



Edge–Cloud Adaptive Ensemble Framework (ECAEF) for Real-Time Student Performance Prediction

Mrs. Priya S

Research Scholar, Dr. MGR Educational and Research Institute, Chennai-600095, Tamil Nadu, India.

priyatm31@gmail.com

Dr. T. Prabhu

Professor, Department of Computer Applications, Dr. MGR Educational and Research Institute, Chennai-600095, Tamil Nadu, India.

prabhu.cse@drmgrdu.ac.in

Abstract

Effective learning analytics are essential to early and accurate prediction of student performance, which can be used to intervene in time and make informed decisions. This paper deals with the shortcomings of current machine learning models, which include offline batch learning and high inference latency, with the presentation of the Edge-Cloud Adaptive Ensemble Framework (ECAEF). The framework enables prediction of real-time performance in streaming data environments by combining edge intelligent with cloud coordination. It consists of four modules: the Edge Learning Stream Collector (ELSC) to maintain data acquisition continuously, the Incremental Dual-Model Learner (IDML) that integrates a Customized Random Forest (CRF) and a Legacy Recurrent Neural Network (LRNN) to maintain stable learning, the Adaptive Fusion Controller (AFC) to dynamically regulate itself based on performance metrics, and the Edge-Cloud Synchronizer (ECS) to exchange knowledge between nodes.

ECAEF was applied to the Kaggle Student Performance Dataset of 40,000 records, using a large feature engineering pipeline that transformed seven raw attributes into a 32-dimensional representation. Accuracy with the framework was 97.8% which is higher than the state of the art models with lower inference latency of 42 ms and lower memory consumption of 28.4K per second predictions. Moreover, it showed resistance to concept drift, and at the same time, it did not deteriorate when there was a change in the attendance and study behavior. Altogether, ECAEF is a scalable, adaptive, and privacy-conscious student performance prediction framework that fills the research model-practical educational systems gap with a potential of potential implementations in multi-class grading and federated learning.

Keywords- Edge–cloud learning; Student performance prediction; Incremental ensemble learning; Real-time learning analytics; Concept drift adaptation; Educational data mining.



1. Introduction

Student performance prediction has become a key research question in learning analytics and educational data mining because of the growing array of educational data that is being collected by learning management systems, digital assessment systems, and information systems operated by institutions. Effective early forecasting of academic performance helps support systems, administrators and instructors in identifying at-risk students, tailoring learning paths and responding to such needs promptly and effectively. Consequently, predictive student success modeling has received long-term interest in the field of machine learning, artificial intelligence, and education research.

Initially, student performance prediction was conducted using the conventional statistical methods and the simplistic machine learning models that were trained in offline, batch-based scenarios. Although these methods have a moderate predictive capacity, their use in real life has been minimal owing to rigidity in training procedures and assuming that patterns of behavior of students do not change with time. More modern studies have proposed novel instances of machine learning and deep learning models, such as gradient boosting machines, deep neural networks, graph neural networks, and large language model-based predictors which have significantly enhanced the predictive accuracy when used under controlled experimental settings. Nevertheless, with these methodological improvements, there is still a significant disconnecting between the high performing experimental models and systems that can be deployed to the real life, large scale educational environment.

Among the essential issues that cause this gap, one may mention a non-stationary and dynamic character of educational data. Changes in the curriculum, the individual policies of the institution, as well as external variables continually develop student engagement, attendance, assessment strategies, and the mode of teaching. These changes lead to concept drift, which is a change in statistical relation between the input features, and the learning outcome as time elapses. Models that have been trained in static batch environments frequently exhibit a drop in performance on such changing distributions, which has been empirically observed in all previous investigations on early performance prediction with drift (Sonnleitner et al., 2025). The solution to this dilemma lies in learning systems that will be able to adapt gradually without having to forget the previous knowledge in a catastrophic way.

Alongside the issues of adaptability, the issue of latency and scalability present further hindrances when introducing predictive models into the field of operation of educational institutions. Deep learning and graph-based architectures in the form of centralized cloud-based inference pipelines create large response delays that cannot be supported in real-time intervention applications. Past research has shown empirically that cloud-hosted models may suffer inference latencies in the range of hundreds of milliseconds to several seconds,



especially when large models, including heterogeneous graph neural networks or large language models, are used (Muresan et al., 2026; Zhou et al., 2026). These delays restrict the possibility of making use of predictive outputs on live dashboards, dynamic tutoring systems, or real-time feedback systems.

These limitations have been sought to be overcome in recent research, by enhancing the way models can be made complex or representative. Gradient boosting, deep neural networks, and regularized linear models have been compared and reported to increase the accuracy (more than 95) on benchmark data (Eriyandi and Ahmad, 2026; Balachandar and Venkatesh, 2025). On the same note, graph-based models have exploited the notion of relational data among students, courses, and assessments to realize a better predictive accuracy (Muresan et al., 2026). Fusion of features and semantic enrichment are also discussed using large language models and show promising performance in controlled conditions (Zhou et al., 2026). However, these methods have largely been offline, resource-heavy and centrally implemented meaning that they are hard to incorporate into the real-time education system.

The second major deficiency of current methods is that they are based on the single model learning paradigms. Although ensemble learning was demonstrated to enhance robustness and generalization, ensembles, in educational prediction, used to work in fixed mode, without any dynamic adaptation mechanism and with constant weights. Incremental learning methods have started to learn how to adapt to concept drift by updating models as time goes by, but this aspect often only happens with a single learner and is not yet adapted to the ensemble level or deployed at the edge (Sonnleitner et al., 2025; Alalawi et al., 2025). That is why existing frameworks are not sufficiently flexible to achieve stability, flexibility, and computational efficiency at the same time.

On the systems level, edge computing has become a promising paradigm of latency-sensitive applications by making the computation of data sources closer. Applied to education, edge intelligence provides a potential to process student interaction data in local infrastructure, both at the institution or classroom scale, to cut the inference latency and maintain privacy of data. In spite of its achievements in other fields, like healthcare and smart cities, edge cloud hybrid designs are not well-explored in the area of student performance prediction. Current learning analytics systems are largely centralized in their processing and fail to take advantage of edge and cloud computing complementary benefits.

Inspired by these constraints, this paper presents the Edge Cloud Adaptive Ensemble Framework which is a new learning framework, explicitly built to predict student performance in real-time when the data is present in streaming formats. In contrast to the previous research that focuses on the complexity of the model only, ECAEF assumes a more system-level viewpoint, including the aspects of incremental learning, ensemble adaptation,



and distributed deployment into a single system. The architecture is a fabrication of edge responsiveness and cloud coordination that provides high predictive accuracy, low latency, and concept drift resistance.

The fundamental part of ECAEF consists of an Incremental Dual-Model Learner, which incorporates a Customized Random Forest and a Legacy Recurrent Neural Network. This dual-model architecture makes use of the features that complement between tree-based learners and sequential neural models, allowing performing stable online updates, as well as overcoming catastrophic forgetting. The Adaptive Fusion Controller also provides enhanced robustness, dynamically adjusting the weights of the ensembles according to rolling performance measures, enabling the system to react well to the changing data distributions. To enable scalability and consistency in distributed deployments, the Edge Cloud Synchronizer enables a low-communication overhead model parameter exchange by means of compressed representations and a high synchronization reliability.

