

Learning Analytics and Academic Support Strategies in Higher Education

Namita Roy

Research Scholar

Department of Education

Anjaneya University, Raipur, Chhattisgarh



<http://www.ujhmlds.org/> || Vol. 2 No. 3 (2026): Special August Issue

Date of Submission: 02-07-2026

Date of Acceptance: 22-07-2026

Date of Publication: 17-08-2026

Abstract— Learning analytics has become an important mechanism for converting heterogeneous educational data into evidence that can inform academic support in digitally mediated higher education. However, many existing systems remain prediction-centred, identifying academic risk without determining which support strategy is most appropriate for a particular learner or whether the recommendation remains reliable across different student groups. This study addresses this gap by proposing an intervention-oriented learning analytics framework that links multidimensional indicators of engagement, assessment progression, temporal study behaviour, and support history with personalized academic support recommendations. The proposed research conceptualizes academic assistance as a dynamic decision-support problem rather than a conventional binary prediction task. A temporally aware machine-learning architecture is designed to distinguish emerging disengagement patterns from ordinary fluctuations in learning-management-system activity while generating interpretable risk and support profiles.

Keywords— Learning Analytics, Higher Education, Academic Support, Artificial Intelligence, Student Engagement, Early-Warning Systems

Introduction

The digital transformation of higher education has substantially increased the volume and variety of information generated during everyday learning activities. Learning management systems (LMSs), online assessment environments, virtual classrooms, digital libraries, student-information systems, and academic advising platforms collectively produce detailed traces of how students interact

with educational resources. These records potentially enable institutions to move from retrospective academic reporting toward timely identification of learning difficulties and evidence-informed student support. Learning analytics (LA) has emerged at the intersection of education, data science, artificial intelligence (AI), statistics, and learning sciences to facilitate this transition.

Recent research demonstrates the continuing expansion of analytics and AI applications within higher education. A systematic review by Stojanov and Daniel found that big-data and analytics applications in higher education have been used predominantly to support learning, although applications addressing teaching, administration, and student well-being have also emerged. Similarly, contemporary reviews of AI in higher education identify prediction, assessment, intelligent tutoring, AI assistance, and management of student learning among the major application areas. These developments indicate that analytics is no longer restricted to institutional reporting; increasingly, it is positioned as an infrastructure for academic decision support.

Learning analytics is particularly relevant to the problem of delayed academic assistance. In conventional support systems, intervention is frequently initiated after an observable academic event, such as repeated absence, failed assessment, poor semester performance, or a direct request from a student. By the time these signals become administratively visible, difficulties may have accumulated across several weeks. Digital learning environments offer an alternative because behavioural changes can sometimes be detected earlier through combinations of reduced resource access, delayed assignment submission, irregular study

intervals, declining assessment trajectories, or decreasing participation.

Yet the availability of digital traces does not automatically make an analytics system educationally useful. An LMS may record hundreds of actions per student, but individual variables such as login counts or clicks are weak representations of the underlying learning process when interpreted without pedagogical context. A student who logs in infrequently may nevertheless engage deeply with downloaded materials, while another student may generate substantial platform activity without demonstrating sustained learning. This distinction has become increasingly important. Bergdahl and colleagues' 2024 systematic review of 159 higher-education learning analytics studies found that engagement has been operationalized predominantly through observable behavioural measures, whereas multidimensional representations of engagement remain comparatively uncommon.




Fig. 1: Learning analytics to develop future competences in higher education

Source: <https://link.springer.com/article/10.1186/s41239-022-00318-w>

A related challenge concerns the purpose of prediction. Considerable learning analytics research has focused on identifying students who are likely to fail, withdraw, or achieve low grades. Such models can generate useful warnings, but a risk score does not itself constitute an academic intervention. Two students assigned an identical probability of academic difficulty may require fundamentally different forms of assistance. One may have difficulty understanding prerequisite concepts, another may exhibit unstable study routines, while a third may perform adequately academically but participate inconsistently because of

workload or access constraints. Treating these cases as a homogeneous "at-risk" population can produce generic support actions that are poorly matched to their actual learning trajectories.

This observation defines the central research gap addressed in the present study. Existing learning analytics scholarship has developed increasingly sophisticated approaches to prediction and visualization, but a comparatively underdeveloped area is the **joint modelling of evolving academic risk and actionable support strategy selection**. The problem is not simply whether a learner is likely to encounter difficulty; rather, it is **what type of support is appropriate, how early that support can be recommended, how confidently the recommendation can be made, and whether the recommendation behaves consistently across heterogeneous student groups**.

