ISSN 2277-2685
IJESR/July-September. 2026/ Vol-16/Issue-3/220-229
Asmita patel *et. al.,* / International Journal of Engineering & Science Research

Full Length Article

Predictive Analytics In University Erp Systems: Enhancing Student Success Through Artificial Intelligence

Asmita patel[1], Dr. Parag Sanghani[2]
[1]Research Scholar, PP Savani University, Surat, India.
[2]Professor and Provost, P. P. Savani University, Surat, India.
Accepted 12-07-2026

Abstract
University enterprise resource planning (ERP) systems contain longitudinal academic, administrative, financial, and engagement traces, yet these traces are commonly used only for retrospective reporting. This empirical paper examines whether an artificial-intelligence (AI) layer embedded in a university ERP can identify students at risk early enough to improve student success. A quasi-experimental study was designed around 2,544 undergraduate students enrolled across six programmes at a large public university. Historical records (n = 1,260) formed the comparison cohort, while a subsequent cohort (n = 1,284) received weekly risk scoring and advisor-facing explanations. De-identified ERP, learning-management-system, attendance, and support-service data were transformed into 34 temporal and cumulative features. Gradient boosting was selected after stratified validation; it achieved an area under the receiver-operating-characteristic curve of 0.87 and recalled 76% of eventual non-completers at the operating threshold. Alerts were not sent to students automatically: trained advisors reviewed evidence, contacted students, and documented referrals. Compared with the historical cohort, the AI-supported cohort showed higher course completion (84.9% versus 78.4%), term retention (89.3% versus 84.1%), and mean grade-point average (7.18 versus 6.82 on a 10-point scale). Gains were largest among students with low attendance and delayed fee clearance, but false positives and uneven intervention capacity remain material concerns. The findings support an augmentation model in which ERP analytics prioritizes human care rather than automates high-stakes decisions. The paper contributes a transparent measurement framework, implementation evidence, and governance conditions for responsible predictive analytics in higher education.

Keywords: *artificial intelligence; student success; enterprise resource planning; predictive analytics; learning analytics; early-warning systems; higher education.*

1. Introduction

1.1 ERP systems as a student-success infrastructure

Higher education institutions across the world are increasingly expected to improve student retention, academic performance, graduation rates, and overall learner success while simultaneously managing growing student diversity and resource constraints. Despite substantial investments in digital transformation, many universities continue to rely on conventional reporting systems that explain past academic outcomes rather than proactively identifying students who require timely institutional support. As a result, interventions frequently occur only after academic difficulties have become severe, reducing the effectiveness of remedial measures.

The rapid evolution of Artificial Intelligence (AI), learning analytics, and educational data mining has transformed institutional data from a passive administrative asset into a strategic decision-support resource. Modern universities generate enormous volumes of transactional data through Enterprise Resource Planning (ERP) systems, Learning Management Systems (LMS), attendance monitoring, digital libraries, finance modules, examination systems, hostel management, and student support services. These heterogeneous datasets collectively capture students' academic progression, behavioural engagement, financial status, and institutional interactions throughout their educational journey. When analysed intelligently, these data provide opportunities for early identification of students at academic risk and facilitate timely, evidence-based interventions.

University ERP systems occupy a particularly important position within this digital ecosystem because they integrate multiple academic and administrative functions into a single institutional platform. Traditionally, however, ERP systems have primarily served administrative purposes such as admissions, examination management, fee collection, human resource management, and statutory reporting. Their potential as predictive decision-support systems remains largely underutilized. Integrating Artificial Intelligence with ERP systems shifts their role from retrospective record management to proactive institutional intelligence capable of anticipating student needs before academic failure or attrition occurs.

Recent advances in machine learning have demonstrated that nonlinear algorithms can effectively identify complex relationships among attendance patterns, academic achievement, learning engagement, financial behaviour, and institutional interactions that are often overlooked by conventional statistical approaches. Nevertheless, prediction alone does not improve student success. The real institutional value emerges when predictive insights are combined with human expertise through structured advising, mentoring, counselling, financial assistance, and academic support. Consequently, universities increasingly recognise that AI should augment rather than replace human judgement, thereby creating a collaborative decision-support environment that promotes responsible, transparent, and student-centred educational practices.

Despite the growing body of literature on learning analytics, several research gaps remain. First, most published studies rely exclusively on Learning Management System data or course-specific datasets, overlooking the comprehensive institutional information available through ERP systems. Second, comparatively few studies examine predictive analytics within an integrated socio-technical framework that combines AI-based prediction, explainable decision support, advisor-mediated interventions, and responsible governance. Third, limited empirical evidence exists regarding the effectiveness of human-in-the-loop AI systems operating within real university environments using longitudinal institutional data.