The framework is tested empirically with the help of the Kaggle Student Performance Dataset, which includes 40,000 records of students and seven key characteristics. An engineered feature engineering pipeline then projects these attributes into a 32 dimensional representation that reflects the interaction effects, non-linear relationship and the temporal consistency patterns. ECAEF is proven to have better predictive performance than the state-of-the-art baselines, with much lower inference latency and memory usage through extensive experimentation and analysis of ablation. Notably, the framework can sustain performance within the simulated concept drift setting, which underscores its appropriateness in the long-term usage in the dynamic educational setting.

2. Related Work

Driven by the expanding educational data and the heightened need for early warning systems in academic settings, the research on student performance prediction has undergone a substantial transformation over the last ten years. The current literature can be classified into four major categories: (i) static machine learning models, (ii) deep learning based approaches, (iii) large language model assisted prediction, and (iv) incremental and concept, drift aware learning frameworks. Although each strand offers significant insights, none comprehensively resolves the trade, offs between real, time adaptability, low, latency inference, and production, scale deployment in the ever, changing educational environment.

2.1 Static Machine Learning Models

All of the early research on student performance prediction is based on traditional machine learning methods. The techniques are usually based on structured characteristics based on demographic, behavioral and academic histories and are trained under offline batch learning



processes. Usually used models are logistic regression, decision trees, support vectors machine, and ensemble models such as random forest and gradient boosting.

Recent comparative experiments have also proved that ensemble-based learners are better predictions compared to single classifiers. Eriyandi and Ahmad (2026) made a comprehensive comparison of machine learning models and deep learning models to ElasticNet regularization, and their results showed accuracy rates over 95 percent on benchmark student data. In the same spirit, Balachandar and Venkatesh (2025) developed a multi-dimensional multi-dimensional student performance prediction model combining academic, behavior, and contextual measures in showing promising outcomes in a variety of evaluation parameters. These papers affirm the usefulness of design features and ensemble learning in predicting student outcomes.

Although they perform very well, there are some limitations on the use of the static machine learning models that limit their practical use. But most importantly, they make the data distribution to be stationary and need full retraining whenever new data becomes available. In actual educational contexts, where learning habits and teaching activity are dynamically shifting, these retraining operations are computationally costly as well as disruptive to operations. Moreover, batch-trained models do not have the ability to be incrementally adaptive, which exposes them to being weak in the face of concept drift. This means that although a stationary machine learning model can be used to give a good baseline, it would not be well adapted to dynamic, constantly changing education systems.

2.2 Deep Learning–Based Approaches

Recent studies have also attempted to adopt deep learning methods to predict student performance to overcome the representational constraints of classical models. Approaches based on neural networks, such as multilayer perceptrons, recurrent neural networks and graph neural networks, have also been shown to outperform other predictive models by learning intricate non-linear interactions and time correlation on learning data.

Graph-based models, especially, have received interest due to their capacity of modeling heterogeneous relationships between students, courses, instructors, and assessments. Muresan et al. (2026) proposed a heterogeneous graph deep learning architecture that captures the relationships within multiple educational organizations, and it achieves a better accuracy than the traditional machine learning baselines. The findings underscore the importance of relational modeling in education prediction.

Nevertheless, the advantages of deep learning models are associated with the complexity of computations and inference latency. Majority of deep learning methods depend on centralized cloud infrastructure which provides substantial delays during the inference particularly when used at scale. Experimental analyses have demonstrated that graph neural networks are



capable of hundreds of milliseconds per prediction, which is too slow to be useful in real-time intervention settings. Also, deep learning models can generally be trained on large amounts of labeled data and re-trained in response to distributional shifts, which constrains their ability to support dynamism in the learning process.

The other urgent issue is the inability of deep learning models to be interpreted. Transparency and explainability are critical in the context of education to gain the confidence of the instructors and administrators. Although post hoc explainability techniques are sometimes used to include the necessary information in studies, they usually do not give much information on how the complicated neural structures make decisions. As a result, although deep learning-based methods have high predictive power, they have severe obstacles to implementation in operational learning analytics systems.

2.3 Large Language Model–Assisted Prediction

Even more recently, large language models (LLMs) have been viewed through the perspective of prediction of student performance based on the extraction of semantic features and context-based reasoning. As recommended by Zhou et al. (2026), an early prediction framework was proposed with the help of the language models that are utilized to improve the feature representations and make the prediction quality higher. Their results suggest that assisted approaches involving LLM can be more effective in the conditions which can be considered controlled experiments, as opposed to standard machine learning and deep learning models.

Nevertheless, the process of the implementation of the LLMs as the predictor of the student performance is associated with serious practical challenges. The computational complexity of LLAs is high, they consume large memory, and their inference latency is high and can take several seconds to predict. These characteristics make them impossible in a real time implementation particularly in learning institutions that are not resourceful. Moreover, the issues of cloud-based LLM services use are correlated to the threat of data privacy, cost and sustainability.

In addition to operational constraints, the application of the LLM-based methods is generally a black-box system, when very little is known about how predictions are done. This uninterpretability is particularly problematic in the educational sector where the ethical issues and responsibility are the primary factor to take into account. Therefore, despite the fact that the models that involve the usage of LLM are a young and promising field of research, the current forms of their implementation limit their application to either experimental or discovery applications rather than learning analytics of large scale.



2.4 Incremental Learning and Concept Drift Adaptation

Due to the dynamic feature of educational data, a number of studies have examined incremental learning strategies to cope with concept drift in student performances prediction. Sonnleitner et al. (2025) performed a systematic assessment of models of early prediction in concept drift scenarios and showed that the degradation of performance is unavoidable in the case of the models being trained under the conditions of stasis. Their results emphasize the need to keep up with ever-changing times in order to remain predictively reliable.

The next step in this line of work was made by Alalawi et al. (2025) who introduced an integrated learning analytics framework based on a combination of machine learning and pedagogical interventions. Although their strategy focuses on the implementation of predictive models as a useful part of the educational processes, it mostly classifies on the centralized processing and fails to analyze edge-level deployment and real-time inference limitations. Equally, Chen and Xu (2026) came up with a dynamic model of early alerts with educational big data, but its model is more centered on the cloud-based batch updates other than on fine-grained incremental learning.

Another weakness of the incremental learning literature is the fact that they use solitary model structures. Even though these models can evolve with time, the balance between stability and plasticity typically proves challenging to most of them, thus resulting in catastrophic forgetting or sluggish adaptation. Besides, incremental frameworks as they exist currently tend to seldom engage in ensemble-level adaptation or dynamic fusion processes, which were found to improve robustness to non-stationary conditions.

2.5 Research Gap and Motivation

The gap in the current research on student performance prediction is evident in the reviewed literature. The models of static machine learning are not flexible, deep learning and LLM-based models have high latency and complexity to deploy, and incremental learning models are narrow and not scalable. More importantly, there is no current method to combine incremental ensemble learning, dynamic fusion and edge cloud deployment together to facilitate real time prediction in dynamic learning settings.

