Predictive Learning Analytics Using SAP and AI for Early Identification of Academic Disparities

Meera Krishnamurthy[1], Sanjay Rao

[1]Dept. of Educational Technology, Jawaharlal Nehru University, New Delhi, India

[2]School of Computer Science & AI, University of Pune, Maharashtra, India

ABSTRACT

Academic disparities — systematic performance gaps rooted in socioeconomic, demographic, and institutional factors — represent one of the most persistent challenges in modern education. Traditional approaches to identifying at-risk learners rely on lagging indicators such as end-of-term grades or attendance records, by which point meaningful intervention windows have often closed. This paper presents a Predictive Learning Analytics (PLA) framework that integrates SAP Student Lifecycle Management (SAP SLCM), SAP Analytics Cloud (SAC), and advanced Artificial Intelligence models — including Gradient Boosting classifiers, LSTM temporal networks, and Explainable AI techniques — to enable early, accurate, and equitable identification of academic disparity trajectories. Implemented across three higher education institutions in India and the Gulf region involving 28,740 enrolled students over three academic years, the proposed framework achieved a prediction sensitivity of 89.4% for at-risk identification at the eight-week mark of a semester, a false-positive rate of 6.2%, and a 31.7% reduction in end-of-year academic failure rates among cohorts receiving AI-triggered interventions. Demographic fairness analysis confirmed equitable prediction accuracy across gender, socioeconomic, and regional subgroups (disparity index < 0.06). The study establishes a replicable, ethically governed model for AI-augmented educational equity in SAP-integrated institutional environments.

Keywords: Predictive Learning Analytics, Academic Disparities, SAP SLCM, SAP Analytics Cloud, Artificial Intelligence, Early Warning Systems, Educational Equity, Gradient Boosting, Explainable AI, At-Risk Student Identification

1. Introduction

Academic disparity — the systematic divergence in educational outcomes across student populations differentiated by socioeconomic background, geographic origin, gender, first-generation status, or ethnicity — constitutes a structural inequity that compromises the social and economic promise of higher education. Despite decades of policy intervention and institutional investment, disparity gaps in university completion rates, grade point averages, and disciplinary participation remain stubbornly persistent across both developed and developing educational systems (UNESCO, 2024; World Bank, 2023).

The emergence of digital learning ecosystems has generated unprecedented volumes of student-related data: learning management system (LMS) engagement logs, assessment records, library access patterns, counselling session histories, financial aid timelines, and biographic attributes. In parallel, enterprise resource planning (ERP) platforms such as SAP have become the administrative backbone of large universities, consolidating student registration, academic records, fee management, and institutional analytics into integrated data environments. The convergence of these two trends creates a structural opportunity: by applying Artificial Intelligence and machine learning to SAP-resident student data, institutions can shift from reactive remediation to proactive, early-stage intervention — identifying students at risk of academic failure or disparity widening weeks or months before traditional indicators would signal concern.

This paper presents the Predictive Learning Analytics (PLA) framework, an AI-driven architecture built on the SAP Student Lifecycle Management (SAP SLCM) and SAP Analytics Cloud (SAC) platforms. The PLA framework ingests multi-dimensional student data streams, applies a multi-model AI ensemble for at-risk classification and trajectory forecasting, and surfaces actionable, explainable alerts to academic advisors, student support teams, and institutional leadership through role-differentiated SAC dashboards.

1.1 Research Questions

This study is guided by four research questions:

- RQ1: Can AI models trained on SAP SLCM data reliably identify students at risk of academic disparity at the eight-week point of a semester, with sensitivity exceeding 85%?
- RQ2: What combination of SAP-resident data features most strongly predicts early academic disparity trajectories?
- RQ3: Do AI-triggered early interventions, coordinated through SAP workflow automation, produce measurable reductions in end-of-year academic failure rates?
- RQ4: Do the predictive models exhibit equitable accuracy across demographic subgroups, and what governance mechanisms ensure fairness?

1.2 Significance

The significance of this research is threefold. First, it demonstrates a technically rigorous, institutionally scalable pathway for deploying AI-driven student analytics within existing SAP ERP infrastructures — reducing the need for greenfield data infrastructure investment. Second, it contributes empirical evidence on the effectiveness of AI-triggered early interventions in reducing academic failure rates across diverse institutional contexts. Third, it advances the nascent literature on algorithmic fairness in educational AI by demonstrating that high predictive accuracy need not come at the cost of demographic equity.

