FICCI 100
100 YEARS OF BUILDING TOMORROW

From AI adoption to readiness: Preparing Indian higher education for the AI era

October 2026



The better the question.
The better the answer.
The better the world works.

EY Parthenon
Shape the future with confidence

Foreword

FICCI

India's higher education sector is at a defining inflection point. Artificial Intelligence (AI) is no longer an emerging technology confined to experimentation. It is rapidly becoming a foundational capability that is reshaping teaching, learning, research, governance and institutional competitiveness. As AI transforms economies and societies across the world, the imperative for Indian Higher Education Institutions (HEIs) is no longer whether to adopt AI, but how to translate adoption into meaningful, responsible and large-scale action.

Building upon the insights of the FICCI-EY Parthenon Knowledge Report 2025, this year's report, "From AI adoption to readiness: Preparing Indian higher education for the AI era" examines how the AI landscape has evolved over the past year and assesses the readiness of Indian institutions to respond to this unprecedented technological shift. While encouraging progress has been made, the report also highlights significant implementation gaps between policy aspirations and institutional realities.

Drawing upon insights from academic leaders, policymakers, industry experts, faculty and students, the report presents an honest assessment of AI adoption across the higher education ecosystem. It explores how AI is influencing classrooms, research, governance, employability and institutional decision-making, while also addressing critical issues such as faculty readiness, digital infrastructure, ethical AI, data governance and equitable access.

The report introduces a practical framework centered around five key pillars—democratization of AI, depth of AI adoption, shared digital infrastructure, stronger industry-academia collaboration and global best practices. It also presents actionable playbooks and recommendations for students, faculty, institutional leaders, employers and policymakers to help accelerate India's transition towards AI-ready campuses.

FICCI remains committed to advancing this transformation through policy advocacy, industry collaboration and thought leadership. We believe that with the right vision, partnerships and coordinated action, India has the opportunity not only to prepare its higher education system for the AI era but also to emerge as a global leader in AI-enabled education.

We extend our sincere appreciation to all contributors who shared their insights and experiences. This report is intended to serve as a catalyst for action thereby guiding institutions, educators and policymakers as they navigate the opportunities and responsibilities of AI in higher education.



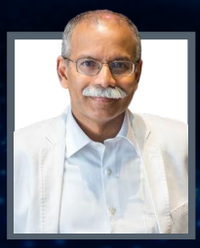
Dr. Vidya Yeravdekar

Chair – FICCI Higher Education Committee & Chancellor, Symbiosis International University



Lt Gen (Dr.) M. D. Venkatesh

Co-Chair – FICCI Higher Education Committee & Chairperson, Manipal University, Jaipur



Prof (Dr). Padmakumar Nair

Co-Chair – FICCI Higher Education Committee & Vice Chancellor, Thapar Institute of Engineering & Technology



Prof (Dr) Rajan Saxena

Hon Advisor – FICCI HE Committee and Former Vice Chancellor, NMIMS

EY-Parthenon

Artificial intelligence is moving rapidly from experimentation to everyday use in higher education. Students are already using AI to learn, create and solve problems; faculty are beginning to incorporate it into teaching and research; and institutions are exploring its potential across academic and administrative functions. Yet adoption alone does not translate into readiness. When students, faculty, administrators and employers move at different speeds, AI can create new forms of misalignment rather than simply improving outcomes.

This is the central challenge for higher education today. Institutions may need to move beyond providing access to tools and ask harder questions about how learning, assessment, research and institutional processes should change in an AI-enabled world. Faculty, too, need to move beyond familiar ways of teaching and evaluating students. The opportunity is not to use AI to replicate existing practices more efficiently, but to reconsider where technology can augment human capability and where judgment, creativity, critical thinking and intellectual ownership must remain distinctly human.

At the same time, the growing global debate around the pace and direction of AI development is an important reminder that technological progress cannot be separated from questions of ethics, accountability and responsible use. For higher education, however, caution cannot mean disengagement. At the grassroots level, institutions must do two things simultaneously: equip students to understand the capabilities, limitations and risks of AI, and enable faculty to use AI responsibly through the provision of tools, training and institutional support.