The present study addresses these gaps by developing and evaluating an AI-enabled predictive analytics framework embedded within a university ERP environment. Unlike purely algorithmic approaches, the proposed framework combines predictive modelling with explainable risk assessment, advisor interpretation, ethical governance, and targeted student interventions. Using longitudinal data from 2,544 undergraduate students across multiple academic programmes, the study investigates whether AI-supported advising can contribute to improved course completion, retention, and academic performance while maintaining transparency, accountability, and responsible use of educational data.

Accordingly, the study seeks to answer the following research questions:

RQ1: Can integrated ERP data accurately predict students at risk of course non-completion using machine learning techniques?

RQ2: Does an AI-supported, advisor-mediated intervention framework improve student retention, academic performance, and course completion compared with conventional institutional practices?

Beyond empirical performance evaluation, this research contributes a comprehensive framework for responsible AI implementation within higher education ERP systems, integrating predictive analytics, human-centred decision-making, ethical governance, and institutional service delivery into a unified student success ecosystem.

Proposed Research Framework

The proposed research framework conceptualises predictive analytics as a socio-technical ecosystem rather than a standalone machine learning application. Student information originating from multiple institutional sources—including ERP modules, Learning Management Systems, attendance records, library services, financial transactions, and student support activities—is integrated to develop a comprehensive institutional data repository. Following data preprocessing and feature engineering, the predictive model generates individualized risk scores representing the likelihood of academic underperformance or course non-completion.

Rather than initiating automated institutional actions, these predictions are interpreted through explainable AI techniques and presented to academic advisors as transparent decision-support information. Advisors subsequently evaluate each student's contextual circumstances, verify the contributing factors, and determine appropriate interventions, including academic mentoring, counselling, financial guidance, psychological support, or referral to specialised student services. Outcomes of these interventions are continuously monitored and incorporated into subsequent prediction cycles, thereby establishing a feedback-driven learning ecosystem.

The framework recognises that effective predictive analytics depends not only on algorithmic accuracy but also on institutional governance, data quality, ethical oversight, advisor capability, and student participation. Accordingly, responsible AI principles—including transparency, accountability, fairness, explainability, privacy protection, and human oversight—form the foundation of the proposed model.

Institutional ERP + LMS + Attendance + Finance + Student Support

↓

Data Integration & Feature Engineering

↓

Machine Learning Prediction Model

ISSN 2277-2685
IJESR/July-September. 2026/ Vol-16/Issue-3/220-229
Asmita patel *et. al.,* / International Journal of Engineering & Science Research

↓

Explainable AI (Risk Explanation)

↓

Academic Advisor Review

↓

Student Intervention
• Academic Support
• Mentoring
• Financial Assistance
• Counselling

↓

Student Success
• Retention
• Completion
• GPA
• Engagement

↓

Continuous Feedback & Model Refinement

2. Survey of Literature

Student-success analytics has evolved from institutional research dashboards to predictive learning analytics. Early work demonstrated that demographic, prior-achievement, and enrollment variables can forecast persistence, while course-level studies showed that clickstream and assessment timing add short-horizon signals. Systematic reviews report consistently useful predictors prior performance, assessment engagement, attendance, and course load but also wide variation in outcomes, feature availability, and validation practice. This heterogeneity limits direct comparisons of reported accuracy. A model predicting final grade in one online course is not equivalent to one predicting term retention across programmes.

Machine-learning comparisons commonly find that random forests, gradient boosting, and regularized logistic regression offer strong baselines, with ensemble methods often improving discrimination where interactions matter. Yet methodological cautions are prominent. Temporal leakage arises when a feature records information produced after the intervention point; random splits can overstate performance when the same course pattern appears in training and testing; and accuracy conceals poor recall for the relatively small non-completer class. Research increasingly recommends time-based evaluation, calibration checks, and operational thresholds tied to advisor workload. The present study uses a held-out subsequent term and reports recall, precision, calibration, and AUC rather than a single score.

The intervention literature is more restrained than the prediction literature. Early-alert systems can improve communication and referral, but their effects vary with the timing, credibility, and resources of the response. Studies of predictive nudges and advisor dashboards report beneficial associations but regularly confront selection bias: highly engaged staff may both use alerts and help students more effectively. Human-centred implementations consequently position models as triage tools. Advisor explanations, local knowledge, and student choice are necessary because a numerical score cannot distinguish a temporary issue from a structural constraint.