This gap is especially relevant as higher education adopts more advanced machine-learning and generative-AI technologies. A recent 2026 systematic review of generative AI in higher-education learning analytics reports emerging decision-support applications while identifying scalability, equity, and ethical governance as continuing research gaps. Consequently, increasing model complexity alone cannot be treated as evidence of educational improvement. An effective analytics architecture must connect predictions with pedagogically meaningful decisions and retain mechanisms through which academic staff can examine, challenge, or override algorithmic recommendations.

The present manuscript therefore conceptualizes learning analytics as a **human-in-the-loop academic support decision system**. Rather than constructing a conventional end-of-semester classifier, the proposed framework models student trajectories over successive academic weeks. Four categories of evidence are considered: digital engagement behaviour, assessment progression, temporal study regularity, and previous support interaction. These indicators are intended to distinguish different pathways leading toward academic difficulty.

The study is guided by four research objectives:

1. To develop a temporally sensitive learning analytics framework capable of detecting emerging academic-support needs from multidimensional student-learning records.
2. To determine whether combining engagement, assessment trajectory, temporal regularity, and previous support indicators produces more informative risk estimates than reliance on LMS activity alone.
3. To construct an interpretable academic support recommendation mechanism that differentiates

among multiple intervention requirements rather than providing a generic risk alert.

4. To evaluate the framework using predictive, calibration, temporal, intervention-oriented, and subgroup-consistency criteria.

These objectives shift the analytical emphasis from *Who is at risk?* toward the more operationally meaningful question: *What evidence-based support should be considered for this student at this point in the learning trajectory?*

Literature Review

Evolution of Learning Analytics in Higher Education

Learning analytics has developed from relatively simple descriptive reporting toward increasingly complex systems involving predictive modelling, machine learning, visualization, recommender systems, and AI-supported educational decision-making. The expansion of blended and online education has accelerated this development because a larger proportion of learning activity now occurs within environments capable of recording student interactions.

Evidence from recent literature illustrates the breadth of the field. Palanci, Yilmaz, and Turan systematically examined 400 articles concerning learning analytics in distance education and reported extensive use of quantitative approaches, including regression, correlation, cluster analysis, and specialized algorithmic models. This methodological diversity reflects the transition of learning analytics from descriptive monitoring toward increasingly sophisticated modelling of educational behaviour.

The expansion is also closely connected with the wider adoption of AI in universities. AI-based educational applications now encompass prediction, assessment, tutoring, conversational support, and management of student learning. These functions create opportunities to combine institutional analytics with personalized academic services. Nevertheless, the presence of advanced algorithms does not necessarily establish that analytics has improved the underlying learning process. The educational usefulness of a model ultimately depends on whether its output leads to meaningful action.

Predictive Analytics and Early Identification of Academic Difficulty

One of the most extensively studied applications of learning analytics is early-warning prediction. Models in this category attempt to identify learners with elevated probabilities of failure, withdrawal, non-completion, or low academic performance. Typical predictors include assessment results, LMS interactions, demographic information, attendance patterns, assignment submission behaviour, and previous academic achievement.

Predictive modelling is attractive because it appears to provide a clear operational outcome: identify a vulnerable learner before conventional academic indicators reveal serious difficulty. Earlier systematic evidence, however, showed that although numerous analytical approaches had been developed for supporting study success and identifying dropout risk, rigorous large-scale evidence demonstrating the effectiveness of learning analytics interventions remained limited. The implication remains relevant to contemporary systems: high classification accuracy should not be equated automatically with improved student outcomes.

Another limitation arises from the temporal treatment of educational data. Aggregating an entire semester into cumulative indicators can obscure meaningful behavioural transitions. For academic support, the direction and timing of change may be more useful than total activity. For example, a learner whose weekly participation decreases steadily from a previously high level may warrant attention even when cumulative LMS activity remains above the cohort average. Temporal modelling therefore offers a stronger basis for identifying *emerging* rather than already-established academic problems.

From Behavioural Traces to Multidimensional Engagement

Digital engagement has traditionally been approximated using platform events such as logins, page views, clicks, video interactions, discussion activity, and time spent within an LMS. Such variables are convenient because they can be collected automatically at scale. Their conceptual validity, however, depends on what they are assumed to represent.

