© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723165

An Ethical Governance Framework for Early-Warning Systems in Indian Higher Education: Moving from Student Risk Prediction to Student Support

ADITYA SAGAR
St. Joseph's University, Bengaluru, Karnataka, India

Abstract- Background: Student dropout in higher education is shaped by academic performance, engagement, family circumstances, financial pressure and the quality of institutional support, often acting together rather than in isolation. Universities increasingly use learning analytics and early-warning systems to flag students who may need help, but these systems can also reduce students to a risk score, gather sensitive data without a clear purpose, and drive high-stakes decisions without adequate human review.

Purpose: This conceptual paper proposes an ethical governance framework for early-warning systems in Indian higher education, arguing that such systems should be designed as support architecture rather than surveillance infrastructure.

Methods: Following an established conceptual-research approach, the paper synthesises literature on student retention, learning analytics, educational data mining and student-data privacy to build a structured model that separates an indicator, a risk signal and an institutional decision, with human judgement placed between the second and the third.

Results: The framework rests on six governance principles — data minimisation and purpose limitation, transparency and student agency, mandatory human oversight, restrictions on high-stakes automated decisions, bias monitoring and accountability, and mechanisms for correction and student recourse. These are operationalised through a seven-stage support architecture and a practical institutional pilot in which academic and engagement indicators trigger human-led outreach rather than automatic academic or disciplinary consequences.

Conclusions: Early-warning systems in Indian higher education should be judged not by how many students they classify as "at risk," but by whether they help institutions notice difficulties sooner, respond with more care, and protect student dignity and agency throughout.

Index Terms- AI governance; early-warning systems; human-in-the-loop systems; learning analytics; responsible AI systems design

I. INTRODUCTION

Dropout is rarely caused by one thing. A student's grades might slip around the same time attendance drops, money gets tight at home, family expectations pull in a different direction, or the institution itself offers too little support to notice. Framing dropout as simply a matter of poor academic ability misses most of what is actually going on.

Take a student from a financially comfortable business family. Money isn't the issue. His parents, though, never finished their own education and expect him to start contributing to the family business sooner rather than later. College slips down his list of priorities. He misses one midterm, then another. His attendance falls below the cutoff. Backlogs pile up. By the time anyone at the institution registers that something is seriously wrong, he has already been quietly filed away as a student who "isn't serious about studies." The academic warning signs were real enough — they just didn't explain what was actually happening.

That gap is the problem this paper is about: an institution can hold more data on a student than ever before and still understand that student less. Viewed as a sociotechnical system, an early-warning system is a pipeline with both computational and human components — data ingestion, pattern detection, human interpretation and institutional action — and a failure at any stage of that pipeline can undo the value of the rest.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723165

Learning analytics and educational data mining have opened up a way to close some of that gap. Universities now sit on records of attendance, grades, assignment submissions and platform engagement, and research has shown that patterns in this data line up with academic failure and dropout closely enough to support earlier intervention [1].

Spotting a pattern, though, is not the same as acting on it well. A model can flag a rise in risk without saying anything about why the risk exists or what kind of help would actually be useful. A student whose attendance is falling might need tutoring, a conversation about money, counselling, a more flexible timetable, disability accommodations, or just someone they trust checking in. The same data point can sit underneath very different stories.

That is where the governance question comes in. Treat a model's output as a verdict, and students risk being watched, labelled or "helped" in ways that don't fit their actual situation. Avoid analytics altogether, and the institution loses a chance to catch problems while they're still manageable.

So the question worth asking isn't just whether universities can predict who is likely to drop out. It's how they can use early-warning information to offer timely support without turning students into data points or handing institutional decisions over to an algorithm.

This paper works through that question. Its central claim is that early-warning systems should function as support architecture, not surveillance infrastructure — systems built to flag patterns worth a human's attention, after which trained staff look into the context and decide what, if anything, should happen next.

That distinction matters a great deal in India specifically. The National Education Policy 2020 sets a target of raising the Gross Enrolment Ratio in higher education to 50% by 2035, alongside a stated commitment to access, equity and inclusion [2]. Growing enrolment without building the retention and support systems to match risks bringing students into higher education who then aren't equipped to finish it.

