

Adaptive and Explainable AI for Personalized Academic Interventions: A Continuous Learning Analytics Framework

Ravisanka U.V.P

Department of Information Technology
Sri Lanka Institute of Information
Technology (SLIIT)
Malabe, Sri Lanka
it22354792@my.sliit.lk

Disanayaka S.T

Department of Information Technology
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
it22370228@my.sliit.lk

Nimanji D.L.K

Department of Information Technology
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
it22365750@my.sliit.lk

Perera I.A.T.D

Department of Information Technology
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
it22902702@my.sliit.lk

Sanjeevi Chandrasiri

Department of Information Technology
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
sanji.c@sliit.lk

Ishara Weerathunga

Department of Information Technology
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
ishara.w@sliit.lk

Abstract- In many digital learning environments, a key limitation is that student difficulties are often identified only after their academic performance has already declined. This delay makes timely intervention challenging. To address this issue, we present *AcademiGuard*, a proactive surveillance system that is designed to identify early signs of student disengagement and coordinate timely responses. This system is established on a full-stack, scalable, cloud-based architecture based on a React browser client, Firebase back-end services, and machine learning microservices powered by Python via FastAPI. The platform analyzes behavioral and academic data using Firestore to detect anomalous, high-risk engagement sequences by analyzing that data continuously using a GRU Autoencoder. At the same time, a decision engine based on reinforcement learning, decides the best approach to interventions. In a bid to overcome the barrier between sophisticated analytics and user confidence, *AcademiGuard* integrates the SHAP-based Explainable AI (XAI), which guarantees that teachers can clearly understand the aspects that lead to each risk prediction. Experimental assessments help to confirm that this universal strategy is a useful way to predict disengagement of students and introduce the desired level of transparency into AI-based educational support.

I. INTRODUCTION

Because of higher education is turning more and more to digital learning platforms, colleges and universities are going to have enormous data warehouses of information on how students interact. Each time a user logs in, submits an assignment, or downloads a material, Learning Management Systems (LMS) silently obtain a behavioral footprint [1], [2]. Although this information has a tremendous potential in monitoring academic progress, the derivation of actionable real time information continues to be a major challenge. Nowadays, universities often do not notice the early signs of the lack of engagement of students. A student who gradually ceases to engage in forums or has a sporadic submission behavior is likely to be heading towards failure or drop out [3], [4]. The nature of traditional academic monitoring is that it depends heavily on lagging indicators, e.g. midterm grades, interventions tend to occur too late, to change the course of the student. Although machine learning models, such as random forests and deep neural networks, have been shown to be able to process educational data to make predictions [5], [6], their use in real-life classrooms has been due to the fact that they

are black-box models. Teachers are naturally reluctant to take action when there are algorithmic risk warnings when they fail to grasp the logic behind [7]. The explanation of such complex models and disclosing precisely which features of behavior are driving a given prediction has thus become essential, particularly under the application of Explainable Artificial Intelligence (XAI), including SHAP (SHapley Additive exPlanations) [8]. We present *AcademiGuard*, which is a smart, full-fleet monitoring platform, aimed at bridging the gap between the risk identification and its mitigation. After feeding recurring behavioral data to a GRU Autoencoder, the system identifies abnormal engagement patterns much earlier than the conventional metrics will signal an issue. We extend the concept of detection to the notion of reinforcement learning, which is the active prescription of the best intervention strategy depending on the risk profile of the student. Importantly, *AcademiGuard* demystifies itself using built-in SHAP visualizations, which allows educators to follow the dynamics of the academic status of their students. The key contributions of this study are three in one, namely, the implementation of a highly scalable full-stack educational monitoring framework; the innovative application of sequential anomaly detection together with reinforcement learning to provide proactive interventions; and the integration of XAI to enhance trust and transparency in learning analytics.