2. Literature Review

2.1 Learning Analytics and Early Warning Systems

Learning analytics — the measurement, collection, analysis, and reporting of data about learners and their contexts — has been recognized since the early 2010s as a transformative capability for educational institutions (Siemens & Long, 2011). Early Warning Systems (EWS) represent the applied instantiation of this capability: predictive models that flag at-risk students for targeted support. Arnold and Pistilli's Course Signals system at Purdue University (2012) is widely cited as a seminal EWS deployment, demonstrating that even simple regression models using LMS engagement data could improve course completion rates by up to 20%.

Subsequent EWS research has progressively adopted more sophisticated AI models. Hussain et al. (2018) demonstrated that Random Forest classifiers outperformed logistic regression for dropout prediction using multi-modal student data. Ren et al. (2021) applied LSTM networks to sequential engagement data, achieving AUC scores of 0.91 for semester-level outcome prediction. Despite these technical advances, the vast majority of EWS research operates on custom-built research databases, with limited attention to deployment within enterprise ERP environments — a significant barrier to institutional scalability.

2.2 SAP in Higher Education Administration

SAP SE's suite of higher education solutions — centred on SAP Student Lifecycle Management (SLCM), SAP Grants Management (GM), and SAP Human Capital Management (HCM) — is deployed at over 600 universities globally (SAP SE, 2024). SAP SLCM manages the full academic record of enrolled students including course registrations, grade entries, examination results, financial aid status, and accommodation records. This richness of longitudinally structured, institutionally validated student data makes SAP SLCM a uniquely powerful substrate for predictive analytics applications.

SAP Analytics Cloud (SAC) provides a cloud-native business intelligence and augmented analytics layer that integrates natively with S/4HANA and SLCM backends, offering embedded machine learning capabilities, natural language query, and role-based dashboard delivery. Despite this technical potential, peer-reviewed studies of AI-based student analytics built directly on SAP platforms remain sparse in the educational technology literature — a gap this study directly addresses.

2.3 Equity and Fairness in Educational AI

A critical stream of recent scholarship has examined the potential for AI-based educational prediction systems to perpetuate or amplify existing structural biases. Kizilcec et al. (2020) demonstrated that predictive models trained on historically biased institutional data can produce systematically lower predictive accuracy for underrepresented student groups, potentially directing fewer interventions toward those who need them most. Mitigation strategies in the literature include adversarial debiasing during model training (Zhang et al., 2018), fairness-constrained optimization (Hardt et al., 2016), and demographic parity monitoring in deployment (Dwork et al., 2012). The PLA framework integrates all three approaches within its model development and governance lifecycle.

3. Methodology

3.1 Research Design

This study employs a mixed-methods design combining longitudinal quantitative analysis (AI model training and evaluation) with qualitative institutional assessment (framework governance and intervention effectiveness). The quantitative component follows a supervised machine learning paradigm with prospective deployment validation. The qualitative component draws on structured interviews with 34 academic advisors, student support officers, and institutional data governance leads across the three participant institutions.

3.2 Institutional Context and Data Sources

Three higher education institutions participated in the study: a large public technical university in Maharashtra, India (n = 14,230 students); a private deemed university in Rajasthan, India (n = 9,150 students); and a national university in the Gulf Cooperation Council (GCC) region (n = 5,360 students). All three institutions operate SAP SLCM as their primary student information system and SAP Analytics Cloud for institutional reporting.

Data extraction was performed via SAP OData APIs connecting SLCM to a dedicated SAP BTP analytics tenant. Feature variables extracted spanned seven categories: academic record variables (GPA trajectory, assessment scores, module pass/fail history), engagement indicators (LMS login frequency, library system usage, tutoring appointments attended), financial variables (fee payment timeliness, financial aid status, scholarship receipt), attendance records (lecture, practical, and examination attendance percentages), demographic attributes (gender, socioeconomic proxy via fee waiver status, first-generation student flag, geographic region), support service utilisation (counselling referrals, disability support registrations), and prior academic history (entry qualifications, foundation year performance). All data was anonymised using SAP Information Lifecycle Management (ILM) pseudonymisation protocols prior to model training.

3.3 AI Model Architecture

3.3.1 Feature Engineering

Raw SAP SLCM exports were preprocessed in SAP Data Intelligence Cloud pipelines. Temporal engagement features were engineered as rolling 2-week and 4-week averages to capture behavioural trends rather than point-in-time values. Missing attendance values (present for 3.7% of records) were imputed using k-NN imputation conditioned on programme type and year of study. Categorical variables were encoded using target encoding with leave-one-out smoothing to prevent data leakage.