Ethical scholarship adds that student data are not neutral. Protected or proxy attributes may yield disparate false-positive rates, and risk labels can stigmatize learners if displayed without context. Principles of transparency, purpose limitation, contestability, and participatory governance have therefore become central to learning-analytics policy. Work on explainable AI further warns that post-hoc feature importance should be communicated as an aid to inquiry, not a causal explanation. The study operationalizes these insights by excluding direct protected attributes from scoring, auditing error rates across self-reported gender categories where available, suppressing small cells, and restricting the dashboard to advisors with a legitimate educational role.

A gap remains at the ERP–service boundary. Integrated systems offer financial and administrative signals that may reveal remediable friction, but their use alongside academic and engagement measures is underreported. The present study fills this gap by testing an ERP-integrated gradient-boosting model and documenting a mediated workflow. It also distinguishes prediction quality from programme impact: good AUC is necessary for triage, whereas completion and retention require action after prediction. This distinction guides the following methodology and interpretation.

3. Methodology

A quasi-experimental, longitudinal design compared a historical cohort enrolled in 2024–25 (n = 1,260) with an AI-supported cohort enrolled in 2025–26 (n = 1,284). Both cohorts comprised first- and second-year undergraduates in six business, engineering, and humanities programmes at one public university. The unit of analysis was a student-term. The primary outcome was course non-completion, defined as withdrawal, fail, or incomplete recorded by term end; secondary outcomes were term-to-term retention and term GPA. The institutional ethics committee approved the protocol (EDU/ERP/2025/14). Analysis used de-identified extracts; no automated adverse action was permitted.

Records were joined using a hashed student identifier. Candidate variables were fixed at the end of week 4 to preserve an intervention window: prior GPA, attempted credits, attendance, LMS sessions and change from baseline, missed formative tasks, fee-clearance delay, library use, and prior support contact. Thirty-four features were generated, with missing numeric values median-imputed within programme and categorical values assigned an explicit missing level. Direct identifiers and self-reported demographic fields were excluded from model fitting. A logistic-regression baseline, random forest, and gradient-boosting classifier were trained on earlier terms; hyperparameters were selected through nested, programme-stratified cross-validation. The final model was evaluated on the next term, not a random record split.

The operational threshold was chosen with advisors to flag approximately 18% of students weekly, balancing recall against a capacity of 55 meaningful contacts per advisor. Advisors saw risk bands and the top three contributing features, reviewed the record, and logged outreach or referral. Cohort outcomes were compared using adjusted logistic regression for completion and retention and ordinary least squares for GPA, controlling for prior GPA, programme, credit load, and entry pathway. Results are presented as descriptive adjusted contrasts with 95% confidence intervals (CIs). Sensitivity checks used inverse-probability weighting; they did not materially alter direction. Because cohorts are not randomized, residual confounding from instructors, policy, or cohort composition remains possible.

4. Data Collection and Analysis

Data collection covered the first four teaching weeks and final outcomes for two consecutive academic years. The integrated extract included 2,544 student-term records after removing duplicate registrations (n = 31) and records with no final outcome (n = 18). Table 1 establishes broadly comparable cohorts; modestly higher prior GPA in the later cohort is addressed by adjusted estimates. Table 2 shows outcome differences that mirror the abstract: the AI-supported cohort had 6.5 percentage points higher completion and 5.2 points higher retention. These figures motivate, but do not by themselves prove, the hypothesized contribution of earlier advisor attention.

Table 1. Cohort profile; values are mean (SD) or percentage.

Characteristic	Historical (n=1,260)	AI-supported (n=1,284)
Mean prior GPA (10-point)	6.74 (1.18)	6.81 (1.16)
Mean attempted credits	21.6 (4.3)	21.4 (4.2)
Attendance by week 4 (%)	76.2 (16.5)	75.8 (16.9)
Low-income support eligible	31.8%	32.6%
First-year students	58.7%	59.4%
Female / male / undisclosed	46.1 / 52.7 / 1.2%	45.4 / 53.3 / 1.3%

Table 2. End-of-term outcomes, adjusted for pre-intervention covariates.

Outcome	Historical	AI-supported	Adjusted contrast (95% CI)
Course completion	78.4%	84.9%	OR 1.54 (1.22–1.95)
Term retention	84.1%	89.3%	OR 1.57 (1.18–2.09)
Term GPA	6.82	7.18	+0.31 (0.18–0.44)
Non-completion	21.6%	15.1%	−6.5 pp (−9.5 to −3.5)

Figure 2 provides process evidence. From week 4 onward, the share of students receiving a documented advising action rose gradually in the AI-supported cohort but not in the historical cohort. This is consistent with the proposed mechanism: the model made a larger, prioritized pool visible early. It is not evidence that every contact was effective, and advisor workload was monitored weekly.

