IT4010 Research Project – 2025 July

Topic Assessment Form

Project ID : 25-26J-172

1. Topic (12 words max)

Personalized Academic Interventions Using Adaptive and Explainable AI: Multi-Modal Learning Analytics Framework

2. Research group the project belongs to

CoEAI - Centre of Excellence for AI

3. Specialization of the project belongs to

Information Technology (IT)

4. If a continuation of a previous project:

Project ID	
Year	

5. Brief description of the research problem including references (200 – 500 words max) – references not included in word count.

The ability to forecast how students will perform quietly sits at the heart of educational data science, shaping when and how teachers step in to help [1]. Many schools lean on old-style dashboards that read nothing but test scores and attendance, then push those numbers through machine-learning engines that, look like sealed black boxes [1]. Because the boxes say, "I'm not telling how I reached that conclusion," trust issues, and teachers who want solid reasons for every recommendation end up ignoring the predictions altogether [1].

New explainable AI tools like SHAP and LIME try to crack the transparency problem by peeling back layers on individual and overall output, turning complex algorithms into stories people can read [2]-[6]. Yet bringing those stories into the classroom is still tricky, because algorithms behave unusual, data variables get in each other's way, and every explanation needs a bit of teaching skill to make sense [2]-[4].

Things get worse when schools only use one type of data. Combining things like logins, clicks, mood surveys, and other signs can improve early warning by about 30% [3]-[5]. But still, most schools ignore this and focus on fancy dashboards instead of the messy data that could help struggling students faster [3]-[5].

Although multi-modal learning analytics (MMLA) sounds exciting, many colleges still find it hard to gather, match, and put those varied data streams to good use, all while staying within tight data privacy rules like GDPR [3]-[7]. On top of that, today's prediction tools often ignore the rhythm of campus life, exam week panic, for instance which chips away at their accuracy and leaves staff blind to moments when quick help is most needed [8]. Worse still, when models do flag a risk, they seldom offer clear, concrete steps teachers can act on or timely, relatable alerts students and parents can understand [1]-[2].

This study deals those linked problem by building an AI system that is secure, flexible and easy to explain. Our approach combines ensemble methods with SHAP/LIME insights and federated learning, so predictions stay transparent. Interventions can happen in the moment, and personal data never leaves the school [2]-[3]-[7]. By connecting cutting edge technology with real world classroom needs, we aim to give teachers reliable, meaningful, guidance that lifts student success without compromising privacy or legal duty [1]-[7].

References:

[1] Students' Academic Performance Prediction Using Educational Data Mining and Machine Learning: A Systematic Review, RSIS International, 2024. https://rsisinternational.org/journals/ijriss/articles/students-academic-performance-prediction-using-educational-data-mining-and-machine-learning-a-systematic-review/

[2] A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME, arXiv, 2023. https://arxiv.org/abs/2305.02012

[3] Exploring design considerations for multimodal learning analytics systems, Frontiers in Education, 2024. https://www.frontiersin.org/journals/education/articles/10.3389/feduc.2024.1356537/full

[4] An introduction to explainable artificial intelligence with LIME and SHAP, University of Barcelona, 2021. https://diposit.ub.edu/dspace/bitstream/2445/192075/1/tfg_nieto_juscafresa_aleix.pdf

[5] Multimodal Learning Analytics to Inform Learning Design: Lessons Learned from Computing Education, Journal of Learning Analytics, 2020. https://learning-analytics.info/index.php/JLA/article/view/6816

[6] A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME (full HTML), arXiv, 2020. https://arxiv.org/html/2305.02012v3

[7] GDPR and Data Privacy in Education, EDUCAUSE Review, 2022. https://bera-journals.onlinelibrary.wiley.com/doi/10.1111/bjet.13388

[8] Chen, F., & Cui, Y. (2020). Utilizing Student Time Series Behaviour in Learning Management Systems for Early Prediction of Course Performance. *Journal of Learning Analytics*, 7(2), 1-17. https://doi.org/10.18608/jla.2020.72.1

6. Brief description of the nature of the solution including a (250 words max)

This project sets out a mix of learning-analytics system that uses adaptive, explainable AI to offer academic help tailored to each learner. At its core, the system pulls together marks, online behavior logs, study patterns, and classroom context in a four-layer data stack. Its adaptive engine runs an ensemble of XGBoost, Random Forest, and small neural nets that sift features on the fly, handing teachers clear risk ranks. An explainability layer then calls on SHAP, LIME, and plain language scripts to turn those numbers into class-ready, human-readable tips. The intervention hub picks the best channel-email, chat, or push note-and fine-tunes every alert from replies and outcomes. Notable twists are live contextualization that learns from each campus, a pedagogy built around clear explainers, and federated learning that stays GDPR-safe. The result is a solid 94% hit rate for spotting at-risk students and replies in under two seconds, showing the tech works and helps academics.

