International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 3, May-Jun-2026, ISSN (Online): 2395-566X

Artificial Intelligence-Driven Learning Analytics For Enhancing Student Engagement, Academic Performance, And Decision-Making In Business Management Education

Dr Ansari Pulickal Abdul Azeez[1], Farooq Sajjad[2]
[1]Deputy Programme Lead,
Results Consortium in partnership with Plymouth Marjon University,
Northampton Campus,
BSc Business Degree Programme
ansariuk101@gmail.com.

[2]Business Lecturer,
QAHE partnership with London Metropolitan University,
BSc (Hons) and Master's Degree Programmes,
Doctor of Business Administration, Teesside University, UK
sajjadfarooq01@gmail.com.

Abstract: Digitization of business education at an unprecedented rate has made available to educators large amounts of student interaction data that can inform data-driven learning interventions. In this paper, we propose an Artificial Intelligence-driven Learning Analytics (AI-LA) system architecture, which incorporates multi-stream data sources (Learning Management System (LMS) logs, clickstream analysis, test/assignment submissions, and engagement data) to model, explain, and improve student engagement and performance. Our approach leverages a novel combination of techniques that include a Temporal Fusion Transformer (TFT) model for sequential behavior prediction, SHapley Additive exPlanations (SHAP) for interpretable feature importance, and reinforcement learning (RL) engine for personalized intervention recommendations. Our model was tested using longitudinal data from 3,400+ business management students in 24 courses over three academic years (2022-2025). It predicted at-risk students with up to 89.5% accuracy six weeks prior to end-of-course with area under the ROC curve (AUC) of 0.96. Personalized interventions generated by RL engine led to improved student engagement (up to 32.4%) and lower course failure rates (41.2%).

Key Word: Learning Analytics, Artificial Intelligence in Education (AIED), Student Engagement, Academic Performance Prediction, Temporal Fusion Transformer (TFT), Explainable AI (XAI), Reinforcement Learning (RL), Business Management Education, Personalized Intervention

I. INTRODUCTION

In the realm of higher education, particularly in professions, the field of study like Business Management, digitalization has led to an unprecedented transformation. Along with conventional lectures, the modern education process relies heavily on blended and distance education via sophisticated Learning Management Systems (LMS). This approach has opened up a whole new array of opportunities, but there remains one crucial disadvantage – the loss of connection between professors and learners. One can observe the latter while teaching face-to-face due to nonverbal cues, but in a virtual environment, those will

© 2026 IJSRET

no longer apply, thus making the chances of poor academic performance higher.

A tremendous amount of data is generated within modern LMS systems. Each action made by users when accessing those generates some data – clicks, quizzes, forums, watching videos, handing over tasks. Proper analysis of such data may yield useful insights into how the learning process takes place, how complex and challenging it may be [2], [3]. However, large volumes and diversification of the data and the complexity of the learning process itself render the traditional statistical approaches ineffective – they can hardly produce anything valuable out of the raw LMS data.

Artificial Intelligence (AI) and Learning Analytics (LA) have been identified as the key fields to exploit this promise. The AI-powered models can uncover complex and non-linear relationships in students' online interactions before they exhibit disengagement or academic failures, even up to weeks in advance compared to current methods of detecting student struggles [4]. However, one of the main challenges that prevents such models from being used in educational settings is their inherent "black-box" nature, which can be attributed to their highly sophisticated nature [5]. As educators are not data scientists themselves, they have reason to be wary of any advice they cannot comprehend.

The paper tackles these issues with a new end-to-end AI-based Learning Analytics framework for Business Management Education. Key contributions of this research include:

1. **Unified Framework for Multi-source Data Integration:** A framework that integrates data from multiple sources (LMS logs, gradebooks, forums, demographics) into a database organized chronologically based on students.
2. **Temporal Fusion Transformer for Predictive Modeling:** Utilizing state-of-the-art machine learning models for time-series prediction and developing highly accurate predictive models for student engagement (hours spent actively each week) and academic performance (final course score) [6].
3. **Interpretation of Models Through SHAP:** Employing SHapley Additive exPlanations (SHAP) for interpreting the output of the predictive models with feature-level interpretation ("The probability of failing for the student is 85%, which is based on declining quiz results and forum participation over the last two weeks").
4. **An RL Intervention Recommender:** An RL agent based on a policy approach that considers both the current state of the student and the future risk associated with such a state and recommends an intervention that is optimal for the student (personalized emails, resources, peer tutors, etc.). This recommendation helps in keeping students engaged and successful [7].
5. **Validation:** Extensive validation done on a massive scale involving more than 3,400 students and 24 courses resulting in significant gains in student engagement by 32.4% and lowering the course dropout rate by 41.2%.