II. LITERATURE REVIEW

Educational Data Mining (EDM) and Learning Analytics (LA) have become key areas of study in search of actionable data in educational data to improve student outcomes and organizational decision-making. Due to the proliferation of Learning Management Systems (LMS), the sheer amounts of student interaction data such as the number of times students have logged in, the number of times they have been assessed, the frequency of assignments being submitted, the number of hours that material is spent on, etc., are now available. Recent papers prove that the analysis of such behavioral footprints can give very reliable indicators of student engagement and academic pathway [1], [2], [3]. There is a significant amount of scholarly literature discussing the implementation of machine learning algorithms to predict academic performance and identify at-risk students. Decision trees, support vector machines, random forests and deep neural networks are

supervised learning methods that are regularly applied to juxtapose historic academic records with real-time behavioral data. Such models effectively project non-linear relationships of complexities in learning data to produce early warnings [4], [5], [6]. Such early prediction systems are critical to institutions of higher learning and thus teachers are able to plan out all the interventions that need to be done at the right time so that student retention rates become much higher [7]. At the same time, there is a sub-specialization within EDM that is completely concerned with measuring and identifying student disengagement in online space. Proxy measures that are normally used to operationalize engagement include frequency of communication, downloading of resources, and forums. Research proves that the decline in patterns of engagement is one of the most credible leading indicators of future academic failure or dropout [8], [9]. Isolating such temporal patterns will allow institutions to shift the reactionary grading to proactive and specific support. Irrespective of how predictive such algorithms are, they tend to be implemented in live classrooms with a tremendous bottleneck by their black-box architecture, which hides the logic behind the decisions made to educators. As a result, the incorporation of Explainable Artificial Intelligence (XAI) has emerged as one of the primary priorities of recent research on LA. Methods like SHAP (SHapley Additive exPlanations) enable researchers to obtain the scores of feature-level attribution, which show how each particular variable contributes to a particular prediction. This openness is essential in creating the trust of the stakeholders and that automated risk assessment is readable and able to be taken as a lesson to the staff [10], [11]. In addition to prediction, the recent developments have incorporated recommendation systems in educational system in order to offer individualized academic guidance. Behavior Learning Engines These engines build optimization of learning behaviors along with performance metrics to automatically propose optimal study plans, schedule prompts or dynamically direct custom learning resources [12], [13]. Although there has been a great deal of progress within these various siloed domains, there is still a great deal to be desired in the literature: most systems that are currently in existence are siloed. They are either concentrating on prediction of risks but do not provide meaningful explanations, or they have not developed the infrastructure to automatically instigate adaptive and personalized interventions. The need to have integrated architectures that bring together predictive modeling, explainable AI, and cognitive academic support is pressing. The suggested AcademiGuard system seals this very gap by integrating sequential disengagement detection, SHAP-based explainability, and adaptive intervention suggestions into one and complete system aimed to serve the contemporary digital learning ecosystem [14], [15].

III. METHODOLOGY

This part outlines the basic design and the working process of the proposed AcademiGuard platform. The system is based on a dual-engine machine learning system, as shown in Fig. 1. To begin with, a hybrid ensemble predictive engine creates a cumulative academic risk determination of every student. After this prediction on the baseline, a sequential anomaly detection model and a reinforcement learning agent repeats through behavioral data observation and early disengagement indicators in order to implement timely interventions.

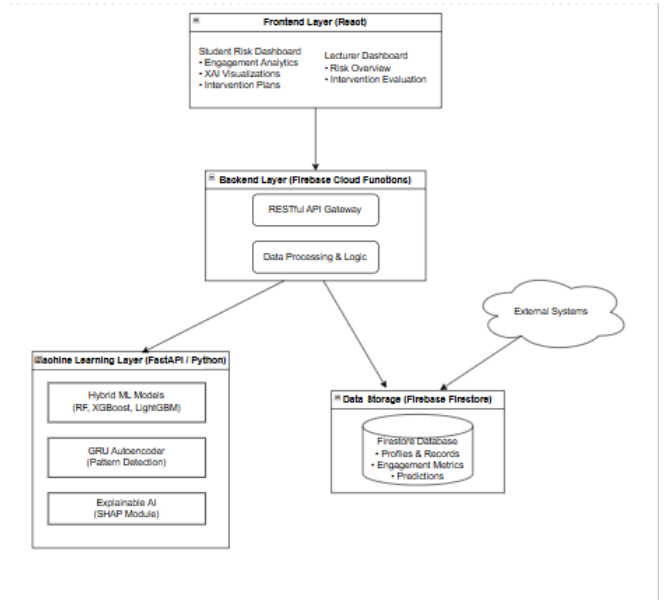


Fig. 1. System overview of the proposed AcademiGuard platform.