The paper's contribution is fourfold. It treats dropout as an institutional problem rather than an individual failing. It draws a clearer line between an indicator, a risk signal and a decision, which clarifies what analytics should and shouldn't be asked to do. It builds a governance framework around privacy, transparency, oversight, fairness and recourse, expressed as a set of design requirements a system architect could implement directly. And it translates all of that into a system architecture and pilot deployment plan that Indian institutions could realistically run without needing to build sophisticated AI infrastructure first.

II. RELATED WORK AND CONCEPTUAL APPROACH

This is a conceptual framework paper, not an empirical study of a predictive model. Following Jaakkola's account of conceptual research [3], it pulls together established ideas from adjacent fields and uses them to build a structured model for a problem that prediction alone doesn't solve. What follows is model-building work: it connects retention research, learning analytics and responsible data governance into an institutional architecture for student support.

A. Student Dropout as a Multidimensional Phenomenon

Research on dropout has long moved past the idea that students leave purely because of their own shortcomings. Tinto's foundational work [4] framed departure as something that happens in the interaction between a student and their institution — a framing that keeps attention on the relationship between the two rather than on the student in isolation.

The learning-analytics literature backs this up. de Oliveira et al. [1] group the variables used in dropout research into academic, personal and external categories, spanning grades, engagement behaviour and features of the institutional environment.

That has a direct implication for how early-warning systems should be built: not around a single number like attendance or GPA. A poor grade on one assignment doesn't mean a student is heading for the exit. Falling attendance combined with incomplete credits and repeated failed assessments tells a fuller story — though even then, the right response is to look closer, not to treat the pattern as a forecast.

The model developed here treats dropout risk as something that depends heavily on context, not as a trait that belongs to a student.

B. Learning Analytics and Early-Warning Systems as Sociotechnical Systems

Learning analytics, broadly, is about measuring, collecting, analysing and reporting data on learners and their environments in order to improve those environments. In higher education this has opened up new possibilities for advising, support and planning. Early-warning systems are one application of that broader field, working on a fairly simple logic: watch for signals, flag elevated risk, step in early, track what happens next.

There's decent evidence that this can work. de Oliveira et al.[1] found widespread use of data-mining techniques to identify students at risk of dropout or failure. More recent work continues to explore machine-learning models and real-time intervention.

But there's a real limit to what that evidence shows: identifying risk is not the same as proving an intervention helps. Sønderlund et al.'s [5] systematic review of learning-analytics interventions found the evidence for their effectiveness thin, and the quality of the underlying research inconsistent.

That's the crux of it. Prediction tells you who might need attention. Intervention is a separate question — what should actually be done. Getting from one to the other requires governance, not just a better model.

III. THE GOVERNANCE PROBLEM IN AI-ENABLED EARLY-WARNING SYSTEMS

More available data creates a pull toward using all of it. But just because a variable can be collected doesn't mean there's a good reason to collect it.

Slade and Prinsloo [6] lay out the ethical challenges here — privacy, informed consent, how data gets interpreted, classification, and the risk of surveillance. Their argument is that institutions need to reckon with the power imbalance between themselves and their students, rather than assuming that knowing more about a student is automatically good for that student. Prinsloo and Slade [7] push further on student vulnerability and agency, warning against paternalistic uses of student data. Roberts et al. [8] found students themselves worried about privacy, inequity, bias and the risk that learning analytics could chip away at their independence if they're not properly involved in how it's used.

So the governance problem isn't really about whether the data is accurate. It's about whether the data should have been collected at all, whether students understand what it's used for, whether the conclusions drawn from it are justified, and whether students have a real way to push back.

A. Data Minimisation and Purpose Limitation

A well-designed early-warning system collects the minimum it needs for a clearly stated support purpose. That reframes the design question from "what could we collect about this student" to "what do we actually need to offer timely, relevant support."

Attendance is plausibly relevant to disengagement. Assessment results are plausibly relevant to academic difficulty. Repeated non-submission speaks to engagement. Unrelated personal information shouldn't be folded into the risk model just because it happens to be sitting in a database somewhere.