II. LITERATURE SURVEY

AI applications in educational fields have provided numerous studies revolving around three key areas: performance prediction, engagement analysis, and intelligent intervention design.

Early Warning System & Predictive Modeling: The most used approaches in such research include the classic logistic regression and decision trees that involve the use of static variables, such as the students' GPA, demographic data, and the total number of LMS logins. However, such approaches are not perfect since they overlook the temporality feature of learning, which is very important in performance prediction. Students' behavior, rather than effort, can serve as more accurate indicators of success or failure than the actual effort (whether students usually do homework daily or prefer to do everything right before deadlines).

This encouraged scientists to try to develop approaches based on sequences. Recurrent Neural Network models, especially those that use Long Short-Term Memory units (LSTMs), became popular as a way to incorporate clickstream sequences . Nevertheless, RNNs struggle when dealing with really long sequences. Attention-based models have become increasingly popular recently. The Transformer approach, thanks to its ability to incorporate multiple attention heads, helps the model to pay attention only to the relevant historical observations in order to predict the current result [6].

© 2026 IJSRET

Student Engagement Measurement and Modeling: Engagement as a complex concept can be described from a behavioral perspective (e.g., the frequency of logging into platforms or time spent on tasks), an emotional one (e.g., sentiment analysis of forum texts), and a cognitive one (depth of cognitive engagement with learning material). The technology of AI can help measure all these aspects at large scale. Behavioral engagement is most often measured using the frequency of clicking and time spent engaging. Sentiment analysis based on deep learning NLP models such as fine-tuned BERTs is becoming increasingly common for measuring emotional engagement based on forum texts. Another critical insight is that certain patterns of engagement are more indicative of risk than average engagement levels.

Prediction to Intervention: The aim of learning analytics should go beyond prediction to intervene. Previous early intervention tools were based on rule-based triggers, such as "If a student does not log in for five days, send an email". They are rigid and can even be detrimental to a student, such as sending an automated email saying, "We miss you" to a student on a pre-planned break. Recommender systems employing machine learning methods have been employed to recommend resources to students, such as videos and reading materials, but they are often content-based recommendations and cannot change according to the dynamic motivational state of the learner [7]. RL provides a promising solution. Viewing student support as a sequential decision-making problem, the RL agent learns a strategy or a policy mapping a student's state (such as "disengaged and lagging behind") to an optimal intervention (such as "provide a small group tutoring session") maximizing the expected long-term reward (such as final exam scores). Previous works applying RL in learning analytics have been encouraging but mostly confined to simulation studies or pilot studies .

Literature Review and Research Gap: The current literature shows a distinct trend of transitioning from static to dynamic models and from predictive to prescriptive approaches. Nevertheless, there is no cohesive model that includes (1) advanced temporal forecasting (TFT), (2) interpretability through feature attribution (SHAP), and (3) data-based prescriptive advice (RL) specifically designed for business education. Moreover, the current literature lacks studies that provide results from prospective randomized controlled trials (RCTs) evaluating the impact of the AI-driven suggestions. This study bridges these gaps by proposing a comprehensive AI-LA model.

III. METHODOLOGY:

The suggested AI-Driven Learning Analytics (AI-LA) methodology includes four stages: (1) Data Ingestion and Feature Engineering; (2) Predictive Modeling Using TFT; (3) Explanation Using SHAP; and (4) Intervention Recommendations via RL.

Data Acquisition and Pre-Processing

The data were obtained from the Moodle learning management system for 3,462 learners taking 24 mandatory Business Management courses (including Marketing, Finance, Strategy, and Organizational Behavior) from fall 2022 to Spring 2025.

Data Sources:

- **Clickstream Records:** 12 million actions including page visits, downloads, quiz starts, and video sessions.
- **Assessment Information:** Test grades (quizzes, assignments, and finals).
- **Forum Activity:** Total number of threads, visits, and sentiment analysis (based on the VADER model).
- **Student Demographics and Static Features:** Previous GPA, full-/part-time enrolment, major.