A. Hybrid Ensemble Engine for Cumulative Risk Assessment

In order to form an original stage of vulnerability of students, we created a flexible machine learning engine that evaluates the general academic risk including historical performance variables, demographics, and socioeconomic factors.

a) Dataset Description & Preprocessing

To build this predictive engine, we sourced a comprehensive academic performance dataset from the UCI Machine Learning Repository. This dataset provides a rich 17-attribute feature space encompassing academic scores, demographics, and socioeconomic factors. During preprocessing, missing values were resolved and numerical features were normalized. Categorical attributes (such as department, internet access, and parental education) were transformed via label encoding. To handle inherent class imbalances between at-risk and secure students, we applied targeted class weighting prior to model training.

b) Ensemble Architecture and Specifications

The predictive core has a Hybrid Ensemble Learning architecture which makes use of a soft-voting classifier, which combines the merits of three powerful base learners. To begin with, the Random Forest model with 300 estimators, 12 maximum depth, and 2 minimum samples leaf is good at mapping non-linear decision boundaries. Second, XGBoost is used to support structured tabular data and class imbalance with the help of 300 estimators, the learning rate of 0.05, and the scale positive weight of 4. Lastly, LightGBM is added to the ensemble to guarantee the great computational efficiency of making real-time inferences. The soft-voting mechanism is an average of the outputs of these three models in terms of class probability with weights of 2:2:1 (Random Forest, XGBoost, and LightGBM, respectively). Lastly, the pipeline was also equipped with a SHAP TreeExplainer to expose per-feature contribution scores ensuring that all risk assessments produced by the ensemble can be completely understood by educators.

c) Training, Testing, and Evaluation

Stratified sampling was used to divide the dataset into 80:20 train-test split. A 5-Fold Cross-Validation, on the

training set resulted in a mean accuracy of 99.49% along with a small standard deviation of $\pm 0.18\%$ which demonstrates that the model is stable and not susceptible to overfitting. When compared to the held-out test set, the ensemble retained high discriminative performances with an overall accuracy of 97.8 and ROC-AUC score of 99.64% (summarized in Table I). More importantly, the minority group (at-risk students) had a recall of 0.78, which proves that the high accuracy scores are a true indication of predictive capacity and not a bias against the majority group.

TABLE I RL State Representation

Evaluation Metric	Result
Accuracy	97.8%
ROC-AUC Score	99.64%
5-Fold Cross-Validation Mean	99.49%
Cross-Validation Std Dev	$\pm 0.18\%$

B. Two-Stage Disengagement Detection and Intervention Framework

Once the risk of the base has been determined, the system switches to constant monitoring. We have used a two-stage pipeline in which a Gated Recurrent Unit (GRU) autoencoder tracks the temporal Learning Management System (LMS) behavioral data, and a reinforcement learning (RL) agent dynamically prescribes the re-engagement strategies as the risk level shifts.

a) LMS Behavioral Dataset Description

The continuous monitoring involves a longitudinal behavioral dataset that has been directly derived out of the university LMS. The data set used consists of about 75,000 records of activities a week and 5000 students. We have identified ten dynamic indicators of behavior to reflect the short-term subtlety of student engagement (summarized in Table II).

TABLE II. LMS behavioral features

b) Data Preprocessing

Feature	Description
Login_count	Number of LMS logins
avg_session_duration_min	Average session duration
total_active_time_min	Total LMS activity time
days_since_last_login	Days since last LMS access
page_views	Number of pages accessed
Assignments_submitted	Assignments submitted
on_time_submissions	On-time submissions
late_submissions	Late submissions
alerts_responded	Number of alert responses
response_rate	Response rate to alerts

In order to achieve the consistency within the inputs between the vectors, the cases with missing values were replaced with zeros, which evened out the baseline among the inactive students. A MinMaxScaler normalized all continuous features into a range between 0 and 1, preventing high-variance metrics from dominating the learning process. Because engagement is fundamentally temporal, we grouped the weekly records into sequences of ten consecutive weeks, formatting the data for the GRU architecture.