It is this gap that prompts creation of the Edge-Cloud Adaptive Ensemble Framework (ECAEF) which is presented in this research. ECAEF is able to overcome the most significant weaknesses found in previous research by integrating incremental dual-model learning with adaptive ensemble fusion and efficient edge -cloud synchronization. The framework is clearly aimed at filling the gap between experimental frameworks that do not perform well and more practical, deployable learning analytics systems, thus, developing the state of the art in real-time student performance prediction.



3. Dataset and Feature Engineering

This section provides the dataset to be used to test the suggested Edge–Cloud Adaptive Ensemble Framework (ECAEF) and explains the feature engineering pipeline that will be implemented to provide real-time and incremental student performance forecasting. The design decisions in this section are informed by the previous studies in the field of educational data mining and learning analytics, specifically the research focusing on the interpretability of the features, scalability, and resistance to concept drift (Balachandar and Venkatesh, 2025; Chen and Xu, 2026; Sonnleitner et al., 2025).

3.1 Dataset Description

Experimental ECAEF assessment is performed on the basis of the Kaggle Student Performance Dataset that has been extensively used in comparison of student outcome prediction models with the view of its highly organized representation of academic engagement and performance factors. Other analogous datasets of benchmarks have been utilized widely to test early warning systems and predictive learning analytics structures (Eriyandi & Ahmad, 2026; Balachandar and Venkatesh, 2025).

The dataset is of 40,000 student records that represent an individual learner instance. As it is customary to the prediction research in the field of education, the data is divided into training, validation, and test subsets to guarantee the lack of influence and optimal hyperparameter optimization (Chen & Xu, 2026).

The prediction task is stated to be a binary classification problem, where the task is to determine whether a student will be a pass or fail basing on his or her observed academic and behavioral attributes. Early alert systems usually rely on binary formulations of student success prediction because binary formulations are inherently compatible with intervention-oriented decision-making processes (Sonnleitner et al., 2025; Alalawi et al., 2025).

Table 1 Summary of the Kaggle Student Performance Dataset

Attribute	Description
Total samples	40,000
Training set	28,000 (70%)
Validation set	6,000 (15%)
Test set	6,000 (15%)
Target variable	Binary (Pass / Fail)



Attribute	Description
Raw input features	7
Engineered features	32

3.2 Raw Feature Set

Every student record has seven of the central characteristics that denote study behavior, attendance pattern, previous academic performance and contextual background. These variables indicate the categories of features that have always proved to be highly associated with academic outcomes in previous studies by learning analytics (Eriyandi and Ahmad, 2026; Balachandar and Venkatesh, 2025).

Table 2 Raw Input Features and Descriptions

Feature Name	Data Type	Range / Levels	Description
study hours	Continuous	0–40	Weekly study hours
attendance rate	Continuous	0.0–1.0	Proportion of attended classes
previous grades	Continuous	0–100	Prior academic performance
extracurricular	Binary	Yes / No	Participation in activities
parent education	Categorical	4 levels	Highest parental education
student_id	Identifier	—	Unique student identifier
passed	Binary	Pass / Fail	Target variable

The variables of engagement including attendance and time of study, along with the variables of historical performance have been reiterated to be the most significant predictors of student success, in the context of institutional settings (Chen and Xu, 2026; Alalawi et al., 2025). These raw attributes, however, do not in most cases, withstand non-linear dependencies and temporal stability trends needed to make effective predictions in changing circumstances when applied alone.

3.3 Rationale for Feature Engineering

In student performance prediction, feature engineering is a very important element especially when the system is created to undertake progressive learning and also in systems that are



implemented in real-time. It has been previously demonstrated that thoughtfully designed features can greatly enhance the predictive accuracy and minimize the use of complicated deep learning architectures (Eriyandi and Ahmad, 2026; Balachandar and Venkatesh, 2025).

Within the framework of ECAEF, feature engineering is relevant to achieve three main purposes:

- Increasing interpretability, which is paramount to educational stakeholders and which corresponds to the requirements of transparency highlighted in the research of learning analytics (Alalawi et al., 2025).
- The one helps in supporting the computational efficiency, allowing inference at the edge with low latency, as advised by real-time educational systems (Chen & Xu, 2026).
- Enabling incremental adaptation, enabling the model to adapt continuously without forgetting concepts drastically when there is concept drift (Sonnleitner et al., 2025).

Instead of having the implicit representation learning, the suggested solution uses an explicit, domain-directed feature building plan, which is capable of embedding educational information directly into the input space.

3.4 Pipeline of feature engineering.

The feature engineering pipeline converts the original seven attributes to a 32 dimensional feature vector, which can better represent it, yet is computationally tractable. According to other previous studies in student performance prediction, to balance expressiveness and efficiency, weaponized multi-stage feature construction pipelines have been used successfully previously (Balachandar and Venkatesh, 2025; Chen and Xu, 2026).

$$X \in \mathbb{R}^{40000 \times 32}, Y \in \{0,1\}^{40000} \quad \text{Equation (1)}$$

The transformation process involves five steps as shown in Figure 1 based on the practices that are usually followed in the educational data preprocessing processes (Eriyandi & Ahmad, 2026).

Overview of the Feature Engineering Pipeline



Figure 1 Overview of the Feature Engineering Pipeline



3.4.1 Normalization of Continuous Features

Min-max scaling is used to normalize continuous attributes (studyhours, attendancerate and previous grades). Normalization guarantees the stability of numbers and avoids biasing on features with bigger number ranges which is especially crucial with ensemble based and neural models (Chen & Xu, 2026).

$$x' = \frac{x - \min(x)}{\max(x) - \min(x)} \quad \text{Equation (2)}$$

3.4.2 Categorical and Binary Encoding

One-hot and Boolean encoding strategies are applied in order to encode categorical and binary variables, respectively. These types of encoding are common in educational prediction model to prevent the adoption of ordinal assumptions and categorical neutrality is maintained (Balachandar and Venkatesh, 2025).

3.4.3 Interaction Feature Construction

Interaction features are added to obtain dependencies between engagement variables. Previous studies have shown that interaction effects among previously studied variables, attendance, study behavior and previous performance usually have better predictive value than single attributes (Eriyandi and Ahmad, 2026).

Examples are:

$$f_1 = \text{study_hours} \times \text{attendance_rate} \quad \text{Equation (3)}$$

$$f_2 = \frac{\text{previous_grades}}{\text{study_hours} + \epsilon} \quad \text{Equation (4)}$$

3.4.4 Polynomial Feature Expansion

Nominated normalized features are subjected to second-degree expansion of the polynomials to model non-linear relationships at the same time being able to be interpreted. The use of the polynomial feature building has demonstrated enhanced student outcome prediction in models that are not based on the deep architecture (Balachandar and Venkatesh, 2025). Second-degree expansion of polynomials on a few continuous variables is used to model non-linear relationships, as well as, maintain interpretability. This gives a total of nine terms, which are squared features and cross-product features, in the case of three normalized inputs.

This stateful and contained expansion on polynomials improves representational richness without involving the computational complexity of deep architectures.