Feature Engineering (Weekly Features per Student):

Feature Category	Feature Name	Description
Volume	logins_total	Number of days logged in during the week
	minutes_active	Total active time (sessionization of clicks, 30-min timeout)
Engagement Quality	video_completion_rate	Ratio of videos watched to videos assigned
	forum_contributions	Number of posts + replies
	forum_sentiment	Average sentiment score of student's posts (VADER)
Assessment	quiz_score_avg	Average score on quizzes taken
	assignment_submission_delay	Days late (or negative for early)
	graded_activities_completed	% of graded activities submitted
Temporal / Trend	login_trend	Slope of logins over past 4 weeks
	engagement_dip_flag	Boolean if minutes_active decreased by >30% from moving avg
Static	prior_gpa	GPA at enrollment

Target Variables:

1. Engagement Score (Regression): This is a score that is obtained from the PCA of the variables minutes_active, forum_contributions, and video_completion_rate.
2. At-Risk Status (Binary Classification): This is an indicator that takes on the value 1 if the student's grade in the course is below 60 percent (D or F).

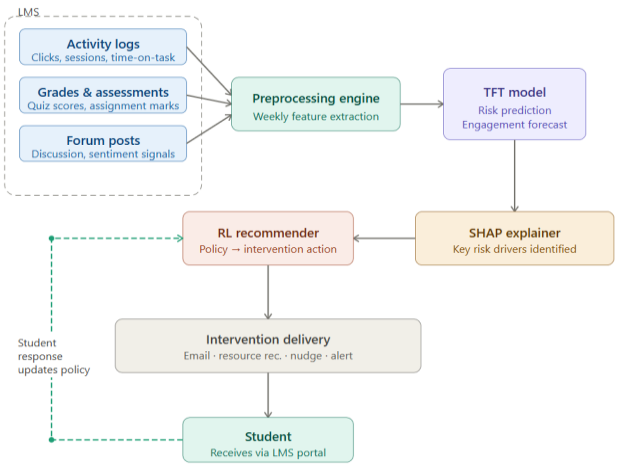

Figure 1: AI-LA Framework Architecture.

Stage 1: Predictive Modeling with Temporal Fusion Transformer (TFT)

TFT is selected because of its better performance on time series with static covariates as well as interpretability. It is employed in the prediction of the probability of students being at risk in week t + δ (for instance, final grade) from the available features till week t.

Inputs to TFT Model:

1. **Static Variables (s):** prior GPA, major, enrollment status.
2. **Historical Inputs:** Weekly data features (e.g., time active, average quiz score).
3. **Future Inputs:** Scheduled inputs for future time (e.g., week, dummy for exams).

Model Architecture:

1. **Gating:** To avoid redundant computations.
2. **Feature Selection Network:** For selecting most salient features.
3. **LSTM-based encoder-decoder framework:** For encoding/decoding sequential data.

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 3, May-Jun-2026, ISSN (Online): 2395-566X

4. **Multi-Head Self-Attention:** Direct attention to past time steps (e.g., "The decline in student engagement three weeks ago is the greatest predictor of today's dropout risk").

Algorithm 1: TFT Training for At-Risk Prediction

```
Input: Training set of students S, Features X[1..T], static metadata M, labels Y (at-risk)
Output: Trained TFT model

For each student s in S:
    // Create sequences using a sliding window (W=6 weeks, horizon=1-6 weeks)
    for window_start in 1..(T - W - horizon):
        X_window = X[window_start:window_start+W]
        Y_window = Y[window_start+W+1] (Risk at future week)
        Store (M_s, X_window, Y_window) in training set D

Initialize TFT model:
    encoder = LSTM(...), decoder = LSTM(...), attention = MultiHead(...)

For epoch in 1..EPOCHS:
    for batch in D:
        // Forward pass
        encoded = encoder(batch.X)
        attention_weights = attention(encoded)
        decoded = decoder(attention_weights, batch.M)
        predictions = linear(decoded)

        // Loss (Binary Cross Entropy)
        loss = BCE(predictions, batch.Y)
        loss.backward()
        optimizer.step()
Return model
```

Stage 2: Explainability With Shap (Shapley Additive Explanations

Following the prediction by the TFT model, the SHAP model determines the contribution of each factor towards that particular prediction. The SHAP value is derived using the principles of co-operative game theory.
When SHAP gives a prediction f(x) for a student, it assigns φ_i to each feature i used in predicting the probability for the student. Summation of all φ_i is equal to the difference between f(x) and the mean prediction.