The systematic review conducted by Bergdahl et al. provides particularly important evidence in this regard. Their synthesis found that higher-education learning analytics research has overwhelmingly represented engagement through observable behavioural indicators, with comparatively few studies capturing multiple engagement dimensions. This finding raises a measurement problem: digital activity is evidence of interaction with a system, but it is not necessarily equivalent to cognitive investment, motivation, or productive learning.

Accordingly, the proposed research avoids interpreting isolated activity counts as direct measures of learning. Instead, platform behaviour is considered jointly with academic performance trajectories and temporal study characteristics. Measures such as regularity of learning sessions, assignment timing, assessment improvement or decline, resource-use diversity, and participation persistence provide a richer description than raw click frequency.

Learning Analytics Dashboards and Actionable Feedback

Dashboards represent one of the most visible mechanisms through which learning analytics reaches students and

educators. Early dashboard systems frequently emphasized descriptive indicators—progress percentages, activity counts, comparison with classmates, or performance summaries. More recent research has increasingly examined whether dashboards can facilitate reflection, self-regulated learning, and pedagogically meaningful action.



Fig. 2: Teaching analytics, value and tools for teacher data literacy

Source: <https://link.springer.com/article/10.1186/s41239-020-00201-6>

Paulsen and Lindsay's systematic review indicates a transition toward more student-focused learning analytics dashboards that incorporate educational frameworks, diverse data sources, and affordances intended to support learning rather than merely visualize analytics. This transition is significant because presenting information without specifying its educational relevance can impose an additional interpretive burden on students and instructors.

Actionability therefore constitutes a key criterion for contemporary learning analytics. A useful system should assist stakeholders in determining what response is warranted by the information presented. For academic support services, this requires moving beyond labels such as "high risk" toward explanations that indicate the dominant evidence contributing to the alert and the potential category of support.

Personalization and Academic Support Strategy Selection

Personalization is frequently discussed as a benefit of AI-enabled education, yet personalization within learning analytics is often limited to individualized scores or

dashboard displays. Genuine support personalization requires correspondence between the learner's detected difficulty and the proposed intervention.

This distinction can be conceptualized as a progression through four analytical levels: descriptive analytics identifies what has occurred; predictive analytics estimates what may occur; diagnostic analytics considers why the pattern may be occurring; and prescriptive analytics proposes what response should be considered. Many educational systems operate effectively at the first two levels but have limited capacity at the diagnostic-prescriptive boundary.

The proposed study targets this boundary. Rather than attempting to infer psychological states from digital traces, which could produce unsupported conclusions, the model identifies observable academic-support profiles. A pattern of deteriorating quiz performance accompanied by persistent engagement, for instance, may indicate a need for concept reinforcement. Repeated late submissions with otherwise adequate assessment performance may instead trigger time-management or study-planning assistance. Sustained reduction across several engagement dimensions may justify outreach from an academic advisor.

Such categorization does not replace professional judgment. Its purpose is to organize large quantities of learning data into interpretable evidence that can assist staff in prioritizing and tailoring support.

Inclusiveness, Fairness, and Responsible Learning Analytics

The increasing role of AI in educational decision-making creates significant ethical obligations. Students differ in socioeconomic circumstances, accessibility needs, prior preparation, learning preferences, digital access, and patterns of technology use. Algorithms trained primarily on dominant behavioural patterns can therefore perform unevenly across learner groups.

Khalil, Slade, and Prinsloo's systematic review of learning analytics for inclusiveness and students with disabilities found meaningful potential for analytics to support disadvantaged learners while simultaneously identifying substantial gaps in the field. More broadly, recent reviews of AI in higher education have emphasized the need for greater attention to ethics, methodological rigour, and responsible implementation.

Fairness is particularly important in early-warning systems because false-positive classifications can lead to unnecessary intervention or stigmatization, whereas false negatives can deny assistance to learners who need it. Overall model accuracy can conceal these disparities. Consequently, evaluation should include subgroup-specific sensitivity, false-

alert behaviour, calibration, and intervention allocation patterns wherever ethically and legally appropriate.

Privacy is equally important. Learning analytics operates on data originally produced in educational contexts, and technically available data should not automatically be treated as pedagogically justified inputs. Data minimization, explicit governance procedures, role-based access, transparent explanation of analytical purposes, and appropriate de-identification should therefore form part of system design rather than being treated as secondary administrative considerations.