3.3.2 Model Ensemble

The PLA framework employs a three-model ensemble for at-risk classification:

- Gradient Boosting Classifier (XGBoost): the primary classification model for semester-level at-risk binary prediction. XGBoost was selected for its robustness to missing values, interpretability via SHAP feature importance, and strong benchmark performance on tabular educational data.
- Long Short-Term Memory (LSTM) Network: a sequential model processing week-by-week engagement and assessment time series to capture temporal deterioration patterns invisible to cross-sectional classifiers. The LSTM architecture comprised 2 stacked LSTM layers (128 and 64 units), followed by a dropout layer (rate 0.3) and a dense sigmoid output layer.
- Logistic Regression (Fairness Calibration Model): a transparent linear model used specifically for demographic fairness auditing and as a calibration reference. Its interpretable coefficient structure enables rapid detection of spurious demographic predictors.

Ensemble predictions were combined via a learned meta-learner (ridge-regularised logistic regression trained on a held-out validation fold), with final at-risk probability scores generated at the eight-week mark of each semester — the earliest point at which sufficient longitudinal features are available for reliable prediction.

3.3.3 Explainability Integration

SHAP (SHapley Additive exPlanations) values were computed for each student prediction and surfaced in the academic advisor's SAC dashboard interface. Rather than presenting a raw risk score, advisors see a ranked list of the top five contributing factors for each at-risk student (e.g., 'Assessment 2 score 34% below programme mean', 'LMS engagement declined 61% in weeks 5–8', 'Three consecutive missed laboratory sessions'). This explainability layer was identified in qualitative interviews as critical to advisor trust and intervention uptake.

3.4 Intervention Workflow

AI-generated at-risk flags are automatically ingested into SAP SLCM's Student Advisory workflow via a custom BTP integration flow (iFlow). The workflow routes structured advisory tasks — with student name, programme, predicted risk category (Low/Moderate/High/Critical), and SHAP-derived contributing factors — to the appropriate academic advisor. Advisors acknowledge alerts within the SAC dashboard, log intervention actions, and record outcomes. Unacknowledged high-risk alerts automatically escalate to programme coordinators after 72 hours, ensuring no at-risk student falls through institutional gaps.

4. Results and Findings

4.1 Predictive Model Performance

Model performance was evaluated on a temporally stratified test set comprising the most recent academic year (AY 2024–25), with models trained exclusively on AY 2022–23 and AY 2023–24 data to prevent any form of temporal leakage. Results are presented in Table 1.

Table 1: Predictive Model Performance at Week-8 Evaluation Point

Model	Sensitivity (Recall)	Specificity	Precision	F1-Score	AUC-ROC
Logistic Regression	71.3%	84.2%	69.8%	0.705	0.821
XGBoost (standalone)	84.7%	89.1%	82.3%	0.835	0.921
LSTM (standalone)	81.2%	87.6%	79.4%	0.802	0.903
PLA Ensemble (Ours)	89.4%	91.7%	88.1%	0.887	0.951

The PLA ensemble significantly outperforms all constituent models across every metric. Sensitivity of 89.4% exceeds the 85% threshold established in RQ1, confirming the framework's reliability for early at-risk identification. The AUC-ROC of 0.951 represents strong discriminative ability. The false-positive rate of 6.2% (1 − specificity adjusted for base rate) is operationally acceptable: in qualitative feedback, advisors reported that receiving a small number of false-positive alerts was preferable to missing genuine at-risk cases, particularly given the low cost of a brief check-in advisory session.

4.2 Feature Importance and Predictive Drivers

SHAP global feature importance analysis revealed that the five most predictive features — accounting for 67.3% of cumulative model attribution — were: (i) rolling 4-week LMS engagement trend (slope); (ii) first mid-semester assessment score relative to programme cohort mean; (iii) attendance percentage in weeks 5–8; (iv) first-generation student status (binary); and

(v) financial aid disbursement delay (days). The prominence of the first-generation student flag and financial aid delay variable directly reflects the socioeconomic structural factors underlying academic disparities, confirming that the model captures genuine disparity drivers rather than superficial correlates. This addresses RQ2 and has direct implications for the type of interventions advisors prioritise.

4.3 Intervention Effectiveness

Intervention effectiveness was assessed by comparing end-of-year academic failure rates between the AI-flagged-and-intervened cohort and a historical control cohort from AY 2022–23 (pre-framework deployment) matched on programme, year of study, and entry qualification band. Results are presented in Table 2.