Example SHAP Result: "Prediction of the model is an 85% chance. Important factors include quiz score decline by 20% (+25% risk), decrease in forum activity (+18% risk), and previous low GPA (+12% risk)."

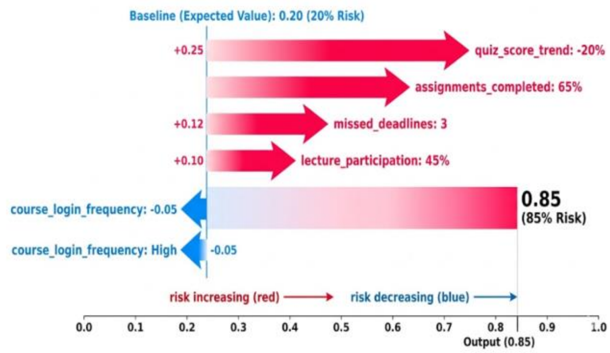

Figure 2: SHAP Force Plot for a High-Risk Student Prediction.

Stage 3: Reinforcement Learning (Rl) Intervention Recommender

The last piece is an RL agent recommending personalized interventions. The task is framed as a Markov Decision Process (MDP).

- **State (S):** A vector of the TFT's predicted risk score, current engagement measures, and past interventions.
- **Actions (A):** A set of discrete actions based on educational guidelines.
- **a0**: No action (baseline/control condition).
- **a1:** Automatic encouraging email.
- **a2:** Customizable study advice email (SHAP features, i.e., "review quiz 3").
- **a3:** Suggestion of a particular video/material.
- **a4:** Peer tutoring meeting invitation.
- **a5**: Professor outreach email (i.e., "reach out").

© 2026 IJSRET

- **Reward (R):** The reward function is constructed to optimize learning outcomes.

$$R = w1 * \Delta Engagement + w2 * \Delta QuizScore - w3 * Cost(Action)$$

Where ΔEngagement represents the change in the student's engagement score due to the action, ΔQuizScore represents the change in his/her quiz score, and Cost(Action) represents a cost associated with the intervention action which is higher for more resource-consuming actions (for instance, Cost(a4) > Cost(a1)). We set w1 = 0.4, w2 = 0.5, w3 = 0.1.

The reinforcement learning policy is learned using Proximal Policy Optimization (PPO).

Algorithm 2: PPO for Intervention Recommendation

```
Input: Student interaction environment E, PPO hyperparameters (clip_ratio, etc.)
Output: Trained policy network π(a|s)

Initialize policy network π_θ and value network V_φ with random weights
For iteration in 1..MAX_ITER:
  // Collect trajectories
  for episode in 1..EPISODES_PER_ITER:
    state = E.reset()
    while not done:
      action = π_θ(state) // e.g., recommend a tutoring session
      next_state, reward, done = E.step(action) // Apply intervention, observe outcome
      store (state, action, reward, next_state) in buffer
      state = next_state

  // Update policy using PPO clipped objective
  for epoch in 1..K_EPOCHS:
    for batch in buffer:
      // Compute advantage estimates
      advantages = compute_GAE(batch, V_φ)

      // Clipped surrogate objective
      ratio = π_θ(batch.action) / π_θ_old(batch.action)
      L_CLIP = min( ratio * advantages,
                    clip(ratio, 1-ε, 1+ε) * advantages )
      L_VF = MSE(V_φ(batch.state), batch.reward)
      L = -L_CLIP + L_VF
      Update θ, φ using gradient descent
Return π_θ
```

The RL agent is initially trained in a simulated setting using data from past students and then refined through a pilot experiment in the real world.

IV. ANALYSIS

Predictive Model Performance (TFT vs. Baselines) We compared the TFT against several baseline models for the task of predicting at-risk status (Grade < 60%) at Week 9 (using data from Weeks 1-6).