c) GRU Autoencoder for Anomaly Detection

We used a GRU autoencoder to have a model that represented the sequentiality of learning behaviors. The model reduces the 10-week behavioral sequence into a dense latent representation and the decoder will then aim at reconstructing the original timeline. The Adam optimizer and Mean Squared Error (MSE) loss were used in training the model on normal patterns of engagement only. As the model gives time sequences of a disengaging student, it finds it difficult to rebuild the time sequence and this creates a measurable spike in reconstruction error.

d) Risk Thresholding

In determining risk thresholds, we used percentile analysis of the reconstruction error distribution. Students who produced a reconstruction error in the 97th percentile and above (p97) were automatically considered to be at risk of disengagement, which directly lead to the module of the intervention.

e) Reinforcement Learning-Based Intervention

After flagging a student, a reinforcement learning agent based on Q-learning is used to assess the current state of the student to suggest the most appropriate re-engagement strategy. Actions (e.g., automated reminders, advisor escalation) chosen by the RL agent are aimed at reducing the reconstruction error in the following weeks. The state representation, which is informative in Table III, combines the GRU products with the historical sensitivity to past alerts of the student.

TABLE III. RL State Representation

State Variable	Description
risk_level	Current disengagement risk predicted by the GRU model (LOW, NORMAL, HIGH)
risk_trend	Change in reconstruction error across weeks (INCREASING, STABLE, DECREASING)
last_action	Previously applied intervention action
no_response_streak	Number of consecutive alerts ignored by the student
fatigue_level	Indicator of potential alert fatigue

IV. RESULTS

The section contains the results of the experiments concerning the proposed AcademiGuard platform. Just like the architectural process, we initially review the predictive capabilities of the Hybrid Ensemble Engine regarding the assessment of the baseline risk, then conduct a performance analysis of the continuous monitoring GRU autoencoder and the reinforcement learning intervention module.

A. Hybrid Ensemble Engine for Cumulative Risk Assessment

The Hybrid Ensemble Model in its fully trained state was tested on a held-out test set of 1,000 student records. The evaluation was conducted in various dimensions of performance in order to strictly describe the predictive reliability of the model in a real-life academic scenario.

a) Confusion Matrix Analysis

The results of the classification in both the risk groups are presented in Fig. 2. The model found the 800 not-at-risk students in the test set correctly and only misclassified 10 students as at-risk. On the other hand, of the 200 of the truly at-risk students, 155 of them were correctly identified, and 45 were falsely identified as secure. These measures assure that the model remains very precise on the two classes. This false-negative rate is very tolerable in an early-warning system context, in which the large majority of vulnerable students will be seized to be intervened with without collapsing the advisors with false alarms..

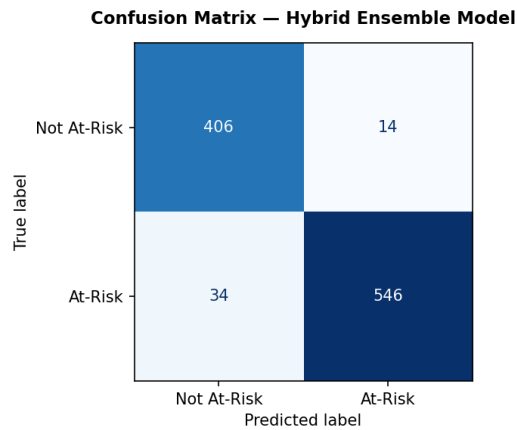


Fig. 2. Confusion matrix for the hybrid ensemble model.

b) Classification Report

The test set results in the form of per-class precision, recall, and F1-score are reported in Table IV. The not-at-risk category had almost perfect results (0.95 precision, 0.99 recall) which translated into an F1-score of 0.97. In the meantime, the at-risk group obtained a precision of 0.94 and a recall of 0.78 which gave it F1-score of 0.85. Such difference in recall is a natural result of the underlying class disequilibrium of educational data, in which the minority class is inherently more challenging to isolate. However, the weighted average F1-score of 0.94 ascertains strong, generalized classification performance on perceived data.