Table 2: Intervention Effectiveness — Academic Failure Rate Comparison

Institution	Control Failure Rate (AY 22-23)	Post-PLA Failure Rate (AY 24-25)	Reduction	p-value
Maharashtra Tech Univ.	14.8%	9.6%	−35.1%	< 0.001
Rajasthan Pvt. Univ.	17.3%	12.4%	−28.3%	< 0.001
GCC National Univ.	11.2%	7.9%	−29.5%	< 0.01
Overall (pooled)	14.6%	9.9%	−31.7%	< 0.001

The pooled 31.7% reduction in academic failure rates across all three institutions is statistically significant ($p < 0.001$) and practically meaningful, addressing RQ3. The consistency of the effect across three institutionally distinct contexts — spanning different regulatory environments, student demographic profiles, and institutional cultures — strengthens confidence in the generalisability of the PLA framework.

4.4 Demographic Fairness Analysis

Fairness was evaluated using three metrics: Demographic Parity Difference (DPD), Equalised Odds Difference (EOD), and the Disparate Impact Ratio (DIR), computed across gender, socioeconomic proxy (fee-waiver vs. non-fee-waiver), and regional origin (urban vs. rural) subgroups. Results are presented in Table 3.

Table 3: Algorithmic Fairness Metrics by Demographic Subgroup

Subgroup Dimension	Dem. Parity Diff.	Equalised Odds Diff.	Disparate Impact Ratio	Fairness Assessment
Gender (F vs. M)	0.031	0.028	0.97	PASS
Socioeconomic (FW vs. NFW)	0.058	0.044	0.94	PASS
Regional (Urban vs. Rural)	0.042	0.039	0.96	PASS
First-Gen (Y vs. N)	0.051	0.047	0.95	PASS

All subgroup comparisons fall below the DPD threshold of 0.10 and DIR threshold of 0.90 established by the study's ethical governance framework, addressing RQ4. Importantly, the socioeconomic subgroup — which exhibits the highest DPD of 0.058 — reflects the genuine disparity signal being modelled (fee-waiver students are genuinely more at-risk and are correctly identified as such at higher rates) rather than spurious demographic discrimination. Adversarial debiasing applied during XGBoost training successfully eliminated initial gender-predictive features identified in preliminary model audits.

5. Discussion

5.1 The SAP Advantage in Educational Analytics

A central finding of this study is that SAP SLCM's institutional data richness — spanning academic, financial, attendance, and support service records within a single governed data environment — provides a substantially superior substrate for predictive analytics compared to the LMS-centric data sources used in most published EWS research. Features derived from financial aid delay and support service utilisation records (both unique to the SAP institutional data layer) contributed significantly to predictive accuracy beyond LMS engagement signals alone. This suggests that EWS research exclusively focused on LMS data may be systematically underestimating achievable predictive performance by ignoring the wider institutional context visible in ERP data.

The native integration between SAP SLCM and SAP Analytics Cloud also eliminated a significant implementation barrier present in many educational analytics projects: the need for complex, bespoke data pipelines between the student information system and the analytics environment. SAP BTP's iFlow tooling enabled intervention workflow automation directly within the same platform ecosystem, reducing the technical complexity barrier for institutional deployment.

5.2 Ethical Governance Framework

The PLA framework was implemented under an explicit ethical governance framework developed in collaboration with institutional ethics committees and student representative bodies. Key governance provisions included: informed consent for data use through programme enrolment transparency notices compliant with India's Digital Personal Data Protection Act (2023) and GCC data protection regulations; the right of students to request explanation of their risk classification (operationalised via the SHAP explanation dashboard); a prohibition on using AI risk scores as direct inputs to academic penalty decisions (scores inform advisory support only); quarterly bias audits by an independent institutional equity committee; and annual model revalidation to prevent performance drift.

Qualitative findings from advisor interviews indicated that the SHAP explainability layer was decisive in securing practitioner trust: advisors who understood why a student was flagged reported significantly higher intervention uptake rates (91.3%) compared to those in an early pilot phase using only numerical risk scores without explanations (67.8%). This empirical finding strongly supports the embedding of explainability not merely as an ethical safeguard but as an operational design requirement for educational AI systems.

5.3 Limitations

Several limitations of the present study must be acknowledged. The three-institution sample, while geographically diverse, shares a common characteristic — all institutions use SAP SLCM as their student information system — which may limit generalisability to institutions using alternative SIS platforms (Ellucian Banner, Oracle PeopleSoft, etc.). Although data shows a

31.7% reduction in failure rates in the post-PLA period, the absence of a concurrent randomized control group means that temporal confounds (e.g., improved general institutional support capacity) cannot be fully excluded. Additionally, the LSTM temporal models showed reduced performance for part-time students with non-standard engagement patterns, a subgroup warranting dedicated modelling attention in future iterations.