This also guards against function creep. Data gathered to support students shouldn't quietly end up being used for disciplinary, commercial or reputational purposes it was never meant for. Prinsloo et al. [9] make a related point: student-data privacy isn't a problem technology alone can solve — it needs institutional governance, accountability and attention to the social context data sits within.

B. Transparency and Student Agency

Students should know an early-warning system exists, roughly what kind of information it draws on, and how that might affect the support they're offered.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723165

That doesn't mean publishing the algorithm. It means students should be able to understand, in plain terms: what kinds of information might be considered; why that information matters; who reviews a risk signal; what kind of support might follow; what the data will not be used for; and how they can get inaccurate information fixed.

This matters because, without it, students are liable to read institutional monitoring as covert surveillance rather than support. Pardo and Siemens [10] treat privacy, trust, accountability and transparency as tied together in learning analytics, and Ifenthaler and Schumacher [11] show students hold specific expectations about controlling their own data and how it's shared.

IV. SYSTEM DESIGN: INDICATOR, RISK SIGNAL AND DECISION

The core contribution of this paper is a three-part distinction.

An indicator is a single observable data point — a sustained drop in attendance, repeated missed assessments, a sudden dip in grades, incomplete credits piling up, or falling engagement with the institution's learning platforms. On its own, an indicator describes something. It doesn't establish that a student is at risk of leaving.

A risk signal shows up when several indicators line up into a pattern. Falling attendance alongside missed assessments and slipping grades is a stronger signal than any one of those on its own. But even here, the right way to read it is as a case for institutional attention — not a forecast that the student will drop out.

A decision is what the institution actually does about it: reaching out to the student, arranging advising, offering tutoring, referring them to counselling, sharing financial-support information, reviewing their workload, or simply talking through attendance barriers.

The framework proposed here inserts a mandatory human-review step between the risk signal and the decision. The chain runs: Indicators → Risk Signal → Human Review → Support Decision — not Indicators → Algorithmic Label → Automatic Consequence, as shown in Fig. 1. That matters because the same risk signal can come from very different underlying situations, and only a human reviewing the context can tell them apart.

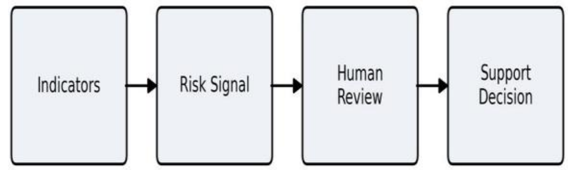

Fig. 1. Governance pipeline from indicator to support decision.

V. PROPOSED GOVERNANCE FRAMEWORK: DESIGN PRINCIPLES

The framework rests on six principles that work together.

A. Design Principle 1: Data Minimisation and Purpose Limitation

Only data with a clear, demonstrable link to the support objective should feed into the system. Institutions need to define why they're collecting data, which variables are allowed, how long records are kept, who can see them, and what counts as an acceptable or unacceptable use. The system should never quietly turn into a general student-surveillance tool.

B. Design Principle 2: Transparency and Student Awareness

Students should be told the system exists and roughly how it works. A university might put it simply: "The university uses selected academic and engagement indicators to identify students who may benefit from additional support. A risk signal does not mean a student is considered likely to leave the university. Any support action is reviewed by a member of the university community, and students may request clarification or correction of inaccurate information." That kind of statement turns hidden monitoring into something accountable.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723165

C. Design Principle 3: Mandatory Human Oversight

No high-impact decision should follow automatically from an algorithm generating a risk signal. A faculty advisor, counsellor, student-support officer or review team needs to look at the context first.

This isn't a box-ticking exercise — it serves a real purpose, which is contextual interpretation. UNESCO's Recommendation on the Ethics of Artificial Intelligence puts human oversight, transparency, accountability, privacy, fairness and non-discrimination at the centre of responsible AI use [12], and this framework applies those same ideas to the narrower setting of higher-education early-warning systems.