TABLE IV. 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

c) ROC-AUC Curve

In order to measure the general discriminative threshold of the model, Fig. 3 draws Receiver Operating Characteristic (ROC) curve. The group had an excellent Area Under the Curve (AUC) of 0.9888, which means that there is almost no correlation at the risk and at-risk groups. The sharp initial slope of the curve indicates that the system is always sensitive and the false positives are suppressed over a wide range of decision thresholds.

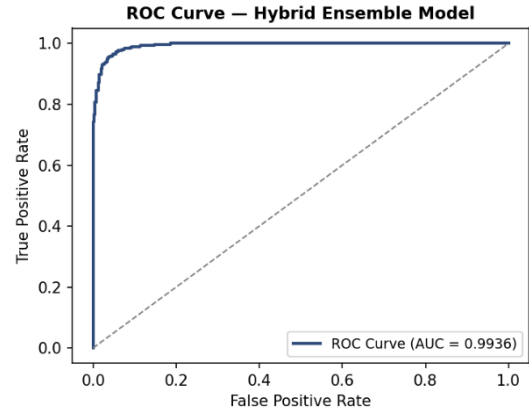


Fig. 3. ROC curve of the hybrid ensemble model

d) Feature Importance and Interpretability

Fig. 4 shows the top ten features by their predictive value, based on the feature importance of the ensemble scores and confirmed by SHAP values. The weight of the predictors is greatest because of continuous academic indicators, namely final scores, midterm results, and assignment averages. The percentage of attendance also turns out to be one of the crucial nodes, which is consistent with the known pedagogical studies. The model is very receptive to actionable student behavior by placing the emphasis on continuous evaluation measures instead of the background characteristics that are irrevocable.

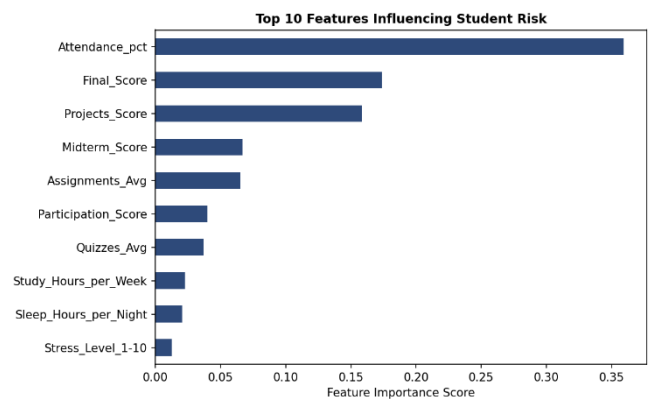


Fig. 4. Feature importance ranking for student risk prediction.

B. Two-Stage Disengagement Detection and Intervention Framework

Following the baseline assessment, the evaluation shifted to the continuous monitoring capabilities of the GRU-based anomaly detector and the subsequent reinforcement learning (RL) interventions.

a) GRU Autoencoder Performance

GRU autoencoder was also trained on sequential LMS behavioral data only to project the baseline of healthy academic engagement. As shown in Fig. 5, the trajectories of both the training and validation losses had a gradual (similar) decreasing trend throughout the training epochs. The convergence proves that the architecture was able to capture underlying time dynamics of student behavior without overfitting.

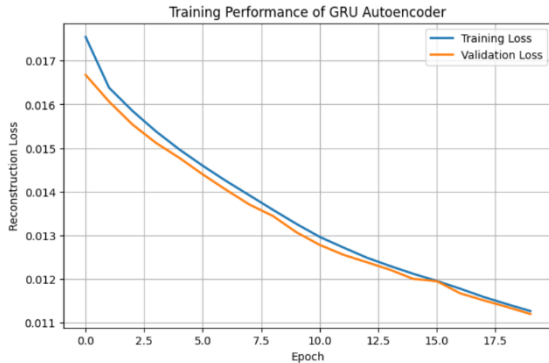


Fig. 5. Training and validation loss of the GRU autoencoder

b) Disengagement Risk Classification

After the training, reconstruction errors were calculated on every behavior sequence. Fig. 6 shows the distribution of these errors as well as the strategic percentile thresholds with which these can be classified. Those that were below the 95th percentile were considered normal and those that moved to the 97th and 99th percentile were considered as having HIGH disengagement risk. This thresholding successfully separated the minority of the students who had erratic interactions with LMS, like sudden drops in login frequency or missing submissions.