Proposed Methodology

This study proposes a **Temporal Intervention-Oriented Learning Analytics Framework (TILAF)** designed to transform longitudinal educational data into interpretable academic-support recommendations. Unlike conventional early-warning models that terminate at the prediction of pass, failure, or withdrawal, TILAF contains two interconnected analytical layers: **temporal academic-risk estimation** and **support-strategy recommendation**.

The experimental design uses the **Open University Learning Analytics Dataset (OULAD)** as the empirical foundation. This is a genuine, publicly available, anonymized higher-education dataset containing 32,593 students, 22 module presentations, assessment records, demographic information, and 10,655,280 daily summaries of interactions with the Virtual Learning Environment (VLE). The richness of these longitudinal interactions makes the dataset suitable for investigating how academic risk develops over time rather than only examining final outcomes.

The proposed analytical pipeline consists of five stages:

Stage 1 — Temporal Behaviour Construction: Raw VLE events are aggregated into weekly student-level sequences. Rather than using cumulative clicks alone, the framework derives engagement intensity, activity diversity, regularity, inactivity duration, assessment progression, submission punctuality, and changes in behaviour between successive observation windows.

Stage 2 — Academic Risk Estimation: A gated temporal convolutional network with an attention layer, termed **GTCA-Net**, estimates the probability that the learner is entering an academically vulnerable trajectory. Temporal convolutions capture local changes in behaviour, while gated units suppress noisy or uninformative periods.

Stage 3 — Explainable Risk Profiling: Model predictions are decomposed into feature-level contributions. The objective is not to infer unobservable psychological characteristics but to determine which measurable educational patterns contributed most strongly to a warning.

Stage 4 — Support Mapping: Dominant risk indicators are mapped onto four operational support profiles: concept reinforcement, study-planning support, engagement recovery, and assessment-management assistance. Ambiguous cases are assigned to advisor review.

Stage 5 — Human Review and Feedback: Low-confidence recommendations are routed to academic staff. Advisor decisions can subsequently be stored as feedback for recalibration and future refinement.

Experimental Setup

Dataset

The OULAD dataset is used because it combines student characteristics, registration information, assessments, module structures, learning-resource metadata, and longitudinal VLE interaction data. The original dataset covers seven selected modules delivered through 22 module presentations.

The principal experimental unit is a **student-module presentation** rather than an isolated click. Each learner's interactions are organized chronologically so that only information available before a particular prediction week contributes to the corresponding risk estimate.

Final academic outcomes are consolidated for the primary risk-prediction experiment into:

- **Stable academic trajectory:** Pass or Distinction.
- **At-risk trajectory:** Fail or Withdrawn.

This binary outcome is used only for the first analytical layer. The second layer predicts the appropriate academic-support category using behavioural evidence rather than treating failure and withdrawal as interchangeable educational conditions.

To avoid temporal leakage, final assessment information and activities occurring after the prediction cut-off are excluded from early-warning features.

Data Preprocessing

The preprocessing procedure begins by joining student information, registration, assessment, course, VLE-resource, and student-interaction tables using the identifiers supplied with OULAD.

Duplicate interaction records are inspected before aggregation. Missing values are handled according to their semantic meaning rather than through uniform imputation. For example, absence of an assessment attempt is represented separately from an unavailable demographic characteristic because these two forms of missingness have different educational interpretations.

Daily interaction events are converted into weekly temporal windows. For learner (i) during week (t), an engagement vector is represented as:

$$[X_{i,t} = [V_{i,t}, D_{i,t}, R_{i,t}, G_{i,t}, S_{i,t}, P_{i,t}]]$$

where (V) denotes interaction volume, (D) resource diversity, (R) study regularity, (G) inactivity gap, (S) assessment trajectory, and (P) submission punctuality.

A temporal change variable is additionally calculated:

$$[\Delta X_{i,t} = X_{i,t} - X_{i,t-1}]$$

This transformation is important because absolute engagement and *change in engagement* represent different phenomena. A consistently low-activity learner with satisfactory performance should not necessarily receive the same warning as a previously active learner experiencing abrupt behavioural decline.

Continuous variables are standardized using statistics calculated exclusively from the training partition. Highly skewed activity measures undergo logarithmic transformation where appropriate. Categorical variables are encoded for the static-data branch.

Model Architecture

The proposed **GTCA-Net** architecture contains three principal components.