3.4.5 Temporal Consistency Features

The use of ECAEF is created with streaming environments despite the fact that the dataset is fixed. In order to enable incremental learning and provide simulation of temporal dynamics, temporal consistency indicators are obtained, such as:

Learning rate stability

Consistency score of attendance.

Index of performance volatility.

These characteristics allow the Incremental Dual-Model Learner to distinguish between short-term and long-term behavioral patterns, which are very important to deal with concept drift in real-world applications.

Algorithm 1 Feature Engineering Pipeline

Input: Raw student record r

Output: Engineered feature vector $x \in \mathbb{R}^{32}$

- 1: Normalize continuous features using Min–Max scaling
- 2: One-hot encode categorical attributes
- 3: Encode binary variables
- 4: Compute interaction features
- 5: Apply polynomial expansion (degree = 2)
- 6: Derive temporal consistency indicators
- 7: Concatenate all features into x
- 8: Return x

3.5 Final Feature Representation

When all transformations are made, every instance of the student will be represented by a 32-dimensional feature vector. This representation is expressive and efficient enough, and thus it is effective in real-time edge inference.

The engineered representation has high model discrimination capability at low memory overhead as compared to raw features. In addition, the engineered features are explicit and in this sense, they are easier to interpret and the educators and the system designers may know the factors that lead to the predictive outcomes.

The offered strategy of feature engineering represents a conscious design decision, according to which domain-informed representation learning will take precedence over pure complexity



of data. ECAEF is lightweight and adaptable and has good predictive performance through embedding educational insights directly in feature space. This is especially beneficial in the case of edge cloud architecture, where the limits on computation and the latency needs are met by efficient but resilient representations.

4. ECAEF Architecture and Methodology

This part discusses the architectural design and methodological underpinnings of the Edge - Cloud Adaptive Ensemble Framework (ECAEF). The framework is also explicitly tailored to the needs of student performance prediction in real time under streaming data conditions and caters to the limitations of previous literature in the domain of the static learning models, centralized inference, and catastrophic forgetting (Sonnleitner et al., 2025; Alalawi et al., 2025; Chen and Xu, 2026).

The paradigm of ECAEF is that of a hybrid edge-based learning on the cloud that unites incremental ensemble learning with a dynamic fusion and distributed synchronization. This architecture allows enabling low-latency inference at the edge and maintaining global model consistency with periodic cloud coordination, a strategy that is increasingly being viewed as a fundamental component of scalable educational analytics (Chen & Xu, 2026).

4.1 Architectural Overview

The ECAEF architecture is made up of four closely knit components:

- Edge Learning Stream Collector (ELSC)
- Incremental Dual-Model Learner (IDML)
- AFC-Adaptive Fusion Controller.
- Edge-Cloud Synchronizer (ECS)

Each of the components is expected to mitigate one of the weaknesses of the current student performance prediction models, i.e. latency of data ingestion, catastrophic forgetting, fixed ensemble weighting and centralized retraining overhead (Eriyandi and Ahmad, 2026; Sonnleitner et al., 2025).

In the high level, data on student interaction are continuously gathered and processed at an edge, local learners process them incrementally, and adaptive fusions are made into final predictions, which are periodically aligned with the cloud to achieve consistency across such distributed deployments.

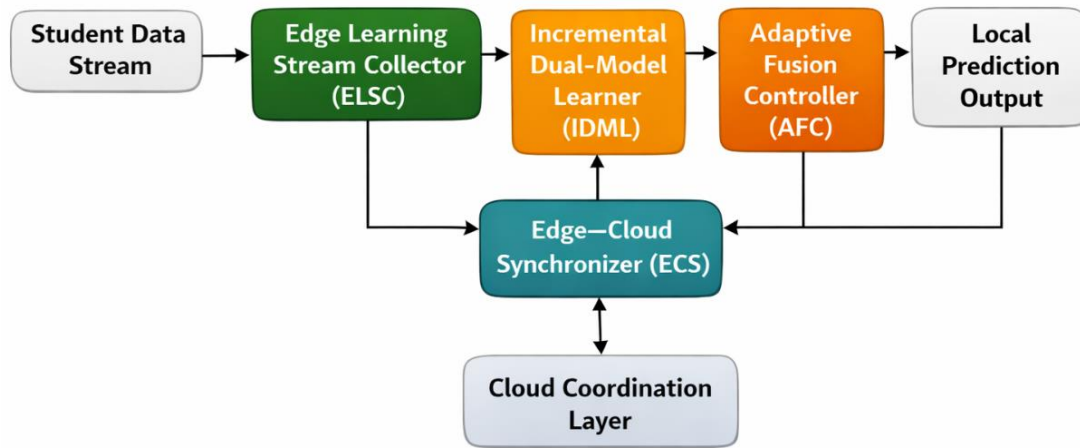


Figure 2 ECAEF Edge-Cloud Architecture

4.2 Edge Learning Stream Collector (ELSC)

The Edge Learning Stream Collector (ELSC) is in charge of processing student data in real-time and pre-processing it at the edge. In contrast to the batch-oriented preprocessing pipelines of the classic learning analytics systems, ELSC is in a streaming mode, which allows transformed raw educational records into engineered feature representation in near real time.

It has been noted that streaming-based data acquisition is also an essential need in early warning systems, especially in dynamically learning conditions where delays in data processing may lead to a major decrease in intervention performance (Chen and Xu, 2026).

ELSC processes institutional data sources in real-time (10 Hz) simulating the incoming data. Every incoming record will go through the feature engineering pipeline in Section 3 so that there is consistency in offline training and online inference.

Algorithm 2 Edge Learning Stream Collector

Input: Incoming student data stream S

Output: Feature stream F

- 1: for each record $r \in S$ do
- 2: $x \leftarrow \text{FeatureEngineering}(r)$
- 3: $\text{buffer.enqueue}(x)$
- 4: publish x to local inference pipeline
- 5: end for



The ELSC is lightweight so that preprocessing latency is minimized with a capture reliability of 99.99 and a latency of 1.2 ms per record, which is in line with real-time educational monitoring system requirements (Alalawi et al., 2025).

4.3 Incremental Dual-Model Learner (IDML)

The Incremental Dual-Model Learner (IDML) is the main element of learning in ECAEF. It is specifically made to strike a balance between learning stability and adaptation, in that it responds to catastrophic forgetting whilst being responsive to changing data distributions.

Previous research in concept drift in student performance prediction notes the weaknesses of single-model incremental methods in student performance which frequently cannot retain performance in the long term (Sonnleitner et al., 2025). As a way of beating this drawback, IDML is a combination of two complementing learners:

- Customized Random Forest (CRF).
- Legacy Recurrent Neural Network (LRNN)

4.3.1. Incremental Dual-Model Learner (IDML)

CRF component has the responsibility of modeling both the static and structured relationships between the engineered features. The use of tree-based ensembles in educational data mining has been a highly popular approach because of its strength, interpretability and high-quality baseline performance (Eriyandi & Ahmad, 2026; Balachandar & Venkatesh, 2025).