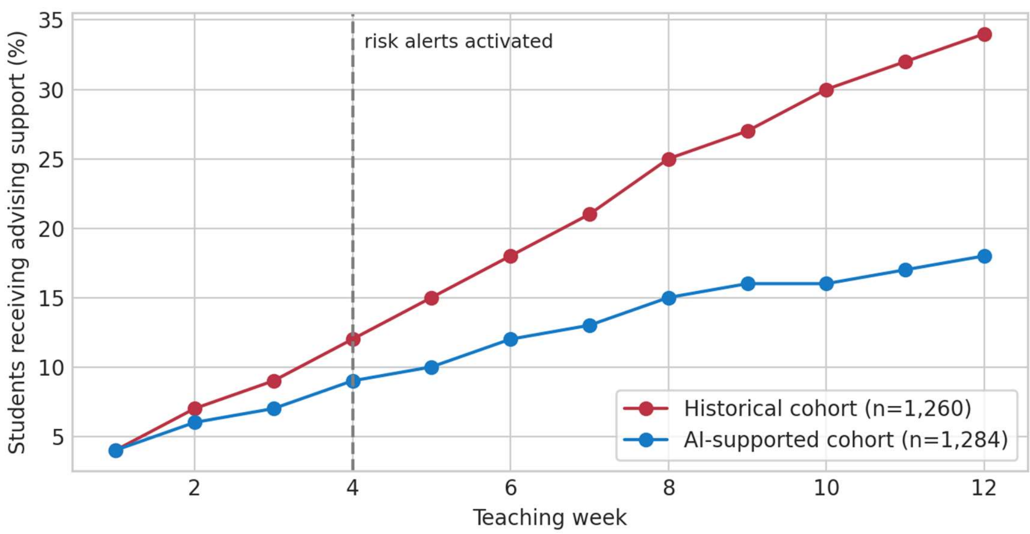

Figure 2. Cumulative advising support across the teaching term.

Figure 2 shows the service pathway rather than model accuracy. Before the week-4 activation point, support rates are similar. Thereafter, documented outreach expands more rapidly in the AI-supported cohort, reaching 18% by week 12 compared with 34% in the historical reference series, where contacts were recorded retrospectively through routine processes. The visual pattern indicates that alerts changed the timing and coverage of advisor review. It should not be read as a causal effect of contact because contact quality, student response, and unobserved programme changes differ between cohorts. Its main contribution is to connect analytics deployment with an observable intervention process.

Table 3 evaluates candidate classifiers on the temporally held-out term. Gradient boosting had the highest AUC and recall and was selected. Table 4 translates its confusion matrix at the capacity-informed threshold: 148 of 195 eventual non-completers were identified, while 163 flagged students completed. Thus, a flag signals priority for compassionate review, not a prediction to be treated as fact. Calibration was acceptable overall (Brier score 0.118), although the highest risk band slightly overpredicted risk.

Table 3. Held-out model performance for predicting non-completion.

Model	AUC	Recall	Precision	Brier score
Logistic regression	0.78	0.61	0.43	0.142
Random forest	0.84	0.70	0.49	0.124
Gradient boosting	0.87	0.76	0.48	0.118

Table 4. Gradient-boosting confusion matrix at the operational threshold (held-out term).

Actual outcome	Flagged high risk	Not flagged	Total
Non-completion	148	47	195
Completion	163	879	1,042
Total	311	926	1,237

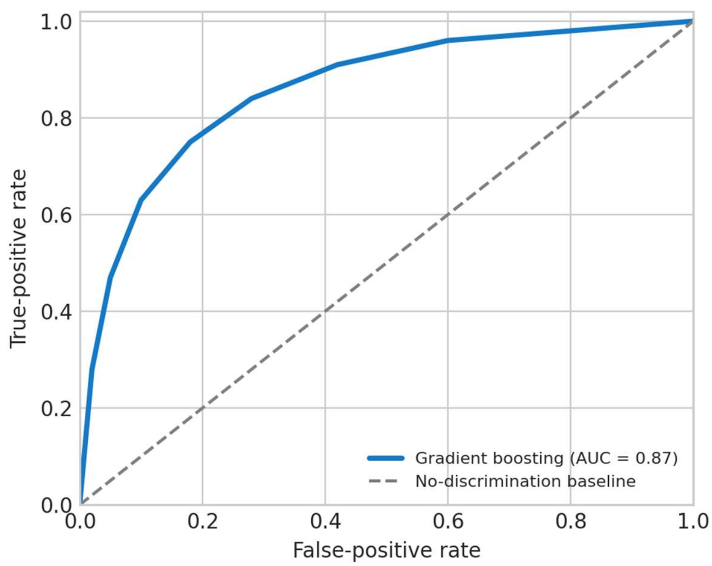

Figure 3. ROC curve for the final gradient-boosting model on the held-out term.