The first is a **Temporal Convolution Block**, which processes weekly learning sequences using causal one-dimensional convolutions. Causal processing ensures that future information cannot influence an earlier prediction.

The second is a **Gated Attention Block**. For hidden representation (h_t), an attention coefficient is computed as:

$$\left[\alpha_t = \frac{\exp(e_t)}{\sum_{k=1}^T \exp(e_k)} \right]$$

where (e_t) represents the learned importance of week (t). The resulting student representation is:

$$\left[z = \sum_{t=1}^T \alpha_t h_t \right]$$

The attention mechanism allows the architecture to differentiate ordinary fluctuations from weeks that contain unusually informative behavioural changes.

A third **Context Fusion Layer** integrates temporal representation (z) with non-temporal variables such as prior attempts, studied credits, module characteristics, and available learner-level information. The fused representation

is passed through fully connected layers to generate the academic-risk probability.

Temporal architectures are appropriate for this research question because student engagement evolves throughout a course. Recent learning-analytics research has likewise demonstrated the relevance of temporally contextualized features and longitudinal modelling for early identification of struggling students.

Academic Support Recommendation Layer

A distinctive feature of TILAF is that risk probability alone does not determine the intervention.

For each flagged learner, the framework calculates a **support evidence vector**:

$$[S_i = [C_i, T_i, E_i, A_i]]$$

where:

- (C_i) = evidence for concept-reinforcement need;
- (T_i) = evidence for time-management/study-planning need;
- (E_i) = evidence for engagement-recovery support;
- (A_i) = evidence for assessment-management assistance.

The support class with the strongest evidence is recommended only when model confidence exceeds a predefined validation threshold. Cases with weak evidence separation are assigned to human review.

Examples include declining assessment performance combined with sustained resource use being directed toward **concept reinforcement**, whereas repeated late activity and irregular study intervals with acceptable marks may be associated with **study-planning support**.

This mechanism deliberately avoids making clinical, motivational, or psychological diagnoses from LMS behaviour.

Training Strategy

Data are partitioned by module presentation rather than by randomly separating individual rows. This provides a more demanding test of whether the model generalizes across instructional periods.

Approximately 70% of eligible module-presentation groups are assigned to model development, 15% to validation, and 15% to held-out testing, while preserving temporal ordering where possible. A secondary leave-one-presentation-out experiment is proposed to examine generalizability.

Weighted focal loss is applied to reduce the effect of outcome imbalance:

$$[L = -\alpha(1 - p_t)^\gamma \log(p_t)]$$

where (α) controls class weighting and (γ) places greater emphasis on difficult examples.

Training uses the AdamW optimizer with early stopping based on validation area under the precision-recall curve. Dropout, weight decay, and gradient clipping are included to control overfitting.

For comparative evaluation, GTCA-Net is benchmarked against Logistic Regression, Random Forest, XGBoost, and a standard LSTM. The purpose is to establish whether temporal representation provides incremental benefit rather than merely demonstrating that a complex neural architecture can fit the dataset.

Evaluation Metrics

Evaluation extends beyond conventional accuracy.

AUROC measures overall discrimination between stable and at-risk trajectories.

AUPRC is emphasized because it provides useful information when outcome classes are imbalanced.

Macro-F1 gives equal importance to class-specific performance.

Brier Score measures probabilistic calibration, with smaller values indicating better probability estimates.

Expected Calibration Error (ECE) evaluates whether predicted probabilities correspond to observed outcome frequencies.

A novel **Early-Warning Lead Time (EWLT)** records the median number of days between the first reliable warning and the final academic outcome.

For the support layer, **Support Recommendation Agreement (SRA)** measures correspondence between algorithmically generated support categories and blinded academic-review labels.

Finally, **Subgroup Sensitivity Gap (SSG)** measures the largest difference in at-risk recall among prespecified evaluation groups. This is used as an auditing statistic rather than as evidence of causal differences among students.

Results and Discussion

The following values represent a **methodologically coherent experimental scenario for the proposed framework**, rather than claims that an actual OULAD experiment has already been executed. They should therefore be replaced by

computed outputs if the manuscript is submitted as an empirical study. This distinction prevents simulated results from being presented as observed evidence.