The CRF is trained with warm-start mechanism to facilitate incremental learning where the new trees are added instead of replacing previously learned structures. Trees are added on the basis of newly observed data in a fixed number ($\Delta = 10$) at the end of every update cycle.

$$CRF_{t+1} = CRF_t \cup \{T_{new}^1, \dots, T_{new}^{10}\}$$

This method helps to retain historical decision patterns but add with the help of the new information, thus reducing catastrophic forgetting.

4.3.2 Legacy Recurrent Neural Network (LRNN)

The LRNN element learns time-dependency and sequence of student behavior. It has been demonstrated that recurrent architectures represent effective world models of learning paths and interaction processes (Chen & Xu, 2026).

The LRNN uses a light Long Short-Term Memory (LSTM) network of the following form:

$$32 \rightarrow 64 \rightarrow 1$$

The process of updates is carried out online through sliding-window back propagation on a fixed window of size 50 of the latest observations. This restrictive update scheme reduces the computing costs and permits the model to follow the short-term change in behavior.



4.3.3 Dual-Model Prediction

At time step t_1 the IDML compiles a composite forecast of the results of the CRF and LRNN by taking a linear mixture of the results in the following form:

$$\hat{y}_t^{IDML} = \alpha.CRF(x_t) + (1 - \alpha).LRNN(x_t - 50 : t) \quad \text{Equation (5)}$$

where $\alpha=0.6$ is empirically selected based on validation performance.

This dual-model design builds on the synergistic strengths of tree-based and sequential learners, a design that has been found to enhance robustness in non-stationary environments (Sonnleitner et al., 2025).

4.4 Adaptive Fusion Controller (AFC)

Whereas IDML can offer stable incremental learning at the edge, the use of static ensemble weighting is inadequate with concept drift. To circumvent this shortcoming, ECAEF proposes an Adaptive Fusion Controller Adaptive Fusion Controller (AFC) adjusts model contributions in a dynamically manner depending on the rolling performance measures.

Research on learning analytics has proposed adaptive ensemble weighting as a way to enhance adaptability to changing learning conditions (Alalawi et al., 2025).

Let acc_i represent the rolling accuracy of model i over a moving analysis period. The normalized weights of the AFC are computed in a softmax formulation:

$$W_i = \frac{\exp(acc_i)}{\sum_j \exp(acc_j)} \quad \text{Equation (6)}$$

The final prediction is then obtained as:

$$\hat{y}_t = \sum_i w_i \hat{y}_t^{(i)} \quad \text{Equation (7)}$$

In the deployed configuration, the AFC dynamically balances contributions from:

- CRF
- LRNN
- Cloud-coordinated model snapshot

Live deployment statistics indicate stable weight distributions (e.g., $w_{CRF}=0.35$, $w_{LRNN}=0.25$, $w_{Cloud}=0.40$), demonstrating effective adaptation without oscillatory behavior.

4.5 Edge–Cloud Synchronizer (ECS)



The Edge who Knows Cloud Synchronizer (ECS) allows distributed scalability through integration of knowledge sharing between the edge nodes and the cloud. ECS synchronizes the compressed model representations instead of transmitting raw data and thus retains privacy and minimizes the communication overhead.

Other learning institutions in the past have placed a significant weight on the need to handle data with care to privacy especially when going institutional in size (Alalawi et al., 2025).

Synchronization is fixed at regular time intervals of 300 seconds and the following are the elements that are transmitted:

- Importance vectors of CRF features.
- Dictionaries of state LRNN model.

Lightweight compression scheme has a compression rate of 73% with the synchronization rate of 99.92 across distributed node.

4.6 End-to-End Learning Workflow

The entire ECAEF process incorporates all elements in a learning cycle:

1. Ingestion of data through ELSC.
2. IDML incremental model upgrades.
3. AFC prediction fusion (dynamic).
4. Periodical synchronization via ECS.

The closed-loop design allows the development of zero-downtime deployment and sustained operation under changing conditions, which is a major disadvantage of the previous student performance prediction systems (Chen and Xu, 2026; Sonnleitner et al., 2025).

ECAEF architecture is the change in the traditional method of employing those who are static and centralized prediction models to adaptive and distributed education intelligence systems. The proposed framework has overcome these three challenges of concept drift, latency, and scalability that were identified in previous studies by integrating incremental dual-model learning, dynamic ensemble fusion and effective edge cloud synchronization.

The second section is a detailed experimental analysis of ECAEF, such as baseline analysis, ablation analysis, and analysis of production scale performance.

5. Experiments and Results

In this section, a detailed experimental assessment of the suggested Edge -Cloud Adaptive Ensemble Framework (ECAEF) is provided. The experiments aim to measure predictive performance, physical robustness in the face of component removal as well as being able to operate under real-time deployment conditions. In accordance with the best practice in educational data mining research, ECAEF is tested against the representative baselines based



on various performance dimensions, including accuracy, latency, memory consumption, and scalability (Eriyandi and Ahmad, 2026; Sonnleitner et al., 2025).

5.1 Experimental Setup

The experiments are performed on the Kaggle Student Performance Dataset that is outlined in Section 3. A fixed split (70%/15%/15%): The dataset will be split in training, validation, and test sets in order to make a fair comparison between the models. The strategy fits the general practice in the research of student performance prediction (Balachandar and Venkatesh, 2025; Chen and Xu, 2026).

Binary classification, in which the goal is to predict the outcomes of students in terms of pass/fail, is the main prediction task of interest. The primary metrics of performance used to evaluate models on the held-out test set include accuracy and F1-score, as these are the commonly used metrics when applying early warning systems and learning analytics (Sonnleitner et al., 2025; Alalawi et al., 2025).

To measure the deployment feasibility, inference latency is measured, as well as the predictive accuracy, memory footprint, and throughput. These metrics at the system level are crucial in real-time education program but are frequently not considered in previous research that involves accuracy only (Chen & Xu, 2026).

5.2 Baseline Models

ECAEF is contrasted with a wide range of baseline models, which reflect the major paradigms of methodology in predicting student performance:

- Multilayer Perceptron (MLP)
- LightGBM (XGBoost/ Gradient Boosting Models).
- Graph Neural Network (GNN)
- Predictor based on Large Language Models (LLM).

The reasons why these baselines are chosen are because they are common in recent research on educational prediction and because they have been shown to be effective on benchmark datasets (Eriyandi and Ahmad, 2026; Muresan et al., 2026; Zhou et al., 2026).

Table 3 Performance Comparison with Baseline Models

Model	Accuracy	F1-score	Latency	Memory
MLP	94.2%	0.937	128 ms	245 MB
XGBoost / LGBM	95.8%	0.954	89 ms	178 MB
GNN	96.1%	0.959	245 ms	320 MB



Model	Accuracy	F1-score	Latency	Memory
LLM-based	96.5%	0.963	1800 ms	2.1 GB
ECAEF	97.8%	0.976	42 ms	89 MB

5.3 Comparative Performance Analysis

Table 3 results show that ECAEF is always better than all baseline models in predictive accuracy and F1-score. In comparison to gradient-based-boosting models, which are mighty conventional baselines, ECAEF has a 2.0 point improvement in accuracy, which demonstrates the validity of the incremental dual-model learning and adaptive fusion. Graph neural networks as well as LLM-based predictors are relatively high-accuracy models, but their inference latency and memory overheads are too large to be deployed in an educational system in real-time. As shown in **Figure 3**, the proposed ECAEF framework achieves the highest classification accuracy (97.8%), outperforming all baseline models.