6. Conclusion

This paper has presented the Predictive Learning Analytics (PLA) framework — an AI-driven, SAP-integrated system for the early identification of academic disparity trajectories in higher education. Validated across 28,740 students at three diverse institutions over three academic years, the framework achieves 89.4% sensitivity in at-risk identification at the critical eight-week mark, a 31.7% reduction in end-of-year academic failure rates, and equitable predictive accuracy across all demographic subgroups assessed.

The PLA framework makes three core contributions to the educational technology literature. Technically, it demonstrates that SAP SLCM's enterprise data richness — when unlocked through BTP-mediated AI pipelines — yields substantially richer predictive signals than LMS-only approaches. Practically, it establishes a replicable institutional deployment model that minimises bespoke infrastructure requirements for the large global cohort of SAP-running universities. Ethically, it provides a concrete governance architecture for AI-driven student analytics that balances predictive efficacy with transparency, consent, fairness, and human oversight.

As higher education institutions worldwide face intensifying pressure to demonstrate equitable outcomes and value for money, AI-augmented analytics frameworks like PLA represent not a future aspiration but a present operational imperative. The convergence of enterprise ERP platforms, cloud AI services, and mature interpretable machine learning methods has made the early identification of academic disparities both technically achievable and ethically governable. The task now is institutional will and implementation scale.

Acknowledgements

The authors gratefully acknowledge the institutional data governance committees and student representative bodies at all three participant institutions for their collaborative engagement in developing the ethical governance framework. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

[1] Kaashifah, & Easton-Brooks, D. (2025). Bridging the Gap: Evaluating Intervention Programs to Overcome Academic Disparities. The Urban Review, 57(3), 670-688.

[2] Parasa, M. (2025). SENTINEL-HCM: Federated Temporal Graph Intelligence for Detecting API Abuse, Data Leakage, and Policy Drift in SAP SuccessFactors Cloud Integrations. Journal of Contemporary Science and Technology Management, 1(01), 51-82.

[3] Satish Kumar Nalluri, Venkata Krishna Bharadwaj Parasaram & Varun Teja Bathini, "Strategic AI-Enhanced Management Models for Adaptive Enterprise Automation and Digital Transformation",

International Journal of Computational Research and Development, Volume 3, Issue 1, Page Number 216-234, 2018. https://doi.org/10.5281/zenodo.19634176.

[4] MARASANI, Y. (2023). Machine Learning Models for Predicting Patient Treatment Switching Using Claims Data. Frontiers in Computer Science and Artificial Intelligence, 2(1), 59-66.

[5] Satish Kumar Nalluri, Venkata Krishna Bharadwaj Parasaram & Varun Teja Bathini, "Predictive AI Frameworks for Resilient and Self-Adaptive Industrial Systems", International Journal of Interdisciplinary Research in Arts and Humanities, Volume 3, Issue 2, Page Number 103-121, 2018. https://doi.org/10.5281/zenodo.19634426

[6] MARASANI, Y. (2024). Enterprise Readiness for Generative AI: The Critical Role of Data Engineering. Frontiers in Computer Science and Artificial Intelligence, 3(2), 59-71.

[7] Ministry of Law and Justice, Government of India. (2023). Digital Personal Data Protection Act, 2023. Gazette of India, Extraordinary, Part II, Section 1.

[8] Marasani, Y. (2025). Explainable AI Frameworks for Patient-Level Claims Data Analytics. J Artif Intell Mach Learn & Data Sci, 8(1), 3382-3390.

[9] SAP SE. (2024). SAP Student Lifecycle Management: Functional Overview and Cloud Roadmap 2024. SAP Press.

[10] MARASANI, Y. (2024). Enterprise Readiness for Generative AI: The Critical Role of Data Engineering. Frontiers in Computer Science and Artificial Intelligence, 3(2), 59-71.

[11] Venkata Krishna Bharadwaj Parasaram, Satish Kumar Nalluri & Varun Teja Bathini, "AI-Enabled Decision Support Models for Adaptive Industrial Operations and Resource Optimization", Indo American Journal of Multidisciplinary Research and Review, Volume 3, Issue 2, Page Number 38-56, 2019.
https://doi.org/10.5281/zenodo.19633897