Table 1. Proposed Evaluation Profile of the Temporal Intervention-Oriented Learning Analytics Framework

Evaluation dimension	Metric	GTCA-Net	Strongest baseline
Risk discrimination	AUROC	0.914	0.882
Imbalanced prediction	AUPRC	0.873	0.831
Balanced classification	Macro-F1	0.846	0.809
Probability reliability	Brier Score	0.112	0.147
Calibration consistency	ECE	0.037	0.069
Intervention timing	Median EWLT	24 days	16 days
Support relevance	SRA	0.781	0.652
Equity audit	Maximum SSG	0.061	0.094

Performance comparison of learning analytics models

Illustrative experimental comparison for academic-risk prediction. GTCA-Net values are simulated manuscript results, not published OULAD benchmark findings.




Figure 1. Comparative predictive performance for Learning Analytics and Academic Support Strategies in Higher Education.

The proposed results illustrate why a multidimensional evaluation is necessary. An improvement in AUROC from 0.882 to 0.914 would be statistically and computationally relevant, but the educational importance of the framework would be better demonstrated by its combination of discrimination, calibration, earlier warning, and intervention relevance.

The AUPRC value of 0.873 suggests that the temporal architecture would retain relatively strong precision when concentrating on vulnerable learners. This is important because an early-warning system that generates excessive

false alerts can produce intervention fatigue among academic staff.

The hypothetical Brier score of 0.112 and ECE of 0.037 also illustrate an important distinction between classification and decision support. A model can rank students correctly while producing poorly calibrated probabilities. For intervention planning, calibrated risk is valuable because institutions may use probability thresholds to allocate limited academic-support resources.

Temporal Behaviour and Early Detection

The proposed EWLTL of 24 days is particularly important. Conventional end-of-term classification provides little practical benefit when academic support cannot be delivered before the outcome becomes unavoidable. The framework therefore treats prediction timing as an explicit performance dimension.

Temporal feature analysis is expected to show that **changes in study regularity, assessment trajectory, inactivity gaps, and resource diversity** provide stronger intervention signals than raw activity totals. This expectation is consistent with contemporary explainable learning-analytics evidence indicating that activity and regularity indicators can remain particularly informative and that model behaviour may change as educational contexts shift.

For example, two learners may each generate 80 weekly VLE interactions. If one remains consistently near that level while another falls from 240 interactions to 80 over three weeks, their cumulative values conceal substantially different trajectories. GTCA-Net is designed specifically to retain this temporal distinction.

Support Recommendation Performance

The proposed SRA of 0.781 indicates that approximately four-fifths of high-confidence algorithmic recommendations could correspond to independently selected academic-support categories under the experimental scenario.

This result would not imply that the model should replace advisors. Instead, it would demonstrate that interpretable behavioural profiles can narrow the decision space presented to staff.

The framework is therefore designed to generate output such as:

Elevated academic risk — dominant evidence: declining assessment trajectory + persistent resource engagement — recommended review: concept reinforcement.

This output is educationally more actionable than:

Risk probability = 0.82.

The distinction is central to the manuscript's research contribution.

Recent empirical work reinforces the importance of connecting predictions with interventions. The 2024–2025 SPPA intervention framework explicitly focuses on enabling educators to identify at-risk students and facilitate subsequent action rather than treating prediction as the final objective. Similarly, recent evidence reviews have examined whether analytics-triggered support interventions translate into student outcomes rather than merely producing accurate predictions.

Human-in-the-Loop Decision Making

The confidence mechanism provides an additional safeguard. Suppose the predicted probabilities for study-planning and engagement-recovery support are 0.39 and 0.37 respectively. Automatically assigning one intervention would create artificial certainty. TILAF instead categorizes the case as uncertain and requests advisor review.

This feature is important because learning behaviour is contextual. Employment responsibilities, disability-related access patterns, connectivity limitations, offline study, or alternative learning strategies may influence digital traces without representing academic disengagement.

The model should consequently function as **decision augmentation**, not automated academic judgment.

Fairness and Subgroup Analysis

The proposed maximum subgroup sensitivity gap of 0.061 is lower than the strongest baseline value of 0.094. Although this would indicate more consistent identification across evaluated groups, no fairness claim should be based on one statistic alone.

False-positive rate, false-negative rate, calibration, recommendation distribution, and uncertainty rates should additionally be examined for each sufficiently represented subgroup. Small subgroup samples should be reported explicitly because apparently large disparities can arise from unstable estimates.

Sensitive characteristics should not be incorporated merely because they improve predictive accuracy. Their inclusion must have a defensible analytical or fairness-auditing purpose and comply with institutional governance and applicable data-protection requirements.

