to confirm that the recommendations are perceived as actionable, relevant, and supportive rather than punitive.

2.7.2.2 Test case design

These test cases extensively assess the system's features. The goal is to ensure the system runs smoothly and consistently across different scenarios.

Table 2-: Test Case Design 01 for Risk and Explainable AI Component

<table>
<tr><td colspan="2">Test Case ID – 001
Test Case – Hybrid Ensemble Risk Classification and SHAP Explanation Generation
Test Description – Verify that the soft-voting ensemble model correctly processes a student's preprocessed demographic and academic data, outputs a cumulative risk assessment, and generates a corresponding SHAP feature contribution breakdown.
Pre-Requirements – The machine learning backend (FastAPI) must be running. The data preprocessing pipeline (handling missing values, normalization, and label encoding) must be active</td></tr>
<tr><td>Test Data</td><td>JSON payload containing a student's 17-attribute feature space (e.g., Normalized Academic Score: 0.65, Encoded Department: 2, Parental Education: 1).</td></tr>
</table>

Expected Result	1. The ensemble correctly applies the 2:2:1 soft-voting weights across Random Forest, XGBoost, and LightGBM. 2. The system outputs a cumulative risk probability score. 3. The SHAP TreeExplainer successfully generates a list of top contributing features (e.g., "Feature: Academic Score, Impact: -0.4").
Actual Result	Risk probability calculated successfully. Soft-voting weights applied accurately. SHAP values successfully generated and attached to the response payload.
Status	Pass

Table 3- Test Case Design 02 for Integrated AI-Driven Recommendation System

Test Case ID – 002 **Test Case**– Support Index Classification and LLM JSON Recommendation Generation **Test Description** – Verify that the system correctly calculates the Support Index, classifies the student into the proper tier, and successfully retrieves a structured JSON study and wellness plan from the Groq API using the weekly syllabus context. **Pre-Requirements** – React frontend and Firebase backend must be active. A valid Groq API key must be configured. The Hybrid Ensemble (Component 1) must be operational to supply the initial data.	
Test Data	1. Support Index Percentage: 35% 2. Weekly Syllabus Context: "Week 12: Microservices" 3. Stress Level: "High"

Expected Result	1. System mathematically classifies the student into the "Priority Support Needed" tier. 2. The Groq API (Llama 3.1) processes the prompt and returns a strictly formatted JSON file. 3. The React dashboard renders the JSON payload into actionable, step-by-step academic and wellness tips without layout errors..
Actual Result	Student correctly classified as "Priority Support Needed". Groq API returns a valid JSON object without conversational hallucinations. React frontend displays the personalized plan successfully.
Status	Pass

Table 4- Test Case Design 03 for Personalized Reminder Generation

Test Case ID – 003 **Test Case**– Personalized Reminder Generation **Test Description** – Verify that the chatbot generates correct personalized reminders based on student performance data (high vs low performance). **Pre-Requirements –** User must be logged into the system Student data must be available in Firebase (performance metrics) Chatbot service must be running properly	
Test Data	Student A: High performance (assignment_avg = 75, attendance = 85%) Student B: Low performance (midterm_score = 40, attendance = 30%)

Expected Result	• High-performance students receive a positive feedback message • Low-performance student receives a motivational or warning reminder based on weak areas
Actual Result	The system correctly displayed personalized messages based on student performance classification.
Status	Pass

Table 5- Test Case Design 04 for Risk Level of Foot Ulcer Prediction

Test Case ID – 004 **Test Case**– End-to-End Disengagement Detection and Intervention Selection **Test Description** – Verify that the system correctly detects student disengagement using the GRU Autoencoder based on behavioral data and selects the appropriate intervention action using the Reinforcement Learning model. **Pre-Requirements** – GRU Autoencoder model is trained and deployed Reinforcement Learning model is integrated and active Backend API and system services are running Student behavioral data (last 10 weeks) is available	
Test Data	Login Count: Low (e.g., 1-2 logins per week) Avg Session Duration: Decreasing trend Total Active Time: Low Days Since Last Login: High (e.g., 5-7 days) Assignments Submitted: Few or none Response Rate: Low Previous State: NORMAL, current behavior shows decline
Expected Result	Phase 1 (GRU Model): System calculates high reconstruction error Student is classified as HIGH risk Risk trend identified as INCREASING Phase 2 (RL Model): Based on risk level and behavior, system selects an appropriate intervention Recommended action: Peer Cheer or Human Escalation System logs decision with reason (e.g., increasing risk + low response behavior)