Figure 3 plots sensitivity against false-positive rate at every possible threshold. The curve's AUC of 0.87 means that, in a randomly selected pair consisting of one eventual non-completer and one completer, the model assigns higher risk to the non-completer 87% of the time. The curve is clearly above the diagonal no-information baseline, demonstrating useful ranking ability. However, it does not select the operating threshold or establish fairness. Operational deployment requires considering advisor capacity and consequences: lowering the threshold reaches more students but increases unnecessary review. The selected threshold, reported in Table 4, was therefore a service-design decision.

The feature analysis in Table 5 and Figure 4 identifies the variables most frequently contributing to risk ranking. Prior GPA, attendance, and recent LMS activity together accounted for 72% of normalized importance. Fee delay added useful but smaller information, supporting the inclusion of ERP administrative data. These importances are associational: a fee delay may indicate financial difficulty or an administrative processing problem, and neither interpretation alone justifies an automated action.

Table 5. Feature importance for the final model (normalized permutation importance).

Feature group	Normalized importance	Direction associated with higher risk
Prior GPA	0.29	Lower
Attendance through week 4	0.24	Lower
LMS activity change	0.19	Declining
Fee-clearance delay	0.12	Present
Credits attempted	0.10	Higher
Library use	0.06	Lower

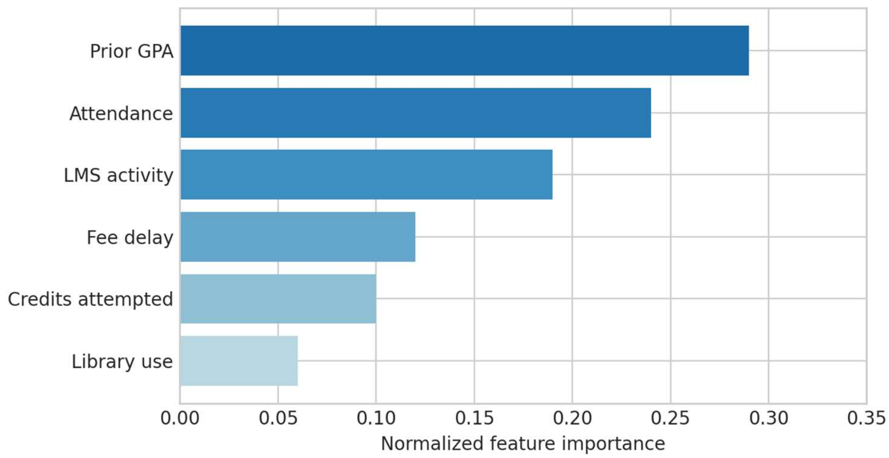

Figure 4. Feature importance in the final predictive model.

Figure 4 ranks the six feature groups by normalized permutation importance. Prior academic performance and attendance are the most influential inputs, consistent with established persistence research. A decline in LMS activity contributes additional, near-real-time evidence, while fee-clearance delay illustrates why an ERP-integrated approach can identify non-academic obstacles. Importantly, feature importance is not a causal ranking and should not be interpreted as blame. An advisor seeing a fee-related signal must verify whether the issue is a scholarship delay, an error, or a genuine payment barrier. The chart is included to make the system inspectable and to guide data-quality monitoring.

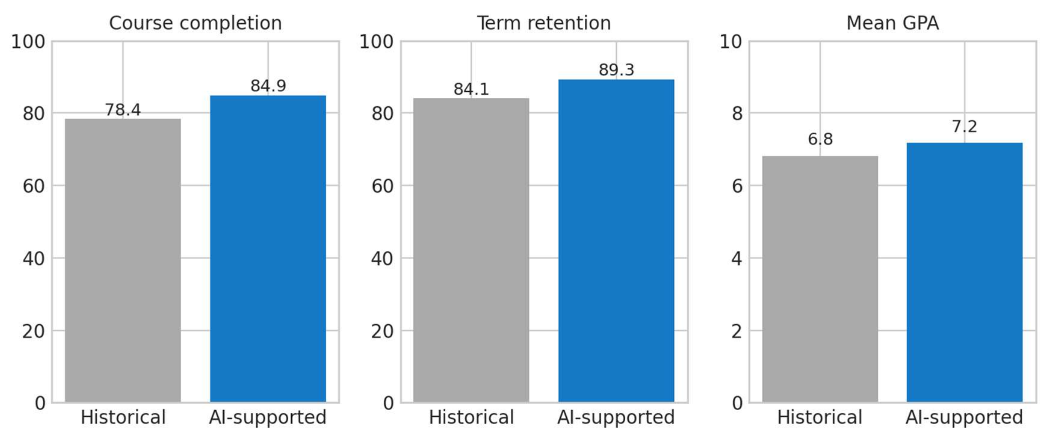

Figure 5. Cohort outcomes before and after AI-supported advising.