The findings of this report suggest that the next phase of AI in Indian higher education will therefore be defined less by whether institutions adopt AI and more by how deliberately they do so. The institutions that make the most of this moment will be those that align students, faculty and leadership around a shared understanding of AI; redesign learning rather than simply digitizing it; and create the governance, capabilities and confidence required to experiment responsibly. The objective is not an AI-driven university, but a human-centered institution that uses AI to deepen learning, expand opportunity and strengthen the capabilities that matter most.



Dr. Avantika Tomar

Partner and Learning
and Development Head,
EY-Parthenon India

Contents

Looking back to move forward

Chapter 01 Executive summary 10

Chapter 02 The India campus story: Unfiltered, uncomfortable, undeniable 12

Chapter 03 What the industry actually wants 36

Chapter 04 India vs. the world: Where India stands in ai-driven higher education 58

Chapter 05 Designing the ethical, effective, AI-enabled Indian campus 76

Chapter 06 India's AI imperatives for higher education, 2026-2030 88

Chapter 07 The playbooks, with a focus on equity 102

Chapter 08 The future we need to build 138



Looking back to move forward

Indian higher education has entered the next phase of its AI journey

The 2025 FICCI-EY-Parthenon report established an institutional baseline for understanding AI adoption in Indian higher education, revealing a landscape characterized by rapid uptake but fragmented implementation across tools, departments and isolated pilot initiatives. Developments since its publication have heightened the urgency of that assessment. Public policy initiatives, national skilling efforts and education focused AI infrastructure have advanced significantly, while the evolution of student behavior, faculty practices and institutional systems has proceeded at markedly different rates.

0.1 What the 2025 report set out to do?

A comprehensive view of AI's evolving role in higher education

Future-ready campuses: Our report "Unlocking the Power of AI in Higher Education" examined the implications of AI across teaching and learning, assessment, research, career services, campus administration, curriculum design, transnational education, faculty development, governance and institutional strategy. Its principal contribution was to move the discourse beyond generative AI tools and demonstrate that AI was beginning to reshape the entire academic and administrative operating model of higher education institutions.

A shift in focus from experimentation to institutional readiness

The 2025 survey of 30 Indian higher education institutions assessed progress across six interconnected dimensions: AI tool adoption, curriculum and employability, research and innovation, data and technology infrastructure, faculty capability development and policy and governance. Nine institutions reported no academic AI applications, while 13 had already implemented three or more. Similarly, 12 institutions had established institution-wide AI visions or roadmaps, five had departmental-level policies and a further 12 were in the process of developing governance frameworks.

A clear strategic agenda for Indian higher education institutions

The report identified AI literacy across disciplines, domain-specific curriculum integration, faculty upskilling, strengthened digital and data infrastructure, ethical governance, industry collaboration and student-centered adoption as critical priorities. AI was presented both as a catalyst for improving quality, access and global competitiveness and as a source of emerging challenges related to privacy, bias, transparency and academic integrity.

An institutional perspective rather than a complete picture of campus-level adoption

The evidence base was primarily institution-facing, drawing on inputs from institutional leaders, faculty members and administrators and complemented by case studies and global student data. While this approach provided valuable insight into institutional priorities, initiatives and future plans, it offered only limited visibility into students' independent use of AI, the interpretation of faculty guidance within individual classrooms, the extent to which formal policies were influencing assessment practices and differences in adoption across institutional tiers. These limitations provide the foundation for the student focus group discussions and stakeholder interviews undertaken as part of the 2026 report, which seek to capture a more comprehensive understanding of how AI is being experienced and applied across the higher education ecosystem.

Source: EY-Parthenon analysis

AI adoption accelerated across campuses, exposing a growing gap between activity, readiness and meaningful implementation

0.2 What actually happened since our last report?

Policy and ecosystem momentum accelerated significantly

Policy and ecosystem momentum accelerated significantly. The notification of the Digital Personal Data Protection (DPDP) Rules in November 2025 transformed data protection from an emerging consideration into an immediate implementation priority for institutions deploying AI-enabled systems and services. At the same time, the IndiaAI Mission launched YUVA AI for ALL, a foundational AI literacy program aimed at reaching one crore citizens, while the Ministry of Education's Bharat Bodhan AI Conclave established the Bodhan AI Centre of Excellence and initiated the development of the open and interoperable Bharat EduAI Stack. Collectively, these developments advanced several priorities identified in the 2025 report, including AI literacy, shared digital infrastructure, education-specific governance and stronger collaboration among government, academia and industry.