D. Design Principle 4: No High-Stakes Automated Decisions

A risk signal should never automatically trigger disciplinary action, exclusion from a programme, suspension, denial of opportunities, financial penalties, or any other consequential decision about a student's standing. The system exists to support students, not to punish them. Nobody should end up academically worse off simply because a model flagged them.

E. Design Principle 5: Bias Monitoring and Institutional Accountability

Historical institutional data carries historical inequalities with it. If certain groups of students have systematically had worse outcomes because of structural disadvantage, a model trained on that history is likely to reproduce the same pattern.

Bias needs checking at every stage — data, model, risk signal, human review, intervention, outcome. The question isn't only "how accurate is this system," but who gets flagged, who gets missed, who actually receives support once flagged, whether certain groups are flagged more than others, and whether different groups end up with different outcomes from the same intervention. That moves the evaluation beyond accuracy and into institutional fairness.

F. Design Principle 6: Correction, Contestability and Student Recourse

Student records aren't always right. An attendance record can be thrown off by an administrative error. A missed assessment might have been formally excused but never properly logged. A student's circumstances can shift in ways the institution's data simply doesn't capture.

Students need a real way to look at the information held about them, flag mistakes, request corrections, add context, and push back on conclusions that don't fit their situation. A risk signal that can't be challenged stops being a temporary prompt for support and starts becoming a permanent label.

VI. SYSTEM ARCHITECTURE

The six principles translate into a seven-stage institutional process.

A. Stage 1: Data Collection

Collect a limited, justified set of indicators — academic performance, engagement behaviour, interactions with institutional support, and voluntary student requests for help. Sensitive personal characteristics shouldn't be pulled in just because they happen to be available.

B. Stage 2: Signal Generation

The system looks for patterns rather than reacting to isolated events. One missed assessment is an indicator. Repeated missed assessments plus falling attendance is a stronger signal. Add academic decline on top of that and you have a signal worth elevating for review. Where exactly the threshold sits should be worked out and validated by each institution rather than borrowed wholesale from elsewhere.

C. Stage 3: Human Review

A designated team looks at the signal and asks the questions a model can't: Is the data actually accurate? Is there an obvious explanation sitting in plain sight? Is the student already getting some kind of support? What might actually help here? Is stepping in likely to do more good than harm?

D. Stage 4: Student Engagement

The first contact with a student should feel supportive, not accusatory. Rather than "you've been flagged as a dropout risk," something closer to: "We noticed you might be facing some academic or attendance difficulties — would you like to talk about what's going on and whether there's support that could help?" That keeps the student's dignity intact and leaves room for them to explain what's actually happening.

E. Stage 5: Targeted Intervention

Support should match the need that's actually there — tutoring, advising, mentoring, counselling, financial guidance, a conversation about workload or timetable, accessibility support, peer support, career guidance. A detection system with no intervention capacity behind it isn't much use to anyone.

F. Stage 6: Outcome Monitoring

The institution should track whether students actually took up the support offered and whether their situation improved. The number of students flagged isn't the metric that matters — whether the intervention actually helped them stay and succeed is.

G. Stage 7: Audit and Review

The whole system needs periodic review: false positives, false negatives, disparities across demographic groups, how many students took up the support offered, student feedback, data quality, how staff are actually making decisions, and any consequences nobody anticipated. That closes the loop — Detection → Intervention → Outcome → Audit → System improvement, illustrated as a closed control loop in Fig. 2.