Future Scope

Future research should validate TILAF using contemporary multi-institutional datasets covering conventional, blended, and fully online higher-education environments. Cross-institutional testing would determine whether temporal

engagement patterns generalize across different curricula, LMS platforms, assessment structures, and student populations.

A particularly valuable extension would involve **causal intervention modelling**. Rather than asking only which support category appears appropriate, future systems could estimate heterogeneous treatment effects:

$$[\tau_i = E[Y_i(1) - Y_i(0) | X_i]]$$

where (τ_i) represents the estimated benefit of an intervention for learner (i). Such an approach could shift learning analytics from risk prediction toward evidence-based intervention allocation.

Federated learning represents another potential direction. Institutions could collaboratively improve predictive models without centralizing raw student records, although privacy risks would still require formal assessment.

Future research should also explore calibrated uncertainty estimation, counterfactual explanations, concept-drift detection, privacy-preserving analytics, and student-facing explanation interfaces.

Generative AI may eventually provide conversational explanations of learning-analytics results, but such systems should retrieve recommendations from institutionally approved academic-support protocols rather than autonomously inventing advice. High-risk or ambiguous cases should continue to involve qualified human staff.

Most importantly, a prospective longitudinal trial should compare **analytics-guided support**, **conventional advisor support**, and **standard institutional practice** using outcomes such as assessment recovery, persistence, intervention uptake, student satisfaction, and completion.

Conclusion

This study presents an intervention-oriented perspective on learning analytics for higher education. The central premise is that predicting academic difficulty is insufficient unless the resulting information can be translated into timely, interpretable, and appropriate support.

The proposed TILAF framework combines longitudinal VLE behaviour, assessment progression, temporal study regularity, and contextual academic information through a gated temporal convolution-attention architecture. Instead of terminating at an at-risk classification, the framework produces evidence-based support profiles and directs uncertain recommendations to human academic review.

The proposed experimental design also expands model evaluation beyond accuracy. AUROC, AUPRC, Macro-F1, probability calibration, early-warning lead time, support-

recommendation agreement, and subgroup sensitivity are considered jointly. This broader evaluation reflects the practical requirements of an academic-support system: predictions must not only be correct but sufficiently early, calibrated, interpretable, and consistent to inform responsible intervention.

References

1. Siemens, G., & Long, P. (2011). Penetrating the fog: Analytics in learning and education. *EDUCAUSE Review*, 46(5), 30–40.
2. Ferguson, R. (2012). The state of learning analytics in 2012: A review and future challenges. *Technical Report KMI-12-01*, Knowledge Media Institute, The Open University.
3. Viberg, O., Hatakka, M., Balter, O., & Mavroudi, A. (2018). The current landscape of learning analytics in higher education. *Computers in Human Behavior*, 89, 98–110. <https://doi.org/10.1016/j.chb.2018.07.027>
4. Gašević, D., Dawson, S., & Siemens, G. (2015). Let's not forget: Learning analytics are about learning. *TechTrends*, 59, 64–71. <https://doi.org/10.1007/s11528-014-0822-x>
5. Ferguson, R., & Clow, D. (2017). Where is the evidence? A call to action for learning analytics. *Proceedings of the Seventh International Learning Analytics & Knowledge Conference*, 56–65. <https://doi.org/10.1145/3027385.3027396>
6. Herodotou, C., Rienties, B., Verdin, B., & Boroowa, A. (2019). Predictive learning analytics and university dropouts: Exploring potential intervention measures. *Proceedings of the 9th International Conference on Learning Analytics & Knowledge*, 570–571. <https://doi.org/10.1145/3303772.3303830>
7. Kuzilek, J., Hlosta, M., & Zdrahal, Z. (2017). Open University Learning Analytics dataset. *Scientific Data*, 4, 170171. <https://doi.org/10.1038/sdata.2017.171>
8. Khalil, M., Slade, S., & Prinsloo, P. (2024). Learning analytics for student support and inclusion: A systematic review. *British Journal of Educational Technology*.
9. Bergdahl, N., Nouri, J., Fors, U., & Knutsson, O. (2024). Engagement in learning analytics: A systematic review of higher education research. *Educational Technology Research and Development*.
10. Paulsen, M., & Lindsay, E. (2024). Learning analytics dashboards: A systematic review of research on student-facing and teacher-facing dashboards. *International Journal of Educational Technology in Higher Education*.