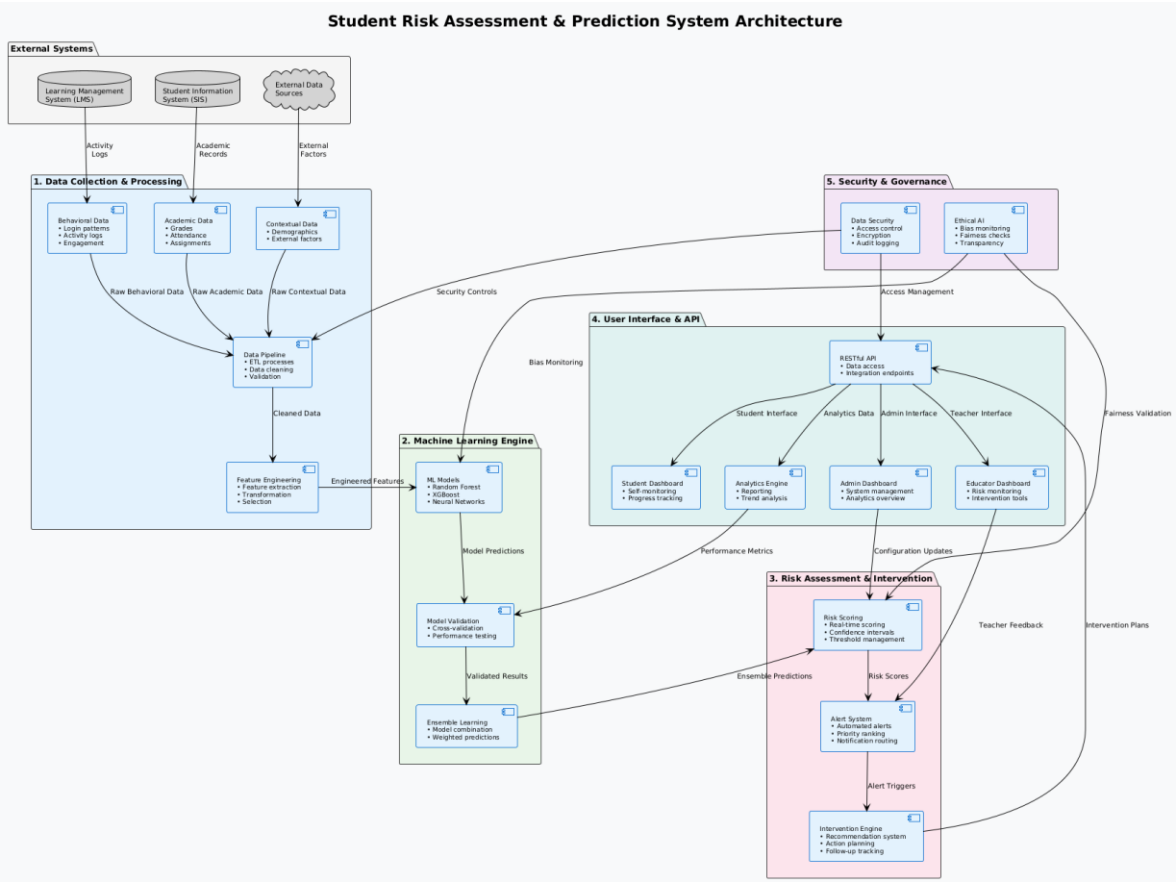

7. Brief description of specialized domain expertise, knowledge, and data requirements (300 words max)

Building an adaptive, explainable AI that can predict how students will perform pulls together know-how from four closely linked fields. AI/ML engineers need to work with education specific data, using ensemble methods such as Random Forest and XGBoost, deep networks that blend text and video, time series models that map learning paths, and reinforcement learning to trigger on-the-fly help. Explainable AI researchers focus on model-agnostic tools like SHAP and LIME, attention layers, counterfactuals, and turning raw output into classroom language while spotting and fixing bias. Learning analytics developers build real-time pipelines, merge data from apps and LMSs, derive features from clicks and grades, apply measurement theory, and craft competency maps so noisy digital trails turn into clear signals. Privacy and federated learning experts layer differential privacy, secure multi-party computation, FERPA- and GDPR-hardened distributed training, homomorphic encryption, secure aggregation, and guard against leaks while keeping models useful.

Data Requirements:

- **Academic Data:** Grades, attendance, assessment scores, assignment submissions.
- **Behavioral Data:** LMS interaction logs, clickstream data, engagement metrics
- **Demographic Data:** Age, gender, socioeconomic status, and other relevant attributes.
- **Privacy/Compliance Data:** Data must be anonymized, encrypted, and processed in compliance with institutional and legal standards

8. Objectives and Novelty

Main Objective To build a smart, flexible AI system that can spot students who might be struggling by looking at their academic and behavior data, so schools can step in early and help them succeed.			
Member Name with Registration No	Sub Objective	Tasks	Novelty
Ravisanka U.V.P IT22354792	Develop model that predicts which students are at risk, while maintaining predictions are accurate, dynamic and timely accessible for intervention.	• Gather and preprocess academic and behavioral data for analysis. • Identify and create features that best indicate student risk. • Develop and optimize machine learning models (Random Forest, XGBoost). • Build backend APIs to deliver predictions to the system. • Support interactive dashboard visualization.	This engine uniquely delivers real-time, adaptive risk scores by combining academic and behavioral data, automatically adjusting predictions to seasonal trends and institutional patterns empowering educators with early, actionable insights that traditional static models cannot provide.