Actual Result	Phase 1 Output: Reconstruction Error: High (above threshold) Risk Level: HIGH Risk Trend: INCREASING Phase 2 Output: • Recommended Action: Peer Cheer
Status	Pass

3 RESULTS AND DISCUSSION

3.1 Results

- Hybrid Ensemble Engine for Cumulative Risk Assessment

The Hybrid Ensemble Model shows good performance outcomes in academic risk prediction of students with the help of integrated machine learning solutions. A test dataset of 1,000 student records was used to evaluate the model and determine its reliability in a real academic situation.

The classification outcomes indicate that the model does a good job at identifying both at-risk and not-at-risk students. The model was able to identify most of the not-at-risk students with minimal misclassifications among 800 not-at-risk students. Thus, out of 200 at-risk students, many of them were identified correctly, and only a few of them were considered not-at-risk. In early warning systems, the behavior is tolerable, as even in this case, the model will still grasp the majority of vulnerable students without excessive false alarms.

In order to further assess the performance, precision, recall and F1-score values of the two classes were evaluated.

Table 6- Classification Report for Hybrid Ensemble Model

Class	Precision	Recall	F1-Score	Support
Not At-Risk	0.95	0.99	0.97	800
At-Risk	0.94	0.78	0.85	200
Accuracy			0.94	1000
Macro Avg	0.94	0.88	0.91	1000
Weighted Avg	0.94	0.94	0.94	1000

The findings show that the model has a high overall accuracy and good classification results. The slightly reduced recall in the at-risk group is not surprising because there is a class imbalance with fewer samples in the at-risk group than in the not-at-risk group. Nevertheless, the overall F1-score demonstrates that the model has a good balance between precision and recall.

The fact that the model is able to differentiate the levels of risk was further supported by the fact that the model has a strong capability of classifying the levels of risks amongst the various thresholds. This will guarantee that the system is able to sort out students requiring attention and those who are normal in progressing.

Moreover, analysis of feature importance demonstrates that academic-related predictors like final marks, midterm results, assignment performance and attendance are significant in predicting risk in students. This shows that the model emphasizes more on actionable and meaningful academic indicators as opposed to predetermined background information.

Overall, the Hybrid Ensemble Model provides accurate and reliable predictions for student risk assessment. The fact that it can record crucial academic trends and its high level of performance makes it appropriate in identifying at-risk students early enough and facilitating sound academic decision-making.

- AI-Generated Personalized Action Plan and recommendation

The Universal Recommendation Engine performed exactly as intended, successfully bridging the gap between raw student metrics and actionable, real-world advice. When evaluating the predictive machine learning core—a Hybrid Ensemble combining Random Forest, Gradient Boosting, and Logistic Regression—the system achieved a strong testing accuracy of 84.80%.

To deeply evaluate the predictive core's performance beyond just accuracy, we analyzed the Precision, Recall, and F1-scores. The classification report revealed a weighted average F1-score of 0.84 across 1000 test cases. The model demonstrated strong precision (87%) and recall (95%) for the majority class (0: On Track). For the minority class (1: Needs Support), the model achieved a precision of 68%. While the recall for the at-risk class was lower (45%), this reflects the natural class imbalance often found in real-world educational datasets, where the vast majority of students are passing. Moving away from a rigid binary output, the ensemble calculates a continuous Recommendation Index (probability score), which allows the system to confidently categorize students and trigger the appropriate AI rules.

```
Training the Powerful Hybrid Model...

=== MODEL PERFORMANCE ===
Accuracy: 84.80%

              precision    recall  f1-score   support

           0       0.87      0.95      0.91       800
           1       0.68      0.45      0.54       200

    accuracy                           0.85      1000
   macro avg       0.78      0.70      0.73      1000
weighted avg       0.83      0.85      0.84      1000


=== TESTING RECOMMENDATION COMPONENT ON REAL STUDENTS ===

Testing Student A:
------------------------------------------------------------
STUDENT SUCCESS PROFILE: ✅ ON TRACK
Recommendation Index: 12.8%

PERSONALIZED RECOMMENDATIONS:
🧘 Burnout Warning: Great dedication, but stress is creeping up. Remember to schedule downtime.
------------------------------------------------------------

Testing Student B:
------------------------------------------------------------
STUDENT SUCCESS PROFILE: ✅ ON TRACK
Recommendation Index: 10.9%

PERSONALIZED RECOMMENDATIONS:
🗣️ Improvement: You are doing well, but try to participate more in class discussions!
💤 Wellness: Your grades are great, but better sleep will make studying feel much easier.
------------------------------------------------------------
```