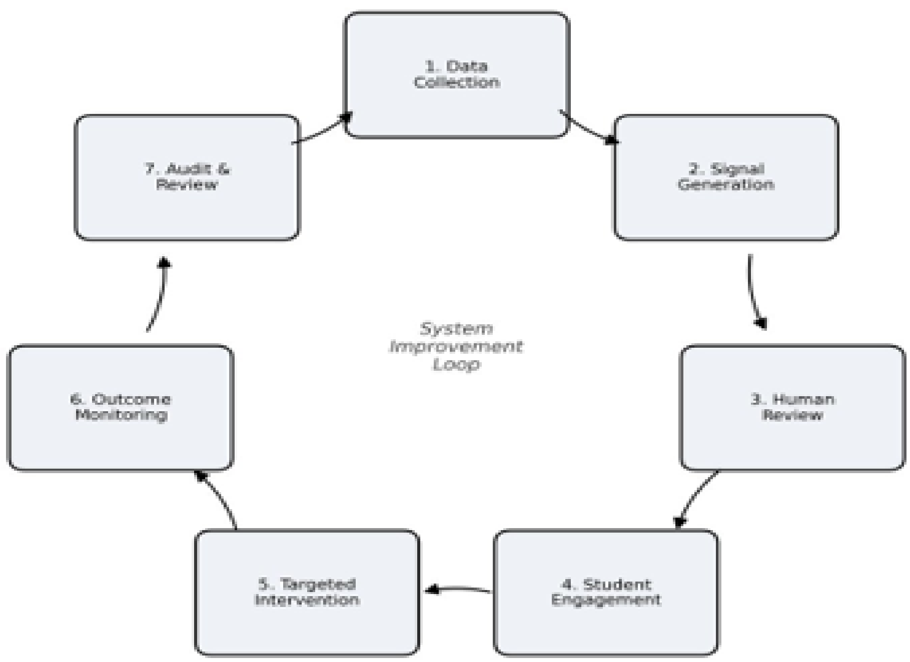

Fig. 2. Seven-stage system architecture as a closed improvement loop.

Table 1. Illustrative indicator categories for an institutional pilot

Category	Illustrative indicators	Purpose
Academic	Assessment performance; failed credits; grade trajectory	Identify academic difficulty
Engagement	Attendance; repeated non-submission	Identify disengagement
Institutional	Use of academic support; repeated administrative issues	Identify unmet support needs

Student-reported	Voluntary requests for support	Capture information unavailable through administrative data

VII. IMPLEMENTATION PLAN: A PILOT DEPLOYMENT FOR INDIAN UNIVERSITIES

None of this requires a full-scale rollout on day one. A university could start with a single department or one incoming cohort.

Phase 1: Define the Purpose. Put in writing that the system exists to support student retention, progression and wellbeing — nothing else.

Phase 2: Select Indicators. Pick a small number of indicators that are academically and institutionally meaningful. Resist the urge to throw in every variable available.

Phase 3: Establish the Human Review Team. A small, multidisciplinary group — faculty, academic advisors, counselling or student-support staff, and someone from institutional data — with clearly defined responsibilities from the outset.

Phase 4: Build the Support Capacity. Work out what support is actually available before generating a single alert. There's little point identifying students who need help if there's nothing to offer them.

Phase 5: Begin With Low-Stakes Outreach. Make first contact voluntary and supportive. The goal is understanding what's going on, not confirming that the algorithm was right.

Phase 6: Track Outcomes. Record how many signals were generated, how many were reviewed, how many led to intervention, what kind, how students responded, whether their indicators improved, what feedback they gave, and where the system got it wrong.

Phase 7: Independent Review. Someone outside the immediate implementation team should periodically check whether the system is actually being fair. This matters because staff who work with the system every day can start treating its outputs as objective fact rather than a starting point for inquiry.

VIII. INSTITUTIONAL AND SYSTEM-LEVEL IMPLICATIONS

For university administrators, this framework turns early-warning systems from a software-procurement question into a governance question. Buying predictive software doesn't, on its own, create a student-support system. Institutions still need governance rules, sound data practices, clearly assigned staff roles, real intervention capacity, escalation protocols, a plan for communicating with students, and ways to monitor how it's all working. How the system is organised may matter just as much as how well the model predicts.

The framework also pushes institutions to bring together support functions that usually sit in separate offices. A student's grades, attendance and administrative history are often scattered across different systems that don't talk to each other. The problem, more often than not, isn't a lack of information — it's the absence of any mechanism for seeing how the pieces relate.

Responsible use of this kind of technology means thinking through its consequences before measuring its accuracy: what kinds of student behaviour become visible, who gets to define risk, who benefits from being classified, who bears the cost when the system gets it wrong, whether students can challenge it, whether staff can override it, and whether the institution can actually explain why a given student was contacted.