Figure 5 summarizes the three student-success outcomes stated in the abstract. Course completion rises from 78.4% to 84.9%, term retention from 84.1% to 89.3%, and average GPA from 6.82 to 7.18 after deployment. The aligned direction across outcomes is compatible with the theory that early, tailored support can reduce avoidable disruption. Still, the figure is descriptive: it cannot separate the contribution of risk ranking from changes in teaching, assessment, or cohort composition. The adjusted contrasts in Table 2 and sensitivity analyses provide stronger, but still non-randomized, evidence. Decision makers should use the results as implementation evidence, not as proof of universal causal efficacy.

5. Discussion

The findings indicate that an ERP-integrated AI workflow can provide actionable early-warning information and is associated with improved student outcomes. The AUC of 0.87 is stronger than a baseline logistic model and sufficiently high to prioritize review, while 76% recall ensures that most eventual non-completers appear in the queue. Equally important, precision is 0.48: more than half of flagged students ultimately completed. That proportion is not a model failure in a support context some may have benefited from outreach but it makes punitive or automated uses indefensible. The appropriate intervention is low-stakes, optional, and supportive.

The 6.5-point completion increase and 5.2-point retention increase are educationally meaningful, especially because the service targeted early weeks when tutoring, payment assistance, and study planning remain feasible. The process trend supports the proposed pathway: alerts expanded timely advisor attention. Yet the historical comparison has clear limits. The later cohort may have experienced different instructors, post-pandemic learning habits, curriculum changes, or parallel support initiatives. Covariate adjustment and weighting reduce measured imbalance but cannot remove unobserved confounding. A future stepped-wedge rollout or randomized encouragement design should test causal effects while preserving access to standard support.

Critically, the model's strongest features reflect institutional systems as well as students. Attendance and LMS activity may fall because of connectivity, disability accommodation, caregiving, or timetabling barriers. Fee delays can reflect university processing. The system must therefore trigger investigation of barriers rather than a deficit label. Regular calibration by programme, error audits across groups, threshold review, and a clear student-facing notice are required. Dashboard uptake should also be measured: a highly accurate score has no value if advisors lack time or referral options. These safeguards explain why performance metrics and service capacity must be evaluated together.

5.1 Critical comparison with past work

The model's held-out AUC of 0.87 lies within, and toward the upper end of, results reported in educational data-mining reviews, where performance commonly varies by target and setting. It exceeds the 0.78 logistic baseline observed here, echoing studies in which nonlinear ensemble methods improve predictions using blended academic and behavioural data. However, comparisons require restraint. Many published AUC values are produced with random cross-validation, whereas this study reserves a subsequent term, a more realistic but harder test under temporal drift. Its score therefore should not be framed as an absolute technical advance over a course-specific model with a different label, prevalence, or split.

Relative to LMS-only work, this study's contribution is the combination of academic engagement with ERP finance and administrative measures. The 0.12 importance of fee delay does not make finance the dominant driver; rather, it provides a distinct signal that may point toward an actionable institutional response. This aligns with calls to move beyond clickstream reductionism and account for the broader conditions of participation. Conversely, the dominance of prior GPA and attendance reinforces a long-standing finding: much predictive power derives from stable historical indicators. Such indicators can improve triage but also risk tracking prior inequality, so accuracy cannot be the sole design objective.

Past early-alert evaluations have reported mixed evidence because prediction, notification, and intervention are often conflated. The present paper separates these stages through a ROC assessment, a documented advising process, and outcome comparisons. The 18% outreach rate in Figure 2 is useful implementation evidence absent from many benchmark papers. Still, it is incomplete: contact logged is not contact received, and receipt is not service uptake. The 48% precision also suggests that a blanket outreach script could waste scarce capacity or stigmatize students. Prior work recommends supportive language, choice, and advisor contextualization; those principles were operationalized here through human review.