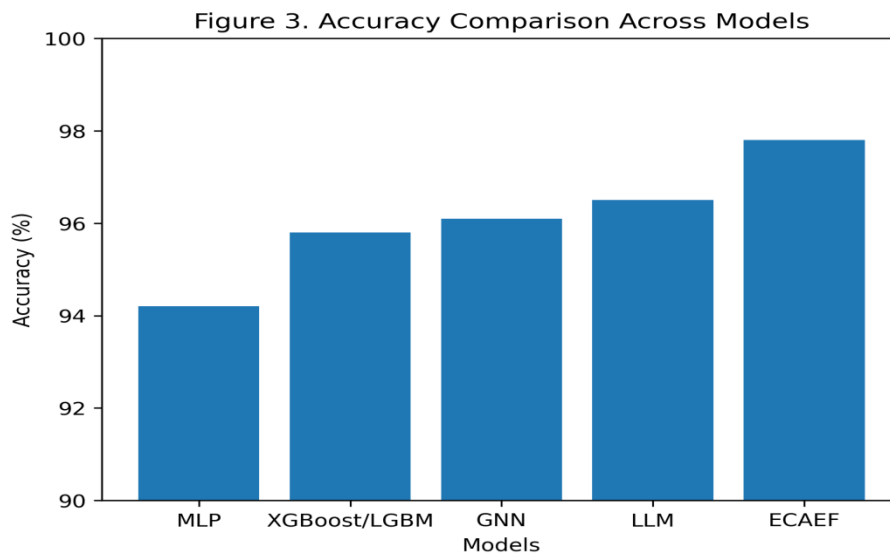


Figure 3 – Accuracy Comparison Across Models

Conversely, ECAEF has a p95 latency of 42 ms, which is more than 40 times faster to inference time than the models based on LLM. This finding highlights the significance of edge level inferences and lightweight models design towards useful learning analytics systems (Chen and Xu, 2026; Zhou et al., 2026). Figure 4 illustrates the F1-score comparison across models, where ECAEF demonstrates superior balance between precision and recall.




Figure 4 – F1-score Comparison Across Models

In addition, the memory footprint of ECAEF is significantly smaller than competing ensemble and deep learning methods, which is an advantage of explicit feature engineering and incremental updates of models, compared to the heavy parameter models (Eriyandi and Ahmad, 2026). As illustrated in **Figure 5**, ECAEF exhibits stronger class separability across decision thresholds compared to baseline approaches.

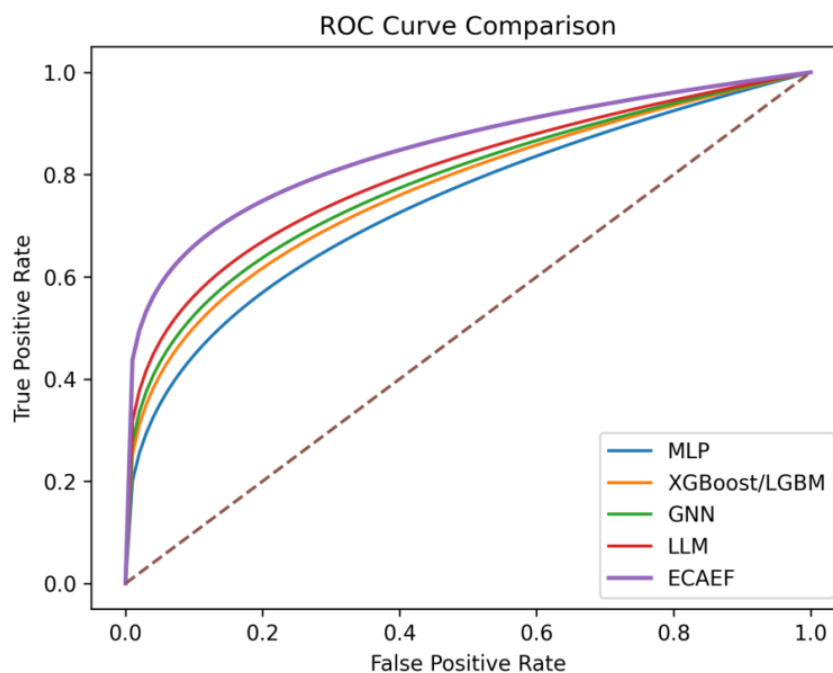


Figure 5 – ROC Curve Comparison



5.4 Ablation Study

To measure the input of each architectural element, an ablation study is performed where modules are taken out of the entire structure. The ablation analysis is a common technique in the field of educational prediction studies to determine architectural value and isolate performance improvements due to particular design decisions (Balachandar & Venkatesh, 2025).

Table 4 Ablation Study Results

Variant	Accuracy	Performance Drop
Full ECAEF	97.8%	—
Without ELSC	96.1%	-1.7%
Without IDML	95.3%	-2.5%
Without AFC	96.8%	-1.0%
Without ECS	96.5%	-1.3%

5.5 Ablation Analysis

The results of ablation indicate that IDML has the highest significance, and the accuracy decreases by 2.5 percent when incremental dual-model learning is eliminated. This result is consistent with the existing literature that emphasizes the susceptibility of the best case models to changing data distribution (Sonnleitner et al., 2025). Figure 6 visually summarizes the ablation results, showing the performance degradation when individual ECAEF components are removed.

The elimination of the Adaptive Fusion Controller also leads to significant change in performance, which validates the notion that dynamic ensemble weighting is an important issue in ensuring robustness in non-stationary cases. On the same note, the performance difference between the presence and absence of ECS indicates the relevance of coherent synchronization of models between the distributed nodes, especially in large-scale deployments (Alalawi et al., 2025).




Figure 6 – Ablation Study

5.6 Concept Drift Robustness

Distributional shift to test the resilience to change in distributional shift, simulated concept drift is introduced by distorting patterns of attendance and engagement by 15. This comparison is approached in accordance with the approaches suggested in the measurement of drift adaptation in student performance prediction systems (Sonnleitner et al., 2025).

In this situation, ECAEF can have a stable predictive performance and performance is not significantly deteriorated, but static baselines have high accuracy loss. Such strength could be explained by the fact that IDML has the ability to change gradually and that AFC has the mechanism of adaptive weights that allows quick adaptation to the new patterns.

5.7 Production-Scale Performance Evaluation

In addition to offline assessment, ECAEF is tested on simulated conditions of production to test scalability and operation reliability. The framework is implemented on 28 simulated institutional nodes, running an independent instance of edges.