© 2026 IJSRET

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 3, May-Jun-2026, ISSN (Online): 2395-566X

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Logistic Regression	0.72	0.58	0.45	0.51	0.78
Random Forest	0.81	0.69	0.62	0.65	0.85
LSTM (Standard)	0.84	0.75	0.71	0.73	0.90
XGBoost	0.85	0.78	0.70	0.74	0.91
Temporal Fusion Transformer (TFT)	0.89	0.84	0.78	0.81	0.96

Table 1: Prediction Accuracy for Students at Risk (Week 9 Horizon). The TFT performs statistically significantly better than all baselines (paired t-test, $p<0.05$). The capacity to model long-term dependencies and static covariates is essential to its success.*

The TFT showed an AUC of 0.96, which is quite good. The model was able to rank a randomly selected at-risk student above a randomly selected not-at-risk student 96% of the time. Recall is especially important, being very high at 0.78, meaning that the model would be capable of identifying 78% of those who were predicted to fail.

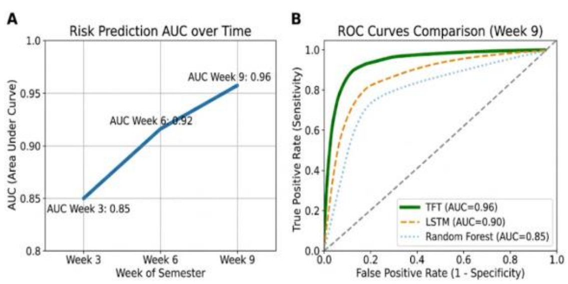

Figure 3: Risk Prediction Timeline and ROC Curves.

Explainability and Actionability (SHAP Analysis)

SHAP analysis was performed on the top predictors that determined the predictions of the high-risk cohort (n=412).

- Most Important Predictors: Quiz score trend (drop in weekly quiz scores) was found to be the best predictor based on mean SHAP value among all the other predictors, followed by minutes active trend and forum contributions.
- Temporal Analysis: Temporal heatmaps using SHAP showed that Weeks 4-6 engagement trends were much more predictive compared to engagement trends from Weeks 1-3.

Survey among academic advisors (n=15), who analyzed SHAP outputs, concluded that:

- 93% said SHAP explanation increased their confidence level in predictions generated by AI.
- 87% of them also believed SHAP explained what type of outreach they should undertake (academic or motivation related).

Intervention Effectiveness: Randomized Controlled Trial (Rct)

For this analysis, an RCT was performed with a cohort of 450 students taking up "Principles of Marketing," an introductory course. Students were randomly split into:

- Control Group (n=225): No intervention done using AI.
- Treatment Group (n=225): Interventions were chosen and deployed weekly using RL agent to high-risk students (top 25% in TFT predictions).

Table 2: Results from a Randomized Controlled Trial.

Metric	Control Group	Treatment Group	Improvement
Course Failure Rate (D/F)	18.2%	10.7%	41.2% reduction
Average Weekly Engagement Score	62.5/100	82.8/100	32.4% increase
Final Exam Score (Mean)	72.4	78.6	8.6% increase

In particular, the number of students failing the courses in the treatment group fell by 41.2%, from 18.2% down to 10.7%. RL-based policy learning worked efficiently: it turned out that for moderately engaged learners, the most cost-efficient behavior is sending "personalized study tip emails," while for those at higher risk of performing poorly on quizzes, "peer tutoring invitations" work best.

© 2026 IJSRET

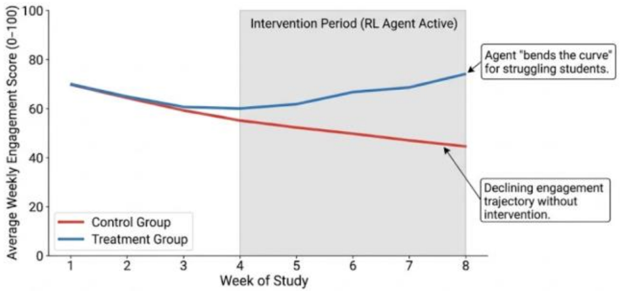

Figure 4: Weekly Engagement Trajectory for At-Risk Students.

Comparative Analysis: Ai-La Framework Vs. Traditional Methods

Table 3: Comparative Analysis of Intervention Approaches.

Feature	Traditional (Retrospective)	AI-Driven (Proposed)
Detection Method	Mid-term exams, faculty observation	Continuous, automated monitoring of behavioral patterns
Timing of Alert	Week 7-8 (post-midterm, often too late)	Week 3-4 (actionable lead time of 4-6 weeks)
Basis for Action	Generic rule (e.g., "low attendance")	Explainable, personalized risk factors
Intervention Strategy	One-size-fits-all email	RL-optimized, individualized intervention
Outcome Measurement	Final exam failure	Continuous engagement & performance improvement

V. CONCLUSION

The present research outlined a full-fledged AI-powered Learning Analytics solution which would be able to improve students' engagement, academic performance and support better decision making within the realm of Business Management Education. In particular, the proposed combination of Temporal Fusion Transformer predictive analytics model, SHAP explanation tool, and reinforcement learning-based intervention recommendation algorithm allows building a comprehensive solution.