Institutional adoption progressed along distinct trajectories

While some institutions transitioned from experimentation to structured implementation through new academic programs, teaching initiatives, AI centers and strategic collaborations, others focused primarily on policies, committees and awareness activities without substantively changing curricula, assessment practices, procurement processes or faculty support mechanisms. A third group experienced predominantly bottom-up adoption driven by students and individual faculty members, often operating beyond formal institutional frameworks, while resource-constrained or risk-averse institutions remained largely inactive. The stakeholder interviews suggest that institutions can be more accurately understood through four broad trajectories: implementation, announcement, informal adoption and stalled adoption, recognizing that not all reported activity reflects the same level of progress or preparedness.

Some of the most significant developments occurred beyond institutional visibility

Students adopted publicly available AI tools for tutoring, writing, coding, translation, research support, career preparation and assignment completion, often independently of formal institutional strategies. Faculty responses varied considerably, ranging from comprehensive course and assessment redesign to prohibition, cautious experimentation and procedural compliance. Comparing student and faculty perspectives provides important insight into whether institutional policies were consistently understood, whether AI literacy was being developed through formal educational channels or informal peer networks, and whether the existence of policy frameworks was translating into meaningful changes in classroom practice.

The core findings of the 2025 report were reinforced, while previously unseen challenges became more prominent

The earlier report accurately identified fragmented adoption, limited faculty preparedness, infrastructure constraints and the growing need for ethical governance frameworks. However, developments over the following year brought greater visibility to issues that could not be fully captured through an institution-level survey. These included gray-zone academic use, the emergence of AI as a parallel tutoring system, concerns regarding dependency and diminished learner confidence, uncertainty surrounding faculty authority in AI-enabled learning environments, and the widening gap between polished AI-assisted outputs and demonstrable student understanding. These emerging realities provide the foundation for the unfiltered campus-level perspectives explored in Section 2.

Source: EY-Parthenon analysis

Students are adopting AI by default, making institutional action an urgent necessity

0.3 Why a 2026 "hard reset" report became necessary and why AI is no longer optional?

In 2025, the primary challenge was how Indian higher education institutions could adopt AI in a responsible and strategically aligned manner. By 2026, however, AI had already become embedded in how students learn, produce academic work and prepare for employment, irrespective of whether their institutions had formally adopted it. The challenge is therefore no longer one of introducing AI into higher education, but of governing, shaping and equitably distributing a capability whose use is increasingly becoming routine.

Three asymmetries now characterize the Indian higher education landscape. Policy development has progressed, but often without translating into course-level implementation. Students have adopted AI at a pace that frequently exceeds faculty confidence and preparedness to guide its use, and technological capabilities continue to evolve more rapidly than academic regulations, procurement processes and assessment systems can adapt. The resulting environment is neither one of orderly adoption nor complete resistance, but of growing divergence between formal institutional frameworks and everyday educational practice.

Institutional inaction is not a neutral stance. In the absence of meaningful guidance, students acquire AI-related behaviors and practices through peers, social media platforms and commercial providers. Where assessment approaches remain unchanged, AI may obscure rather than address underlying learning gaps. Where access depends on private subscriptions, existing advantages are amplified, enabling some students to progress more rapidly than others. Similarly, institutions that adopt AI-enabled services without robust data governance frameworks expose themselves to privacy, security and vendor lock-in risks that may not be fully understood or managed. The consequences are reflected in reduced confidence in assessment integrity, widening capability disparities, unmanaged student data and graduates whose polished outputs may exceed their capacity for independent judgment.

The benefits of action extend well beyond risk mitigation. When implemented thoughtfully, AI can provide personalized and multilingual academic support, accelerate feedback cycles, reduce repetitive faculty and administrative workloads, expand access to research and global knowledge resources, and better prepare students for AI-enabled workplaces. The critical distinction is therefore not between adoption and prohibition, but between unmanaged use that substitutes for capability and governed use that strengthens and extends it.