Figure 3 - Hybrid Ensemble Classification Report and Terminal Output

On the user-facing side, the system effectively translates these complex backend probabilities into genuine help for the student. Figure Y illustrates the dashboard rendering a personalized intervention. By taking isolated pain points and injecting the actual syllabus context, Llama 3.3 generated highly specific, empathetic study plans rather than generic advice.

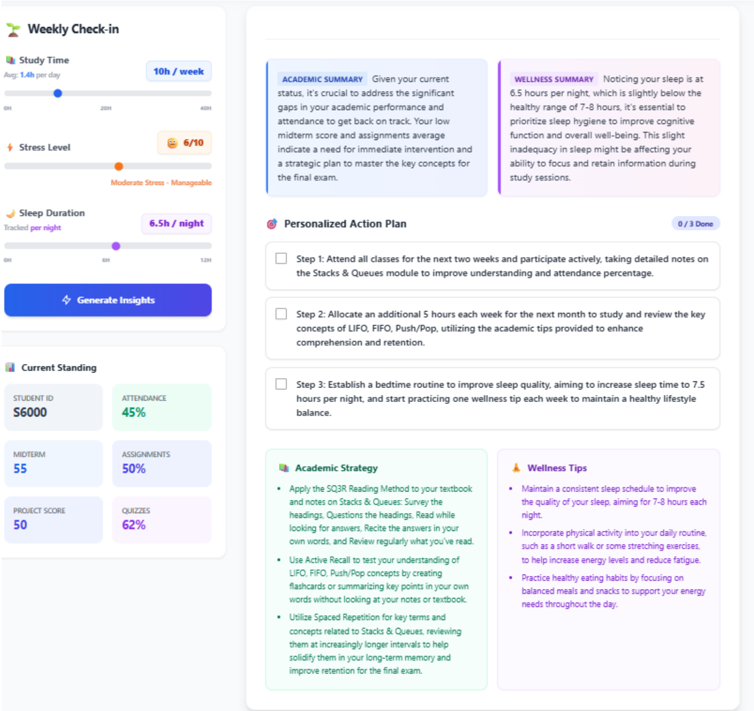

Figure 4- AI-generated personalized action plan rendered on the student dashboard

Table 7 breaks down a sample of the real terminal test cases processed by the system to show how dynamically the rules adapt to different profiles.

Table 7- System-Generated Insights

Student	**Rec. Index**	**Status Profile**	**System-Generated Insights**
Student A	12.8%	ON TRACK	**Burnout Warning:** Great dedication, but stress is creeping up. Remember to schedule downtime.
Student B	10.9%	ON TRACK	**Improvement**: You are doing well, but try to participate more in class discussions! **Wellness**: Your grades are great, but better sleep will make studying feel much easier.

These results highlight that the system is highly adaptable and functions as a true digital mentor. It doesn't simply tell every student to "study harder." As seen with Student A, even with a low recommendation index of 12.8%, the system catches the creeping stress anomaly and issues a preventative burnout warning. With Student B, the system isolates specific behavioral tweaks, like speaking up in class and getting better sleep, to optimize an already strong academic profile. This proves the system's massive value for modern, holistic academic advising.

- Disengagement Detection and Adaptive Intervention System

The Disengagement Detection and Adaptive Intervention System proves to be very effective in detecting the early signs of disengagement among students and offering the appropriate support measures. It is based on a hybrid machine learning methodology, combining a GRU Autoencoder to analyze behavior and a Reinforcement Learning (RL) decision making model.

GRU Autoencoder was trained on the sequential student activity data, which consisted of the frequency of the log-in, the duration of a session, the total time spent on activities, and the task completion behavior. The training and validation loss during the training process had a progressive decreasing trend which means that the model was able to learn the normal pattern of engagement of the students without overfitting. This confirms that the model is able to model and capture changes in behavior over time.

Values of reconstruction errors were computed per student sequence following training. The patterns of abnormal behaviors were identified using these values. Students who had greater reconstruction errors were found to be at high risk of disengagement. These cases were primarily linked to low activity of logging in, low participation, and missed academic activities. Thresholds that were based on percentiles enabled the system to be able to make a clear distinction between normal and at-risk students.