For Indian higher-education policy more broadly, the sector's expansion and digitalisation bring an opportunity alongside an obligation. NEP 2020 aims for a 50% Gross Enrolment Ratio by 2035 while emphasising access, equity and inclusion [2]. The policy goal shouldn't just be encouraging universities to adopt predictive analytics — it should be encouraging student-support infrastructure that treats analytics as one part of a larger, fundamentally human system.

IX. DISCUSSION

The argument running through this paper is that an early-warning system shouldn't be judged mainly on how accurately it predicts.

A model with strong predictive performance can still produce a bad outcome for the institution — if students get mislabelled, if there's no support to offer once they're flagged, if staff trust the model blindly, if sensitive data gets collected without a clear reason, if students are kept in the dark, if a "high risk" label ends up feeding into punitive decisions, or if some groups are disadvantaged more than others by how the system works.

A much simpler system, on the other hand, can generate real value if it just helps a university notice a struggling student sooner and connect them to the right kind of help.

That lines up with the wider literature. Sønderlund et al. [5] point to how thin and uneven the evidence base is for learning-analytics interventions. de Oliveira et al. [1] show how much data and how many predictive techniques get used in dropout research, while flagging privacy concerns of their own. The ethics literature, more broadly, keeps coming back to the same point — student-data systems sit inside institutional power relationships, not in some neutral technical space.

The framework here leans on a human-in-the-loop model, a design pattern increasingly standard in responsible AI systems engineering wherever automated outputs feed into consequential decisions. The analytical side handles detection, prioritisation and documentation. The human side handles context, judgement, relationships and accountability. That division matters. The point isn't to remove human judgement from the process — it's to give that judgement better information, earlier.

X. CONTRIBUTION TO RESPONSIBLE AI SYSTEMS DESIGN

This framework contributes to the learning-analytics conversation along three lines.

It moves the conceptual endpoint from prediction to support. Much of the existing literature asks whether models can identify students likely to have a bad outcome. This paper pushes the chain a step further, toward: Prediction → Human interpretation → Support → Outcome, rather than stopping at Prediction → Classification.

It reframes early-warning systems as governance mechanisms rather than purely technical ones. That puts data collection, transparency, accountability, human review and student recourse at the centre of how these systems get designed, rather than treating them as add-ons.

And it reframes what a risk signal actually tells an institution. A risk signal doesn't only say something about the student — it can say something about the institution too. Repeated non-submission might reflect a disengaged student, but it could just as easily point to weak teaching support, an assessment design that's hard to access, timetable clashes, thin advising, fragmented services, or resources stretched too far. Read this way, a risk signal is a prompt for institutional inquiry as much as it is a flag on a student.

XI. LIMITATIONS

This is a conceptual paper. It doesn't validate the framework empirically. No institutional dataset was used, and no predictive model was built or tested, so the paper can't say anything about how well any particular combination of indicators would actually predict dropout, or whether the proposed framework would reduce it in practice.

The framework should be read as a governance and implementation proposal that still needs empirical testing. Future work could pilot it across different types of Indian institutions and look at which indicators produce useful early signals, how often false positives crop up, how students experience being

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723165

monitored this way, whether human review actually improves the quality of interventions, whether interventions move the needle on dropout or progression, whether particular groups of students get flagged disproportionately, how institutions balance privacy against support, and what kind of governance builds the most trust with students.

Because dropout patterns and institutional support structures differ so much from one university to another, longitudinal, multi-institutional research would be especially useful here.

XII. FUTURE SCOPE

A few directions for empirical work follow naturally from the framework:

RQ1: Which combinations of academic and engagement indicators give the most useful early signal of student support needs in Indian higher education?

RQ2: Does mandatory human review actually cut down on inappropriate interventions generated by automated early-warning systems?

RQ3: How do students experience transparency, privacy and agency when universities use learning analytics for retention purposes?

RQ4: Do early-warning interventions produce different outcomes across socioeconomic and demographic groups?

RQ5: Which institutional characteristics determine whether an early-warning system produces meaningful support outcomes rather than just flags?

RQ6: Can explainable predictive models build more trust and produce better interventions than opaque, high-performance ones?