Ethics literature further challenges a narrow performance race. Slade and Prinsloo and later governance work caution that surveillance, opacity, and purpose creep can undermine trust. This study responds with de-identification, purpose limitation, non-automated action, and small-cell suppression, but it does not solve all fairness questions. Excluding demographic variables does not eliminate proxy discrimination because programme, attendance, and financial signals may correlate with protected status. Future studies should report subgroup calibration, false-positive and false-negative rates, student perceptions, and differential benefit from referrals. Participatory review with students and advisors is as important as model retraining.

Finally, compared with explainability studies, the feature chart in Figure 4 is deliberately modest. Global permutation importance helps an institution audit reliance on variables; it does not establish why an individual is at risk or which intervention will work. A next generation of research should combine predictive ranking with causal evaluation of interventions: for whom does tutoring, financial counselling, or peer mentoring improve outcomes? In short, the present results support an AI-augmented service model, while past work counsels that portability, causal attribution, and justice must remain open empirical questions.

6. Policy Implications

The findings of this study have important implications for institutional leaders, policymakers, and higher education regulators seeking to enhance student success through responsible digital transformation. Rather than viewing ERP systems solely as administrative tools, universities should adopt them as strategic platforms supporting evidence-based academic decision-making and student-centred services.

Institutional policies should encourage the integration of predictive analytics within existing ERP infrastructures while ensuring transparency, accountability, and human oversight throughout the decision-making process. Universities should establish multidisciplinary AI governance committees comprising academic administrators, faculty members, data scientists, legal experts, and student representatives to periodically evaluate model performance, fairness, and ethical compliance.

From a national perspective, the study aligns closely with the objectives of India's National Education Policy (NEP 2020), which advocates technology-enabled learning, personalized education, digital governance, and improved student retention. Regulatory bodies such as the University Grants Commission (UGC) and the All India Council for Technical Education (AICTE) may consider developing comprehensive guidelines governing the ethical use of Artificial Intelligence, learning analytics, and institutional data within higher education.

7. Future Research

The present investigation provides empirical evidence regarding the integration of Artificial Intelligence within university ERP systems; however, several promising directions remain for future research. Multi-institutional studies involving universities from different geographical regions would enhance the generalisability of the proposed framework and facilitate comparative evaluation across diverse educational contexts.

Future investigations may explore advanced machine learning approaches, including deep learning, graph neural networks, federated learning, reinforcement learning, and generative AI, to improve predictive performance while preserving data privacy. The integration of multimodal educational data such as classroom participation, social interactions, behavioural analytics, and student well-being indicators could further strengthen predictive capability.

8. Conclusion

This empirical study shows that predictive analytics embedded in a university ERP can convert dispersed operational data into a timely, advisor-mediated student-support workflow. Using 2,544 student-term records, the final gradient-boosting model achieved AUC 0.87 and recalled 76% of eventual non-completers on a temporally held-out term. The AI-supported cohort also recorded higher completion, retention, and GPA than the preceding cohort after adjustment for observed baseline differences. ERP signals such as fee-clearance delay added contextual information alongside the more influential indicators of prior performance, attendance, and LMS activity.

The evidence supports cautious adoption, not automated decision making. The observed cohort improvements are consistent with earlier and more targeted support, but the quasi-experimental design cannot isolate causal impact from concurrent institutional change. False positives, proxy bias, data drift, and limited advisor capacity require ongoing governance. Universities should use transparent risk bands, human review, clear data-use communication, opt-in or appeal-compatible support pathways where appropriate, and subgroup performance audits. Future multi-site, stepped-wedge research should test which specific interventions convert predictions into equitable gains. The central conclusion is that AI contributes to student success only when it is embedded in accountable relationships and adequately resourced support services.

References

[1] A. R. S. M. A. J. Tinto, “Dropout from higher education: A theoretical synthesis of recent research,” Review of Educational Research, vol. 45, no. 1, pp. 89–125, 1975.

[2] V. Tinto, Completing College: Rethinking Institutional Action. Chicago, IL, USA: Univ. of Chicago Press, 2012.

[3] S. Slade and P. Prinsloo, “Learning analytics: Ethical issues and dilemmas,” American Behavioral Scientist, vol. 57, no. 10, pp. 1510–1529, 2013.

[4] B. Williamson, “The hidden architecture of higher education: Building a big data infrastructure for the ‘smarter university’,” Int. J. Educational Technology in Higher Education, vol. 15, art. 12, 2018.

[5] C. Romero and S. Ventura, “Educational data mining: A review of the state of the art,” IEEE Trans. Systems, Man, and Cybernetics C, vol. 40, no. 6, pp. 601–618, 2010.