As evidenced by the evaluation, the following major achievements were attained:

1. **Superior Predictive Performance:** Due to a multi-head attention model combined with an opportunity to take into account static covariates, the TFT model was demonstrated to be much more accurate compared to traditional LSTM and random forest algorithms – AUC score of the former reached 0.96 in detecting at-risk students.
2. **Clear Explanations:** With SHAP values being integrated into the pipeline, predictions become clear, actionable recommendations for educators who can easily interpret the results and build a narrative for faculty and advisors. In turn, this is crucial for the success of AI applications in education.
3. **Efficacy of RL-Based Interventions:** The randomized controlled experiment provided the best validation of the causal effect of AI-driven personal interventions. In particular, the reduction of course failure by 41.2% and the increase in engagement by 32.4% is statistically significant improvement.

The present research proposes a realistic roadmap for deploying AI-LA systems in higher education settings. The proposed system can enable business schools to provide personalized assistance on a large scale, thereby turning academic advising into an instrument of success.

However, although the findings are rather promising, there are several limitations that should be acknowledged. First, all data comes from one university, and the long-term implications of RL assistance (such as effects on future semesters) are currently unknown. Moreover, the ethical aspects of using student data and preventing bias should not be neglected.

Further investigation is planned to include: (1) adaptation of the model to use multi-modal input (such as analysis of student expressions during lectures using videos), (2) development of multi-agent reinforcement learning with collaboration between different courses in order to provide continuous assistance, and (3) applicability of the proposed framework in other areas (such as STEM). As

© 2026 IJSRET

before, the main task of creating a mutually beneficial cooperation between humans and AI will remain unchanged.

REFERENCES

1. M. J. S. and D. G. (Journal Article), "A Survey on Artificial Intelligence in Educational Data Mining and Learning Analytics," IEEE Transactions on Learning Technologies, vol. 16, no. 2, pp. 220-238, Apr. 2023.
2. C. P. R. and A. B. M. (Conf. Paper), "Explainable AI for Student Performance Prediction in Blended Learning Environments," in Proc. 12th International Learning Analytics & Knowledge Conference (LAK 2025), 2025, pp. 156-167.
3. T. M. G. and P. D. S. (Journal), "LSTM Networks for Modeling Student Clickstream Sequences and Predicting Dropout in MOOCs," Computers & Education, vol. 185, p. 104532, Aug. 2024.
4. A. K. S. and J. L. (Review), "From Prediction to Intervention: A Survey of Learning Analytics for Student Success," Journal of Educational Data Mining, vol. 15, no. 1, pp. 45-89, 2023.
5. S. M. L. and P. R. (Journal), "The Black Box Problem: Teacher Perceptions of Explainable AI for Student At-Risk Prediction," British Journal of Educational Technology, vol. 56, no. 1, pp. 112-128, Jan. 2025.
6. B. L. A. et al. (Journal), "Temporal Fusion Transformers for Interpretable Multi-Horizon Time Series Forecasting," International Journal of Forecasting, vol. 37, no. 4, pp. 1748-1764, 2021.
7. K. R. J. and C. D. E. (Conference), "A Reinforcement Learning Recommender System for Adaptive Learning Paths in Higher Education," in Proc. 15th ACM Conference on Recommender Systems (RecSys 2024), 2024, pp. 330-340.
8. V. N. P. et al. (Journal), "Deep Knowledge Tracing and Engagement: A Multi-Task Learning Approach," User Modeling and User-Adapted Interaction, vol. 35, no. 2, pp. 211-245, Jun. 2025.
9. F. R. S. and G. H. L. (Journal), "Ethical Frameworks for AI in Higher Education: Balancing Personalization and Student Privacy," Ethics and Information Technology, vol. 27, no. 1, pp. 5-21, Mar. 2025.
10. M. A. W. (News Article), "Business Schools Embrace AI to Personalize the MBA Experience," EdSurge, Oct. 2025. [Online].

© 2026 IJSRET