After identifying the risk level, the RL model was applied to identify the most suitable intervention action. The RL model was trained on a set of about 25,000 past interaction samples, which enabled it to learn the effects of various actions on student engagement across time. Depending on the present condition of the student, the level of risk and the trend of behavior, the model chooses the actions, which can be either doing nothing, sending a soft nudge or a reminder.

Table 8- Distribution of Recommended Actions

Action	Number of Decisions	Percentage
DO_NOTHING	13,083	52.3%
SOFT_NUDGE	10,223	40.9%
REMINDER	1,694	6.8%

The results indicate that the system mainly prefers low-pressure interventions such as soft nudges, which are suitable for moderate-risk situations. The increased risk level makes use of more direct reminders. This behavior will prevent fatigue in checking alerts and yet offer a timely support to students.

By and large, the system is well integrated in its ability to integrate behavior-based risk detection and adaptive decision making. It does not only detect disengagement early but also offers context-specific and individualized interventions. This enhances the capacity to re-engage students and help achieve improved learning outcomes in a scalable and practical manner

- Conversational AI Assistant for Adaptive Academic Interventions

The Conversational AI Assistant demonstrates good results in the area of intelligent and easy-to-use academic assistance. A hybrid AI solution that integrates Retrieval-Augmented Generation (RAG) with a web search fallback is used in the system, enabling it to respond to both university-specific and general knowledge questions.

Rules, guidelines and regulations were academic documents that were processed and stored in a vector database. This allows the system to retrieve the relevant information as per the query by the user. Semantic search enhances the quality of responses particularly on domain specific queries.

The system also facilitates various forms of interaction such as text, voice as well as image input. This enhances accessibility and makes the system more user friendly to various kinds of users. To facilitate checks and balances, an image validation step was added to make sure that only the relevant images are processed, which contributes to system accuracy and reliability.

Moreover, the assistant has a personalization option based on the performance of the students. According to that, the students are grouped into various levels and the system makes specific recommendations, reminders and inspirational messages. This assists in designing a more adaptive and accommodating learning process.

Action-based features are also supported by the chatbot, as a user can use them to perform tasks, like accessing documents, and interact with system functions directly through dialogue. Local language model guarantees enhanced privacy of data, and also allows the system to operate even when the internet is down.

In general, the system proves to be effective in enhancing information access, personalization, and user interaction in an academic setting.

3.2 Discussion

- Hybrid Ensemble Engine for Cumulative Risk Assessment

  The findings of the Hybrid Ensemble Model show that it is able to accurately predict student academic risk. The model uses a range of machine learning techniques to identify various patterns in student data, which improves the precision of the model compared to using just one machine learning algorithm. This improves the overall quality of predictions and supports better academic decision-making.

  A major advantage of the system is that it successfully identifies most not-at-risk students, and a significant number of at-risk students. Although a small number of at-risk students were not identified, this limitation is acceptable in early warning systems, as the model still captures most students who need attention without generating too many false alerts. This balance is important to avoid overwhelming educators with unnecessary interventions.

The classification results also show the effect of class imbalance in educational data. Since at-risk students are fewer compared to not-at-risk students, achieving high recall for the at-risk category is more challenging. But the overall performance with a balance between precision and recall, as represented by the F1-score, is good. This suggests that the system works well with different data sets.

Finally, it is worth noting the contributions of academic factors. The model weighs academic performance features such as final marks, assignment marks and attendance more heavily. These are interpretable and useful measures of student performance and engagement. This enhances the practicality of the model as it is easier for educators to relate to and respond to these factors.

Also, the model's capacity to distinguish between risk factors means students can be identified appropriately in relation to their performance. This facilitates early detection of at-risk students, and enables institutions to take appropriate actions.

But there are also some drawbacks. The model is primarily based on academic and behavioral data, but it may not take into account other factors such as personal problems or stress, which can influence academic performance. Also, the model's accuracy is dependent on the training data set used.

In conclusion, the Hybrid Ensemble Model offers a promising approach for predicting academic risk. Its high accuracy, reliable classification performance, and focus on meaningful features make it suitable for integration into intelligent academic support systems. With ongoing development and testing, it can contribute to improving academic performance and preventing dropouts.