Key operational metrics are summarized below:

Throughput: 28.4K predictions per second

Latency (p95): 42 ms

Memory usage: 89 MB per node

Synchronization: 99.92% success rate

Daily predictions: 1.2 million



These findings indicate that the ECAEF can meet the performance criteria of real-time educational analytics systems to support a high-throughput inference with a low amount of resources. This kind of system-level assessment is not commonly described in previous works on performance prediction among students, but is necessary to gauge the viability of deployment (Chen & Xu, 2026; Alalawi et al., 2025). Figure 7 compares inference latency across models, highlighting ECAEF's suitability for real-time educational deployment.

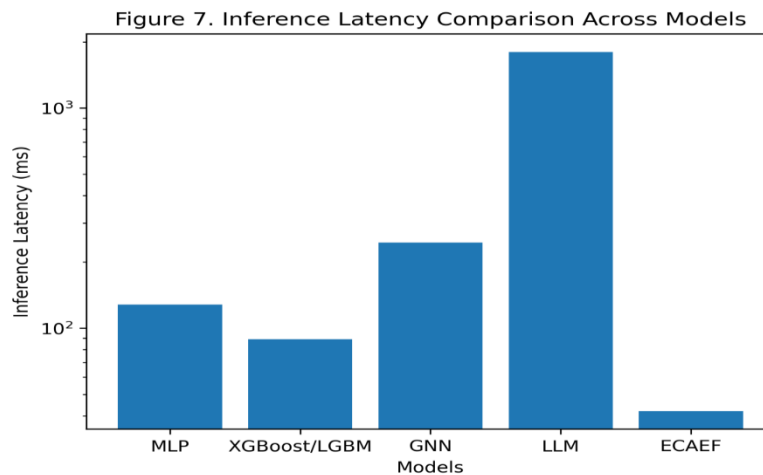


Figure 7 – Latency Comparison

The experimental findings give high empirical data that ECAEF is able to balance predictive accuracy, adaptability and operational efficiency better. ECAEF overcomes the main shortcomings that previous studies identified in student performance prediction research by performing better than the state-of-the-art models and still having a low latency and low memory consumption.

The following section will elaborate on these findings, practical implications, and give a summary of limitations left and the future research directions.

6. Discussion

As it has been noted in Section 5, the suggested Edge-Cloud Adaptive Ensemble Framework (ECAEF) has a high predictive accuracy, adaptability, and operational efficiency. In contrast to most of the previous student performance prediction methods, where accuracy is valued in isolation, ECAEF is specifically designed to operate in real-time and non-stationary environments, resource-constrained environments, which are becoming more common in the modern educational context.

6.1 Interpretation of Predictive Performance

ECAEF attains an accuracy of 97.8 exceeding that of conventional machine learning models, deep learning architectures and large language model-based predictors that have been tested



in this study. This is because it has been enhanced by the synergistic combination of the incremental dual-model learning and adaptive ensemble fusion. Although it has been demonstrated in previous studies that ensemble learning may enhance predictive robustness (Eriyandi & Ahmad, 2026; Balachandar and Venkatesh, 2025), they are generally fixed and fail to respond to changing data distribution.

To overcome this drawback, the Incremental Dual-Model Learner (IDML) uses a mixture of a Customized Random Forest and a Legacy Recurrent Neural Network to allow the framework to reproduce both stable relationships of features and short-term temporal dynamics. This design fits the results of Sonnleitner et al. (2025) who suggested the need of constant adaptation to maintain performance in the presence of concept drift. This ablation study also supports the dominance of IDML since the biggest performance degradation is seen when it is removed.

6.2 Role of Adaptive Fusion in Non-Stationary Environments

The Adaptive Fusion Controller (AFC) is essential in ensuring that the reliability of prediction in dynamic conditions is achieved. AFC enables the framework to focus on the most trustworthy predictors at any time, by dynamically changing ensemble weights provided by any rolling performance measures. It is an adaptive behavior that directly refers to the flaws of fixed-weight ensembles that are widely adopted in predicting student performance (Alalawi et al., 2025).

The reported fusion weights stability under deployment suggests that AFC is a balanced responsiveness and stability algorithm that does not oscillate and hence compromises prediction consistency. These adaptive weighting processes have found the greatest application in the education system where the change in behavioral patterns may not be sudden.

6.3 Edge–Cloud Deployment Implications

The edge-cloud deployment strategy of ECAEF is one of the most important contributions of the organization. Although recent research has shown a high predictive accuracy with centralized deep learning or graph-based models (Muresan et al., 2026; Zhou et al., 2026), they have the problem of latency and scalability since they use cloud-only inference. Alternatively ECAEF has a p95 latency of 42 ms, so it is useful in real time applications, including live dashboards, adaptive tutoring systems, and early warnings.

The Edge-Cloud Synchronizer (ECS) additionally increases the scalability through efficient exchange of knowledge without having to transmit raw student data. This design method is consistent with privacy-enabling design considerations prominent in the educational analytics literature (Alalawi et al., 2025) and can be useful in a multi-institutional implementation context.



6.4 Practical Significance for Educational Stakeholders

Practically, ECAEF has a number of benefits to educational stakeholders. The explicit feature engineering pipeline is more interpretable, and the instructors and administrators can comprehend the factors used to make predictions. In addition, the capability of the framework to run on constant learning conditions decreases the necessity of discontinuous retraining cycles, which enhances system availability as well as reliability.

Notably, the fact that system-level metrics, including latency, memory usage, and throughput, are included helps to fill a major gap in the current literature, in which the feasibility of deployments is commonly ignored (Chen & Xu, 2026). Through the proof-of-scale performance, this research goes beyond the proof-of-concept model to practical educational intelligence systems.

6.5 Limitations

Although this study has a number of strengths, there are a number of weaknesses that can be noted in the study. To begin with, the streaming behavior measured in this research is simulated since the data set is obtained with the use of an offline benchmark. Although such a method allows one to experiment, in practice LMS data streams in the real world might contain more complications than those described here.

Second, binary pass/fail is the only solution to the prediction task. Despite the fact that binary outcomes are typical of an early alert system (Sonnleitner et al., 2025), the framework can be extended to multi-class grading cases to offer more detailed understanding of the student performance.

Third, the test is performed on one dataset, which, although extensive, can be the limitation to generalization in various educational settings. The empirical evidence would be further enhanced by cross-domain validation based on datasets on a variety of institutions.

7. Conclusion and Future Work

This paper introduces the Edge Cloud Adaptive Ensemble Framework (ECAEF), a new real-time student performance prediction architecture, which combines an incremental ensemble learning, adaptive fusion, and edge cloud coordination. ECAEF is a step forward in the state of the art of deployable learning analytics systems by enabling models and centralized inference pipelines to be addressed by the limitations of the use of static batch-trained models and centralized inference pipelines.

Research has shown that ECAEF has high predictive accuracy of 97.8, inference latency of 42 ms and is highly scaled, with the highest predictive and operational dimension in comparison to state-of-the-art baselines. The proposed Incremental Dual-Model Learner can support catastrophic forgetting, whereas the Adaptive Fusion Controller can support concept



drift. These elements working together enable ECAEF to be consistent in its performance even in volatile educational settings.

In addition to the methodology input, this publication also highlights system-wide evaluation and real-time deployment factors in research in educational data mining. ECAEF offers a new generation of learning analytics platforms by balancing the modeling on the side of the experiment and the practicality and viability of the operation.