The focus must shift from AI adoption to AI action

Accordingly, the 2026 report moves beyond measuring the number of tools, policies, courses or pilot initiatives in place. Horizontal adoption considers how widely and equitably AI-related capabilities are distributed across students, faculty, institutions and regions, while vertical adoption examines the depth of AI integration within learning, assessment, research, administration and institutional decision-making. The starting point for this reassessment is a simple but important recognition: **AI use has already scaled across the higher education ecosystem, whereas institutional readiness has not yet scaled at the same pace.**

Source: EY-Parthenon analysis

The report is built on perspectives from institutions, faculty, students, administrators and employers across India



Methodology

The 2026 FICCI–EY-Parthenon report is based on a combination of primary and secondary research. The analysis draws on primary interviews with higher education institutions, industry representatives and students, supplemented by a review of global literature, policy documents, institutional case studies and publicly available datasets. Insights from the 2025 FICCI–EY-Parthenon survey have also been incorporated where relevant to track developments in AI adoption and institutional readiness across India's higher education ecosystem.

Primary research covered 23+ higher education institutions across Tier 1, Tier 2 and Tier 3 locations, spanning different institution types, disciplines and regions and included 61+ stakeholder interactions. These comprised one-to-one interviews with faculty, administrators, university leaders and employers, as well as student focus group discussions with four to eight students in each group. The sample included 16 private and seven public/government-aided institutions, of which nine were leading private universities or Institutions of National Importance (INIs) and 14 were regional institutions. It covered nine multidisciplinary institutions, 12 –STEM[1] focused institutions and two non-STEM institutions, with faculty and students across disciplines represented wherever possible. These perspectives have been used alongside secondary evidence to develop and test the report's point of view on the on-ground realities of AI adoption, use and institutional readiness.



AI disclosure

To practice the ethical and effective use of AI that this report advocates, we have also used generative AI as a research and writing aid in select areas, including synthesis, structuring and drafting. Its use has been supported by human review and validation throughout. Key structures and arguments were developed through discussions with the committee and stakeholders across the higher education ecosystem in India, after which AI was leveraged to support research and provide additional depth to expert insights. All analysis, interpretations and points of view presented in the report remain the authors' own and responsibility for the final content rests with FICCI and EY-Parthenon.

1. STEM: Science, technology, engineering and mathematics

Source: EY-Parthenon analysis



Chapter 1

Executive summary

Artificial Intelligence (AI) has moved beyond experimentation to become an embedded reality across India's higher education ecosystem. The central question for institutions is no longer whether AI should be adopted, but how its use can be governed, scaled and aligned with educational outcomes. While policy momentum, institutional announcements and technology adoption have accelerated, readiness across students, faculty, administration and governance systems remains uneven.

AI adoption has outpaced institutional preparedness. Students have integrated AI into their learning workflows, employers increasingly expect graduates to work alongside it, and governments are investing in AI infrastructure and literacy. Yet many institutions still operate with legacy pedagogies, limited faculty preparedness, unclear governance frameworks and insufficient attention to data protection, equity and assessment redesign.

The report identifies five shifts shaping the future of higher education in India.

1 Student AI adoption has become widespread, but learning outcomes remain uncertain

AI has rapidly become part of how students research, write, learn, revise, prepare for examinations and seek academic support. Students increasingly use AI not only for productivity but also as a tutor, language assistant, brainstorming partner and personalized learning companion. For many learners, AI now functions as an always-available support system that complements or, at times, substitutes for traditional educational channels.

However, student usage exists along a continuum rather than within a simple distinction between ethical and unethical behavior. Most student activity falls within a "gray zone" where institutional policies provide limited clarity regarding acceptable use. As a result, students are often defining their own norms of AI usage in the absence of structured guidance.

The report highlights growing concerns around dependency, diminished critical thinking, erosion of confidence in independent work, widening differences between polished outputs and genuine understanding, and the emergence of AI-enabled shortcuts that may obscure underlying skill gaps. These challenges are increasingly becoming questions of pedagogy and assessment design rather than purely academic integrity issues.