Working through these questions would take this conceptual framework toward something empirically testable.

XIII. CONCLUSION

Dropout isn't purely a prediction problem. It's a problem of understanding, institutional responsiveness and acting on that understanding in time.
Universities now have the technical means to spot patterns tied to student disengagement. The harder question is whether they can use that capability responsibly.

An early-warning system shouldn't try to tell a university that a student is destined to leave. Its job is smaller and more useful than that: to tell the institution it might be time to ask what's going on.

That's why this framework insists on a boundary between analytics and institutional decisions. Indicators generate signals. Signals prompt human attention. Human attention builds understanding. Understanding is what makes real support possible.

Data can show patterns, but it can't show circumstances. That distinction matters especially for Indian higher education right now, as institutions expand access, roll out digital systems, and work toward NEP 2020's goals. A responsible early-warning system can genuinely support retention and student success — but only inside an institutional structure that protects privacy, stays transparent, watches for bias, allows for correction, and keeps a human ultimately responsible.

The measure of success, in the end, shouldn't be how many students get classified as "high risk." It should be whether more of them get the right support before it's too late.

REFERENCES

[1] C. F. de Oliveira, S. R. Sobral, M. J. Ferreira, and F. Moreira, “How does learning analytics contribute to prevent students' dropout in higher education: A systematic literature review,” *Big Data and Cognitive Computing*, vol. 5, no. 4, Art. no. 64, 2021, doi: 10.3390/bdcc5040064. MDPI

DOI: https://doi.org/10.64388/IREV10I3-1723165

[2] Ministry of Education, Government of India, *National Education Policy 2020*. New Delhi, India: Government of India, 2020.

[3] E. Jaakkola, "Designing conceptual articles: Four approaches," *AMS Review*, vol. 10, pp. 18–26, 2020, doi: 10.1007/s13162-020-00161-0. Springer

[4] V. Tinto, "Dropout from higher education: A theoretical synthesis of recent research," *Review of Educational Research*, vol. 45, no. 1, pp. 89–125, 1975, doi: 10.3102/00346543045001089. SAGE

[5] A. L. Sønderlund, E. Hughes, and J. Smith, "The efficacy of learning analytics interventions in higher education: A systematic review," *British Journal of Educational Technology*, vol. 50, no. 5, pp. 2594–2618, 2019, doi: 10.1111/bjet.12720. Wiley

[6] S. Slade and P. Prinsloo, "Learning analytics: Ethical issues and dilemmas," *American Behavioral Scientist*, vol. 57, no. 10, pp. 1510–1529, 2013, doi: 10.1177/0002764213479366. SAGE

[7] P. Prinsloo and S. Slade, "Student vulnerability, agency and learning analytics: An exploration," *Journal of Learning Analytics*, vol. 3, no. 1, pp. 159–182, 2016, doi: 10.18608/jla.2016.31.10. Journal of Learning Analytics

[8] L. D. Roberts, J. A. Howell, K. Seaman, and D. C. Gibson, "Student attitudes toward learning analytics in higher education: 'The Fitbit version of the learning world,'" *Frontiers in Psychology*, vol. 7, Art. no. 1959, 2016, doi: 10.3389/fpsyg.2016.01959. Frontiers

[9] P. Prinsloo, S. Slade, and M. Khalil, "The answer is (not only) technological: Considering student data privacy in learning analytics," *British Journal of Educational Technology*, vol. 53, no. 4, pp. 876–893, 2022, doi: 10.1111/bjet.13216. Wiley

[10] A. Pardo and G. Siemens, "Ethical and privacy principles for learning analytics," *British Journal of Educational Technology*, vol. 45, no. 3, pp. 438–450, 2014, doi: 10.1111/bjet.12152. Wiley

[11] D. Ifenthaler and C. Schumacher, "Student perceptions of privacy principles for learning analytics," *Educational Technology Research and Development*, vol. 64, no. 5, pp. 923–938, 2016, doi: 10.1007/s11423-016-9477-y. Springer

[12] UNESCO, *Recommendation on the Ethics of Artificial Intelligence*. Paris, France: UNESCO, 2021. UNESCO