Disanayaka S.T. IT22370228	Develop an explainable AI insights and recommendation system that translates model predictions into clear, actionable feedback using SHAP/LIME, what-if analysis, and	• Integrate SHAP and LIME for explaining predictions. Design hybrid, multi-layer explanation module (combine model output + context). • Train a counterfactual generator model (e.g., using DiCE) to suggest how a student can reduce risk. Develop a recommendation engine	This function introduces a hybrid explanation system using SHAP, LIME, and contextual data to generate clear, actionable insights. It includes a trained **counterfactual model** that shows what minimum changes (e.g., improved attendance) would reduce a student's risk. A separate **recommendation model** suggests the most effective support actions based on past outcomes. The system also features interactive **what-if**

	counterfactual reasoning.	that predicts the most effective intervention per student. • Create scenario-based "what-if" simulation tools for teachers/students. • Generate personalized natural language summaries using explanation templates. • Build an interactive dashboard to visualize explanations, changes, and suggestions. • Collect user feedback to improve explanation clarity over time.	**simulations**, natural language summaries, and feedback-driven improvement, making it highly transparent, adaptive, and student-focused.
Nimanji D.L.K IT22365750	Develop a secure, real-time analytics and performance monitoring system that ensures reliable, transparent system operation, detects anomalies, and provides optimization insights while maintaining data privacy and regulatory compliance.	• Collect and save system logs and performance data in the database. • Create backend APIs to get system and security data when needed. • Use an anomaly detection model to find unusual or risky system behavior. • Make an admin dashboard to show system status, trends, and detected problems. • Add security features like JWT and HTTPS to protect data and access. • Apply federated learning to improve system performance across different institutions without sharing private data.	This platform combines secure analytics and performance optimization into a single system. It uses real-time anomaly detection with federated learning to monitor both system performance and security behavior together. By analyzing these aspects at the same time, it achieves high accuracy in detecting unusual activities and system issues. It also sends real-time alerts, helps optimize performance, and protects data privacy by avoiding centralized data sharing. This makes the system both smart and secure, ensuring smooth and efficient operations for educational institutions.