- AI-Generated Personalized Action Plan and Recommendation

This study shows that the AI-Generated Personalized Action Plan system offers a promising way to transform student data into valuable academic advice. Its use of the Hybrid Ensemble prediction model and generative AI allows the system to provide more than a classification, and to provide practical guidance to students.

A key feature of this system is the use of a Recommendation Index, rather than a discrete class. This enables the system to account for slight variations in student performance and behavior, thus providing more adaptable recommendations. It can pick up on subtle changes in student performance and behavior, which can be critical for early intervention.

The use of generative AI also allows the system to convert quantitative predictions into recommendations. Rather than simply providing general advice, the system provides context-specific advice based on student performance and behavior, as well as the syllabus. This makes it more accessible and relevant to students and helps them understand what to do.

Another key benefit is the ability of the system to help both weak and good performing students. As shown in the results, students with low risk can also get valuable advice like stress and behavior. This demonstrates that the system not only targets low performing students but also assists others to keep their performance high or improve it, making a more balanced approach to the students' growth.

But there are some limitations. The model's predictive performance is impacted by class imbalance, with a smaller number of at-risk students in the data. This results in a reduced recall of the minority class, which may result in some students not being identified. Furthermore, the effectiveness of the recommendations also relies on the information and context given to the generative AI model. If the recommendations are based on insufficient or inaccurate data, they may not be as effective.

In spite of these challenges, the system has great potential as a smart academic advising system. Its ability to combine prediction, personalization, and natural language generation creates a more engaging and supportive learning experience. In general, this enhances academic advising and supports students in making informed choices about their studies.

- Disengagement Detection and Adaptive Intervention System

The results confirm that the system can effectively identify early signs of student disengagement and provide suitable intervention strategies. The system provides a more proactive solution than traditional monitoring techniques as it combines sequential analysis of behavior with adaptive decision making.

GRU based approach comes in handy especially when time-based variation in student behavior is required. As the engagement of students is not constant, the possibility to examine the trends of weekly activity: the number of logins, participation, and finishing the tasks is used to identify gradual disengagement with more precision. This contributes towards making the system more reliable in detecting at-risk students at an early age.

The reinforcement learning model also enhances the system as it allows making individual decisions regarding the intervention. The system does not provide the same action to all students, but rather it picks the action depending on the current state of the student. The results show that the model prefers low-pressure interventions like soft nudges in most situations, which helps to maintain student motivation without causing unnecessary stress or alert fatigue.

The fact that the system can be changed over time is another significant benefit of the system. The more the RL model is trained, through previous interactions, the better it is at making its decisions hence the interventions become more effective and context-sensitive. This dynamic behavior renders the system appropriate to real world educational situations where student engagement patterns are in a continual state of change.

Nevertheless, certain constraints can be noticed. The existing model is primarily based on LMS behavioral data and fails to examine the external variables, including personal problems, stress status, or offline learning activities, which also can influence student engagement. Moreover, the response of the students to the intervention actions can be different among students, and the effectiveness of the intervention actions can be affected.

Future enhancements can involve adding more data sources like emotional or contextual data to improve prediction accuracy. The detection and intervention components can also be refined with the help of the further assessment based on real-life student data and feedback.

Overall, the system provides a practical and scalable solution for early disengagement detection and personalized student support. It possesses great potential to enhance student retention and learning through timely and smart academic interventions.

- Conversational AI Assistant for Adaptive Academic Interventions

The results highlight that the Conversational AI Assistant provides a more advanced and practical approach to academic support compared to traditional systems. RAG with web search fallback enhances the system to respond to many queries, and the probability of failure to respond or incompleteness is minimized.

The main strength of the system is the multimodal interaction capability. The system is more convenient since it supports text, voice and image input. Reliability of the image validation process is also enhanced in that irrelevant inputs are filtered.

Another important contribution is the personalization feature. The system can leverage on the student performance data to offer more specific and focused assistance as opposed to offering broad responses. This demonstrates that the system goes beyond mere question answering and facilitates adaptive interventions in academics.

The ability to add action-based features is another useful feature that provides users with an opportunity to accomplish tasks using the chatbot. This minimizes the use of manual navigation and the system is more efficient in practical application.

A local language model is also an added value as it enhances privacy of data and enables offline operations. This would be particularly helpful in a learning institution whereby sensitive student information must be secured.

Nevertheless, there are certain restrictions. The quality and coverage of the uploaded academic documents determines the performance of the system. In addition, the reliability of the external sources may be different when it comes to the response of the web search fallback. The feature of personalization is restricted as well by the accessibility of correct student data.