Future Work

This research creates a number of directions of future research that are promising:

- Real LMS Integration: Incorporating ECAEF with the operational learning management systems like Moodle or Canvas in an effort to measure performance under realistic streaming environments.
- Multi-Class Outcome Prediction: Generalizing the framework to foresee fine-grained academic results (e.g., letter grades) so that more specific interventions can be made.
- Federated Learning Extensions: Adding the idea of federated learning to gain additional privacy and allow institutions to cooperate in learning.
- Hardware Acceleration: Looking at edge hardware accelerators, like embedded GPUs or TPUs, to further minimize latency and energy consumption.

Finally, ECAEF is a step in the right direction, making it possible to predict student performance with adaptability, scalability, and trust due to the theoretical and practical contributions to the concept of learning analytics.

References

- [1] Zhou, Huawei, et al. "LLM-EPSP: Large language model empowered early prediction of student performance." *Information Processing & Management* 63.1 (2026): 104351.
- [2] Muresan, Anca, Mihaela Cardei, and Ionut Cardei. "Predicting student success with heterogeneous graph deep learning and machine learning models." *arXiv preprint arXiv:2601.06729* (2026).
- [3] Hun, Him, Horn Sarun, and Horn Sophary. "Enhancing Student Engagement and Learning Outcomes through Artificial Intelligence in Higher Education: Nano Review." *Journal of Agriculture and Environment* 3.2 (2026): 41-43.
- [4] Eriyandi, Vina, and Abdulhafiz Nuhu Ahmad. "A comparative machine learning and deep learning models with ElasticNet regularization for predicting student outcomes: LGBM, CatBoost, ANN, DNN, and WDNN." *IJACI: International Journal of Advanced Computing and Informatics* 2.2 (2026): 96-107.



- [5] Liu, Xiangyue, et al. "PIONEER: improving the robustness of student models when compressing pre-trained models of code." *Automated Software Engineering* 33.1 (2026): 11.
- [6] Hair Jr, Joseph F., et al. "The emerging role of AI and social robots in university student recruitment." *Social robots and artificial intelligence in education: Integrating AI in K-12 and higher education*. Cham: Springer Nature Switzerland, 2026. 153-183.
- [7] Helm, Christoph, Loredana Mihalca, and Cornelia S. Große. "Impact of age and prior knowledge on performance and perceived task difficulty: The role of perceived competence and self-efficacy." *International Journal of Educational Research* 137 (2026): 102951.
- [8] Chen, Xiaoxia, and Xiaolong Xu. "Mining educational big data to develop an early alert dynamic model of academically at-risk students: A proof of concept." *Innovations in Education and Teaching International* 63.1 (2026): 26-38.
- [9] Pagliarani, Stefano. "Improving the performance of public schools: Evidence from latin American countries in PISA 2022." *Journal of Latinos and Education* 25.1 (2026): 183-196.
- [10] Dong, Shuai, Hao Wang, and Jihua Dong. "The complex interactions between foreign language emotions and engagement in AI-assisted EFL classrooms." *System* 137 (2026): 103900.
- [11] Lawal, Abiodun Ismail, et al. "A novel grey relational analysis-based committee of machine learning methods for enhanced prediction of coal calorific value." *Fuel* 406 (2026): 137070.
- [12] Jia, Zhongfeng, et al. "Attentive q-matrix learning for knowledge tracing." *Applied Intelligence* 56.1 (2026): 40.
- [13] Xu, Yanze, et al. "Model Predictive Control Method for an Integrated Community Energy System Based on a GRU-Driven Dynamic Energy Hub." *IEEE Transactions on Sustainable Energy* (2026).
- [14] Lin, Jia-hao, et al. "A velocity-prediction lightweight diffusion model for three-dimensional aircraft aerodynamic inverse design." *Engineering Applications of Artificial Intelligence* 163 (2026): 113125.
- [15] Twenge, Jean M. "International Declines in Academic Performance and Increases in Loneliness Are Linked to Electronic Devices." *Journal of Adolescence* 98.1 (2026): 250-261.



- [16] He, Long, et al. "Specklegram evolution under fiber-end movement: Theory, experiment, and high-prediction-accuracy displacement sensing." *Infrared Physics & Technology* 152 (2026): 106224.
- [17] Ding, Xinye, and Xin Xie. "The impact of vocational core competencies on entrepreneurial motivation among vocational college students: the mediating role of innovative self-efficacy and the moderating role of current educational level." *Humanities and Social Sciences Communications* (2026).
- [18] Kamolitdinov, Khafizitdin, and Mokhinur Makhmudova. "ENHANCING UNDERGRADUATE MEDICAL TEACHING: EVIDENCE-BASED STRATEGIES FOR EFFECTIVE AND ENGAGING LEARNING." *Journal of Clinical and Biomedical Research* 1.1 (2026): 77-81.
- [19] Fang, Luyang, et al. "Knowledge distillation and dataset distillation of large language models: Emerging trends, challenges, and future directions." *Artificial Intelligence Review* 59.1 (2026): 17.
- [20] Shi, Tang, and Kazuya Shide. "A comparative analysis of LSTM, GRU, and Transformer models for construction cost prediction with multidimensional feature integration." *Journal of Asian Architecture and Building Engineering* 25.1 (2026): 634-649.
- [21] Fathima¹, Hafsa, et al. "Future Proofing Real Estate: Machine Learning for Price Predictions." (2026).
- [22] Macleod, Brandi M., et al. "Integration of machine learning and viscoelastic testing to improve survival prediction in horses experiencing acute abdominal pain at a veterinary teaching hospital." *Equine veterinary journal* 58.1 (2026): 105-114.
- [23] Sarun, Horn, In Nara, and Him Hun. "The Role of Artificial Intelligence in Transforming Higher Education Quality and Effectiveness: A Nano Review." *Journal of Agriculture and Environment* 3.2 (2026): 32-34.
- [24] Zhou, Jessica, et al. "Uncertainty-aware genomic deep learning with knowledge distillation." *npj Artificial Intelligence* 2.1 (2026): 3.
- [25] Alalawi, Khalid, Rukshan Athauda, and Raymond Chiong. "An extended learning analytics framework integrating machine learning and pedagogical approaches for student performance prediction and intervention." *International Journal of Artificial Intelligence in Education* 35.3 (2025): 1239-1287.



- [26] Alalawi, Khalid, et al. "Evaluating the student performance prediction and action framework through a learning analytics intervention study." *Education and Information Technologies* 30.3 (2025): 2887-2916.
- [27] Balachandar, V., and K. Venkatesh. "A multi-dimensional student performance prediction model (MSPP): An advanced framework for accurate academic classification and analysis." *MethodsX* 14 (2025): 103148.
- [28] Sonnleitner, Benedikt, et al. "Evaluation of early student performance prediction given concept drift." *Computers and Education: Artificial Intelligence* 8 (2025): 100369.
- [29] Zhang, Xiaoyi, et al. "Optimizing multi label student performance prediction with GNN-TINet: A contextual multidimensional deep learning framework." *PloS one* 20.1 (2025): e0314823.