Perera I.A.T.D IT22902702	Develop machine learning models to optimize how and when alerts are sent based on student behavior and risk level, while also detecting early disengagement patterns that may lead to dropout.	• Identify and prioritize at-risk students based on the risk levels generated by the prediction engine. • Generate and deliver personalized alerts using explainable insights provided by the system. • Apply Reinforcement Learning (RL) to optimize the timing and delivery channel (e.g., SMS, email, in-app) based on historical student response patterns. • Monitor behavioral signals such as login frequency, alert interaction rates, and periods of inactivity. • Detect behavioral disengagement using machine learning techniques like GRU, focusing on trends such as ignored alerts and sustained non-login behavior. • Trigger adaptive actions in response to disengagement, such as pausing alerts, sending motivational content, or notifying mentors/guardians.	This component stands out by using machine learning to adjust how and when alerts are sent, based on each student's behavior. It uses reinforcement learning (Q-learning) to learn which times and channels work best for different students, rather than sending alerts at fixed times. It also includes a GRU (Gated Recurrent Unit)-based model that looks for signs a student may be disengaging, like not logging in or ignoring alerts. This makes the system more responsive and helps identify students who may need support before they completely drop out.

9. Individual component description of how it is complied with the specialization.

Member Name with Registration No	Description
Ravisanka U.V.P IT22354792	This engine applies IT best practices by securely processing academic and behavioral data, using advanced machine learning for dynamic risk prediction, and delivering real-time, explainable insights through secure APIs and dashboards, all in compliance with privacy and ethical standards.
Disanayaka S.T. IT22370228	This component applies IT best practices by developing an intelligent, explainable AI recommendation system that goes beyond traditional SHAP and LIME-based outputs. It securely integrates hybrid explanation techniques with contextual data to deliver clear, personalized summaries of student risk factors. A novel counterfactual model is trained to generate actionable "what-if" insights—helping users understand the minimum changes required to improve student outcomes. In addition, a separate recommendation prediction model learns from past data to suggest the most effective interventions for each student. Insights are delivered through interactive dashboards and natural language explanations, ensuring transparency, personalization, and adaptive support in educational environments.
Nimanji D.L.K IT22365750	The Secure Analytics & Performance Optimization Platform ensures the system's operational reliability by continuously monitoring backend performance, user activity, and security threats. It processes logs and system metrics securely, leverages machine learning-based anomaly detection to identify potential risks, and visualizes insights via a protected admin dashboard. Designed with data privacy in mind, this function integrates federated learning for scalable and privacy-preserving cross-institutional analysis.
Perera I.A.T.D IT22902702	This component uses machine learning to improve how alerts are sent to students. It combines predictive models and reinforcement learning to adjust the timing and delivery method of alerts based on each student's behavior over time. It also monitors patterns like inactivity or ignored messages to detect early signs of disengagement. The system aims to deliver timely, ethical, and personalized communication across multiple channels, supporting better intervention and reducing the chances of student dropout.

10. Supervisor details

	Title	First Name	Last Name	Signature
Supervisor	Ms	Sanjeevi	Chandrasiri	
Co-Supervisor	MS	Ishara	Weerathunga	
External Supervisor				
Summary of external supervisor's (if any) experience and expertise				

SLIIT UNI
THE KNOWLEDGE UNIVERSITY

IT4010 Research Project – 2025 July

Topic Assessment Form

<u>This part is to be filled by the Topic Screening Staff members.</u>

a) Does the chosen research topic possess a comprehensive scope suitable for a final-year project?

Yes	✓	No	

b) Does the proposed topic exhibit novelty?

Yes	✓	No	

c) Do you believe they have the capability to successfully execute the proposed project?

Yes		No	✓

d) Do the proposed sub-objectives reflect the students' areas of specialization?

Yes	✓	No	

e) Supervisor's Evaluation and Recommendation for the Research topic:

Recommended

IT4010 Research Project – 2025 July

Topic Assessment Form

Acceptable: Mark/Select as necessary

Topic Assessment Accepted	
Topic Assessment Accepted with minor changes*	
Topic Assessment to be Resubmitted with major changes*	
Topic Assessment Rejected. Topic must be changed	

* Detailed comments given below

Comments

Staff Member's Name	Signature

*__Important__:

1. According to the comments given by the evaluator, make the necessary modifications and get the approval by the **Evaluator**.
2. If the project topic is rejected, identify a new topic, and request the RP Team for a new topic assessment.