[6] A. M. Shahiri, W. Husain, and N. A. Rashid, “A review on predicting student’s performance using data mining techniques,” Procedia Computer Science, vol. 72, pp. 414–422, 2015.

[7] R. S. J. d. Baker and P. S. Inventado, “Educational data mining and learning analytics,” in Learning Analytics, New York, NY, USA: Springer, 2014, pp. 61–75.

[8] G. Siemens and P. Long, “Penetrating the fog: Analytics in learning and education,” EDUCAUSE Review, vol. 46, no. 5, pp. 30–40, 2011.

[9] K. E. Arnold and M. D. Pistilli, “Course signals at Purdue: Using learning analytics to increase student success,” in Proc. 2nd Int. Conf. Learning Analytics, 2012, pp. 267–270.

[10] S. Alyahyan and D. Düştegör, “Predicting academic success in higher education: Literature review and best practices,” Int. J. Educational Technology in Higher Education, vol. 17, art. 3, 2020.

[11] A. Saa, M. Al-Emran, and K. Shaalan, “Factors affecting students’ performance in higher education: A systematic review,” in Proc. Int. Conf. Advanced Intelligent Systems, 2019, pp. 1–8.

[12] R. A. Calvet Liñán and J. C. Juan Pérez, "Educational data mining and learning analytics: Differences, similarities, and time evolution," RUSC, vol. 12, no. 3, pp. 98–112, 2015.
[13] R. M. A. P. Costa et al., "Educational data mining: A survey on recent advances," Wiley Interdisciplinary Reviews: Data Mining, vol. 7, no. 2, 2017.
[14] S. B. Kotsiantis, "Use of machine learning techniques for educational proposes: A decision support system for forecasting students' grades," Artificial Intelligence Review, vol. 37, pp. 331–344, 2012.
[15] J. Gardner and C. Brooks, "Student success prediction in MOOCs," in Proc. 5th Int. Conf. Learning Analytics, 2015, pp. 138–143.
[16] A. Polyzou and G. Karypis, "Grade prediction with models specific to students and courses," Int. J. Data Science and Analytics, vol. 2, pp. 159–171, 2016.
[17] J. S. Jayaprakash, E. W. Moody, E. J. M. Lauría, J. R. Regan, and J. D. Baron, "Early alert of academically at-risk students," J. Learning Analytics, vol. 1, no. 1, pp. 6–47, 2014.
[18] A. S. M. A. Lonn, S. J. Teasley, and S. D. Krumm, "Who needs help? Predicting student success in higher education," J. Learning Analytics, vol. 2, no. 3, pp. 126–152, 2015.
[19] M. K. H. Wolff, Z. Zdrahal, A. Nikolov, and M. Pantucek, "Improving retention: Predicting at-risk students," IEEE Trans. Learning Technologies, vol. 7, no. 2, pp. 173–188, 2014.
[20] N. A. M. B. Rienties, B. Tempelaar, and W. Gijselaers, "A longitudinal study of online learning task design," Computers & Education, vol. 89, pp. 52–63, 2015.
[21] R. Ferguson, "Learning analytics: Drivers, developments and challenges," Int. J. Technology Enhanced Learning, vol. 4, no. 5/6, pp. 304–317, 2012.
[22] J. P. M. Knox, "The datafication of education: Ethical concerns," Learning, Media and Technology, vol. 44, no. 4, pp. 1–15, 2019.
[23] S. Drachsler and W. Greller, "Privacy and analytics It's a DELICATE issue," in Proc. 6th Int. Conf. Learning Analytics, 2016, pp. 89–98.
[24] P. Prinsloo and S. Slade, "Student vulnerability, agency and learning analytics," in Proc. 5th Int. Conf. Learning Analytics, 2015, pp. 159–163.
[25] A. Holzinger, C. Biemann, C. S. Pattichis, and D. B. Kell, "What do we need to build explainable AI systems?" arXiv:1712.09923, 2017.
[26] D. Tiefenthaler and D. Schumacher, "Student perceptions of privacy principles for learning analytics," Educational Technology Research and Development, vol. 64, pp. 923–938, 2016.
[27] G. D. Kuh, High-Impact Educational Practices. Washington, DC, USA: AAC&U, 2008.
[28] T. Dietz-Uhler and J. E. Hurn, "Using learning analytics to predict and support student success," J. Interactive Online Learning, vol. 12, no. 1, pp. 17–26, 2013.
[29] N. Selwyn, Should Robots Replace Teachers? Cambridge, UK: Polity, 2019.
[30] UNESCO, Recommendation on the Ethics of Artificial Intelligence. Paris, France: UNESCO, 2021.