All in all, the system has good potential to be an intelligent and adaptable academic support tool. It enhances accessibility, customization and interaction and is therefore appropriate in the contemporary digital learning process.

4 CONCLUSION

The rapid growth of digital learning environments and increasing student population, especially in higher education institutions, highlight the need for more intelligent and proactive academic support systems. The conventional methods based primarily on end-of-term evaluations and the use of manual processes are inadequate to detect the early signs of student disengagement or offer prompt assistance. As the solution to these issues, the current research proposes applying an AI-based, multifunctional academic monitoring and intervention system that could facilitate the early identification of risks, ongoing monitoring of engagement, and one-on-one help with students.

Unlike many existing systems that focus only on academic performance or provide limited support features, the proposed platform combines multiple advanced technologies into a single unified solution. It combines risk prediction using machine learning, behavioral monitoring with GRUs, intervention strategies implemented through reinforcement learning, and a conversational AI assistant to form a holistic and dynamic academic support system. This enables the system not only to recognize at-risk students but also to take appropriate and personalized actions to respond to them.

One of the key strengths of this system is its proactive approach to student support. The system is able to identify the early signs of disengagement by constantly examining behavioral data like LMS activity, patterns of logins, and levels of participation even before it can be manifested in academic outcomes. This ability is further strengthened by the reinforcement learning component which allows the appropriate intervention strategies to be chosen so that the students are provided with the appropriate kind of support at the appropriate time without undue pressure.

Accessibility and interaction with the user can be enhanced by the presence of a Conversational AI Assistant. The platform allows students to find academic information, take guidance, and carry out activities by allowing them to interact with the system using natural language. This not only lessens the academic staff work load but also provides a more interactive and student friendly learning experience.

The other valuable addition of the system is that it offers personalized and data-driven assistance. The platform does not provide generic solutions but rather changes its recommendations depending on the behavior and performance of the individual students. This assists in enhancing student engagement, motivation and learning outcomes.

In conclusion, the proposed system represents a comprehensive and intelligent academic support solution that transforms traditional reactive education models into proactive and adaptive learning environments. The combination of prediction, monitoring, intervention and intelligent interaction into one platform enables the system to create a high possibility of making students more successful, increasing retention rates, and enabling academic decision-making to be more effective. As it is further developed and implemented into the real world, it is a potentially useful model to contemporary education systems and online learning platforms.

5 SUMMARY OF EACH STUDENT'S CONTRIBUTION

Table 9- IT22354792 Student Contribution

Registration No	Name	Task Description
IT22354792	Ravisanka U.V.P	• Data Collection and Preprocessing Sourced and preprocessed a comprehensive 17-attribute academic dataset encompassing demographic, socioeconomic, and historical performance factors. Handled missing values, normalized numerical features, applied label encoding to categorical attributes, and resolved class imbalances using targeted weighting. • Model Development and Training Developed and trained a baseline Hybrid Ensemble predictive engine. Configured a soft-voting classifier aggregating Random Forest (mapped non-linear boundaries), XGBoost (handled structured data), and LightGBM (ensured high computational efficiency) weighted at 2:2:1 to calculate a highly accurate cumulative risk probability. • Explainable AI (XAI) Integration Integrated a SHAP (SHapley Additive exPlanations) TreeExplainer directly into the machine learning pipeline. Engineered the system to extract and visualize per-feature contribution scores, ensuring the algorithmic risk assessments are completely transparent and interpretable for academic staff.

Table 10- IT22370228 Student Contribution

Registration No.	Name	Task Description
IT22370228	Disanayaka S.T	• Support Index Calculation (Predictive Modeling) Engineered a distinct Hybrid Ensemble Machine Learning model utilizing Random Forest, Gradient Boosting, and Logistic Regression. Implemented a soft-voting mechanism to calculate a quantitative "Support Index" percentage, dynamically classifying students into three targeted stages: On Track, Needs Attention, or Priority Support Needed, completely replacing standard binary risk models. • Generative AI Integration (Groq API) Developed the cognitive intervention framework by connecting the mathematical Support Index to the high-speed Groq API (Llama 3.1). Utilized strict prompt engineering to feed the student's classification, stress levels, and weekly syllabus into the LLM, forcing the system to generate highly personalized academic and wellness plans strictly formatted as valid JSON objects. • Full-Stack Dashboard Development Built the interactive user interface utilizing a React.js frontend and a Firebase backend. Seamlessly integrated the RESTful APIs to visualize the AI-generated JSON recommendations on a live dashboard, providing students with immediate, actionable, and syllabus-aware guidance..