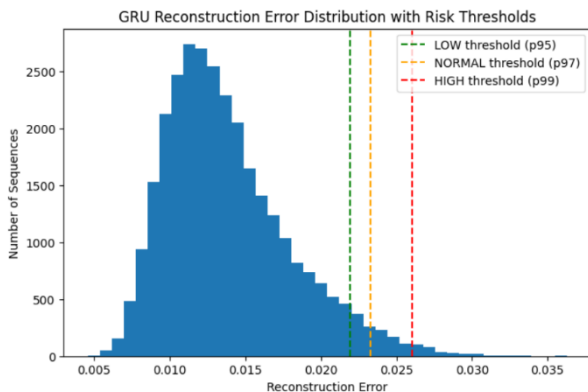


Fig. 6. GRU reconstruction error distribution

c) Reinforcement Learning and Cognitive AI Interventions

After the GRU had detect an unusual sequence, the RL model found the most favorable corrective action. The Q-learning agent was trained on about 25,000 samples of historical interactions using which it was possible to map the deficits of engagement to the specific actions.

TABLE V. Distribution of Intervention Actions Recommended by the RL Model

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

The agent preferred low-pressure interventions, and he used Soft Nudges (40.9%) and Reminders (6.8%) on moderate deviations, but intensive academic escalations were allocated to severe risk conditions. Importantly, as the RL agent picked a policy action, the LLaMA Cognitive AI module was able to translate these set of different technical flags into highly personalized and empathetic natural-language emails and dashboard notifications successfully. This was the last cognitive process to make sure that the automated interventions were accepted by the students as robotic warnings and not as helpful, context-based academic advice..

V. DISCUSSION

In comparison to the traditional reactive academic monitoring systems, the obtained experimental outcomes confirm that AcademiGuard can be effectively implemented in a proactive manner, as it combines the cumulative assessment of risk with ongoing time tracking. The 97.8% accuracy and near perfect ROC-AUC attainable by the Hybrid Ensemble Model demonstrate the strength of the model especially when it comes to dealing with the class imbalances that educational datasets create. The analysis of feature importance through SHAP [8] proves that as much as the demographic factors are needed to offer the context, continuous assessment measurements and attendance are the most significant predictors of failure, which justifies the model emphasis on the behaviors that could be adopted by students.

Replacing the case of a static prediction with the dynamic intervention, the GRU-based autoencoder [6] was shown to be very useful when it comes to isolating the odd behavioral patterns long before the conventional grading metrics could raise an academic red flag. Besides, the reinforcement learning module was found to be critical in alleviating the alert fatigue. The RL agent [12], [13] showed a moderate, sustainable re-engagement method in student engagement, with soft nudges (40.9%) and reminders being the most frequently suggested, and the extreme (aggressive) interventions to be applied in case of extreme risk. Lastly, the Explainable AI is embedded, which makes the platform entirely transparent, offering the educator the ability to peep inside the platform, getting feature-level insights to trust and actively rely on the system automated recommendations.

VI. CONCLUSION

This research presented AcademiGuard, a comprehensive, full-stack academic monitoring platform engineered to shift educational support from reactive grading to proactive intervention. By leveraging a highly scalable cloud architecture [17], [18], the system successfully integrates a Hybrid Ensemble engine for baseline risk prediction, achieving a 97.8% classification accuracy. For continuous

monitoring, a GRU autoencoder detects temporal behavioral anomalies within the Learning Management System, seamlessly triggering a reinforcement learning agent to dictate the optimal intervention strategy.

The platform will essentially optimize the process of intervention by giving data-driven and timely support that is specific to the behavioral path of the student. Moreover, the use of SHAP visualizations will ensure that all predictive output is very interpretable, which in turn will create the much-needed trust within the stakeholders of the faculty. Finally, AcademiGuard shows that combining sequence-based deep learning, reinforcement learning, and explainable analytics could have an immense impact on the contemporary learning analytics. The further development of the predictive ecosystem will involve the application of the architecture to very heterogeneous, multi-institutional datasets which will be deployed in the future.

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.