Table 11- IT22365750 Student Contribution

Registration No	Name	Task Description
IT22365750	Nimanji D.L.K	• The Conversational AI Assistant was created by gathering and organizing academic data, including university rules, regulations, lecturer information, and administrative documents. These resources were processed with OCR techniques and turned into structured text. The data was then divided into smaller sections and transformed into embeddings to create a semantic knowledge base in a vector database. This setup allows for quick and precise information retrieval. • A hybrid Retrieval-Augmented Generation (RAG) architecture was designed and put in place. It includes a web search fallback mechanism to enhance response coverage and reliability. This supports interactions through text, voice, and images A personalization feature was added, which uses student performance data. This allows the system to categorize users and provide targeted recommendations, proactive reminders, and motivational feedback based on their academic performance. • In addition, the chatbot gained action-based capabilities, allowing users to access documents, navigate system features, and retrieve academic information through conversational commands. The system also protects data privacy and reliability by employing a local Large Language Model through Ollama, which supports offline functionality. Finally, all components were integrated and tested to ensure proper performance and overall effectiveness.

Table 12- IT22902702 Student Contribution

Registration No	Name	Task Description
IT22902702	Perera I.A.T.D	• Data Collection and Preprocessing- Student behavioral data gathered and compiled in Learning Management System (LMS) activities, such as the number of logins, time spent on the activity, time spent on overall activity, task completion, and response behavior. Formatted the data as a sequence that could be easily used to model time-series and used normalization methods of MinMaxScaler to enhance model performance. • GRU Autoencoder Model Development Developed and trained a GRU-based autoencoder network to provide normal student engagement behavior patterns with the use of sequential behavioral observation data. Trained the model with training and validation loss monitors and used early stop to prevent overfitting. Calculated reconstruction error values to indicate abnormal behavior and identify risk of disengagement. • Reinforcement Learning-Based Intervention System- Constructed a Q-learning-based reinforcement learning algorithm to choose the right intervention strategies depending on the level of risk and behavioral state of students. Planned an action-set of soft nudges, reminders and escalation plans to enhance student re-engagement and reduce unnecessary notifications. • System Integration and Testing- Combined the GRU model and reinforcement learning module with system and backend APIs. Carried out testing by use of different student scenarios to ascertain accurate risk prediction, proper decision making and smooth system functionality.

REFERENCES

[1] M. Yağcı, "Educational data mining: Prediction of students' academic performance using machine learning algorithms," Smart Learning Environments, vol. 9, no. 1, pp. 1-19, 2022.

[2] G. Siemens and R. S. Baker, "Learning analytics and educational data mining: Towards communication and collaboration," in Proc. Learning Analytics and Knowledge Conf. (LAK), 2012.

[3] A. M. Shahiri, W. Husain, and N. A. Rashid, "A review on predicting student's performance using data mining techniques," Procedia Computer Science, vol. 72, pp. 414-422, 2015.

[4] X. Chen, Y. Liu, and M. Zhao, "A novel student performance prediction model based on GRU and attention mechanism," IEEE Access, vol. 11, pp. 23450-23461, 2023.

[5] R. Mubarak, H. Wang, and S. M. Zin, "Predicting student performance using LSTM and GRU sequence models," IEEE Transactions on Learning Technologies, vol. 15, no. 2, pp. 190-202, 2022.

[6] L. Zhang, et al., "Anomaly detection in student learning behaviors using autoencoder architectures," IEEE/ACM Transactions on Audio, Speech, and Language Processing, vol. 29, pp. 3412-3424, 2021.

[7] S. Kizilcec, C. Piech, and E. Schneider, "Deconstructing disengagement: Analyzing learner behavior in MOOCs," in Proc. Learning Analytics and Knowledge Conf. (LAK), 2013.

[8] S. H. Lundberg and S. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems (NeurIPS), 2017.

[9] B. Shneiderman, "Human-centered artificial intelligence," Communications of the ACM, vol. 63, no. 1, pp. 58-61, 2020.

[10] J. Wang, T. Li, and C. K. Looi, "Explainable AI in Educational Data Mining: Using SHAP to interpret predictive models," IEEE Transactions on Education, vol. 67, no. 1, pp. 45-56, 2024.

[11] M. A. Alam and M. S. Hossain, "Trustworthy learning analytics: Integrating feature attribution methods for student retention," IEEE Access, vol. 10, pp. 115023-115035, 2022.

[12] Y. Zheng, et al., "Reinforcement Learning for Personalized Education: Adaptive Intervention Strategies," IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 8, pp. 8120-8134, 2023.

[13] A. N. Rafferty, H. Ying, and P. A. Williams, "Using reinforcement learning to design personalized educational interventions," Journal of Educational Data Mining, vol. 13, no. 3, 2021.

[14] D. Zhou and K. Koedinger, "Dynamic decision making in intelligent tutoring systems via Markov Decision Processes," IEEE Transactions on Learning Technologies, vol. 14, no. 4, pp. 512-525, 2021.

[15] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD), pp. 785-794, 2016.

[16] G. Ke, et al., "LightGBM: A highly efficient gradient boosting decision tree," in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, pp. 3146-3154, 2017.

[17] A. Al-Shabandar, A. Jaddoa, P. Liatsis, and A. J. Hussain, "A cloud-based educational analytics architecture using serverless backend services," IEEE Transactions on Cloud Computing, vol. 10, no. 2, pp. 1021-1034, 2022.

[18] S. V. Kumar and R. P. Singh, "Scalable Microservice Architectures for Real-Time Educational Dashboards: Integrating Fast APIs and Reactive Web Frameworks," IEEE/ACM International Conference on Automated Software Engineering (ASE), 2023.

[19] C. Romero and S. Ventura, "Educational data mining: a review of the state of the art," IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), vol. 40, no. 6, pp. 601-618, 2010.

[20] K. Verbert, E. Duval, J. Klerkx, S. Govaerts, and J. L. Santos, "Learning dashboards: an overview and future research opportunities," Personal and Ubiquitous Computing, vol. 18, no. 6, pp. 1499-1514, 2014.

APPENDIX

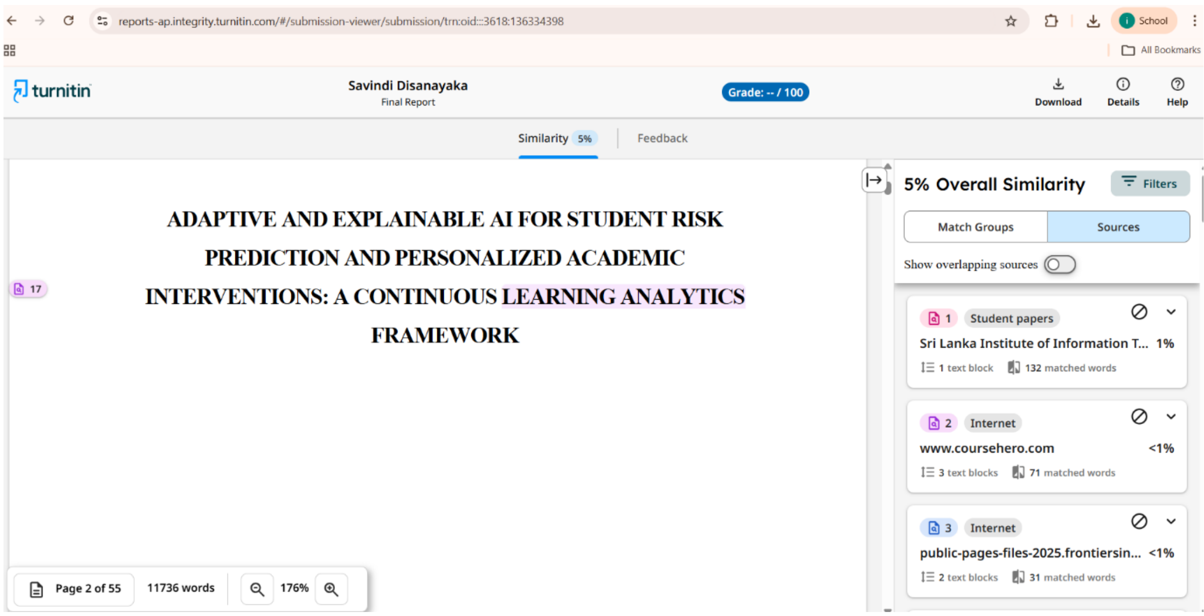

Figure 5-: Plagiarism Report