2 Assessment and faculty practice are being rethought in response to AI

Assessment is emerging as one of the most important areas of institutional change. As generative AI makes traditional take-home assignments less reliable as evidence of student learning, faculty are experimenting with viva voce examinations, presentations, project-based learning, reflective exercises, in-class assessments and continuous evaluation. The focus is shifting from the final submission toward evidence of understanding, reasoning, application and problem-solving.

Faculty attitudes are also evolving. Early responses were often shaped by concerns around plagiarism and the erosion of foundational skills, but educators increasingly recognize that students will enter workplaces where AI is part of professional practice. The challenge is therefore not to preserve an AI-free classroom, but to determine where AI improves learning, where independent capability must still be demonstrated and how the two can coexist within the same course.

This requires more than tool familiarity. Faculty need to redesign assignments, establish clear boundaries for AI use, develop appropriate rubrics and help students verify and challenge AI-generated outputs. The emerging model is one of guided adoption, in which students can use AI while remaining responsible for the reasoning, evidence and judgment behind their work.

3 Institutional readiness has not kept pace with adoption

A significant challenge lies not only in resistance to AI, but in the appearance of readiness without substantive change. Institutions may report faculty training, AI policies and new initiatives while classroom practice, assessment design and student guidance remain largely unchanged. Workshop participation and policy publication provide limited evidence of meaningful integration.

At the institutional level, enthusiasm for AI frequently exceeds implementation capacity. Governance frameworks, procurement systems, infrastructure, funding mechanisms and leadership capabilities are often insufficiently developed to support responsible adoption at scale. The next stage of readiness may therefore depend less on the number of AI initiatives launched and more on whether universities can build the operating structures that connect policy, classroom practice, data governance, procurement and accountability.

4 The labor market is redefining what it means to be employable

Employers increasingly view AI literacy as a foundational capability rather than a specialized technical skill. Yet the ability to use AI tools is becoming less differentiating as access becomes widespread. Greater value is being placed on capabilities that determine whether AI is used well: problem framing, reasoning, verification, judgment, communication and the ability to apply AI within a specific professional context.

AI is also changing the nature of entry-level work. As routine research, drafting, coding and analysis become increasingly AI-assisted, graduates may lose some of the repetitive tasks through which judgment was traditionally developed, even as employers expect them to exercise that judgment earlier in their careers.

The way forward

India's higher education system stands at an inflection point. AI adoption is already widespread across campuses and workplaces, but widespread use should not be confused with readiness. The task ahead is to convert adoption into improved learning, stronger faculty practice, more credible assessment, responsible governance and better preparation for an AI-enabled labor market.

This may require institutions to move beyond isolated policies, one-off training programs and tool-centric strategies toward a more integrated approach spanning pedagogy, assessment, governance, infrastructure, faculty development, employability and student outcomes. For policymakers, the challenge is to create the shared infrastructure, standards and funding mechanisms needed to extend these capabilities beyond the country's best-resourced institutions.

The defining transition for Indian higher education is therefore from AI adoption to AI readiness, with success ultimately determined not by how extensively AI is used, but by whether it strengthens human capability across the system.

Traditional proxies for talent are also becoming less sufficient. Grades and degrees continue to matter, but the ability to explain how a problem was approached, what evidence was trusted and how an output was verified is becoming a more credible signal of capability.

5 Governance, infrastructure and equity will determine the next phase of AI adoption

The next phase of AI adoption will depend on institutional capabilities. Data governance has emerged as a critical concern following the notification of the DPDP Rules, requiring institutions to reassess how learner data is collected, stored, processed and shared through AI-enabled systems.

At the same time, infrastructure constraints, bandwidth limitations, uneven device access and the recurring costs of enterprise AI licenses present significant implementation barriers, particularly for resource-constrained institutions.

Future success is likely to depend on developing scalable institutional models that combine infrastructure, governance, faculty development, industry engagement and equitable access. Industry-led Centers of Excellence, shared AI infrastructure and hub-and-spoke collaboration models offer promising pathways to scale capability beyond a small group of elite institutions.



The India campus story: Unfiltered, uncomfortable, undeniable

Chapter 2

AI is reshaping how students learn, blurring the lines between support, shortcut and substitution in everyday academic work

2.1 Students and AI: The ground reality no one wants to talk about

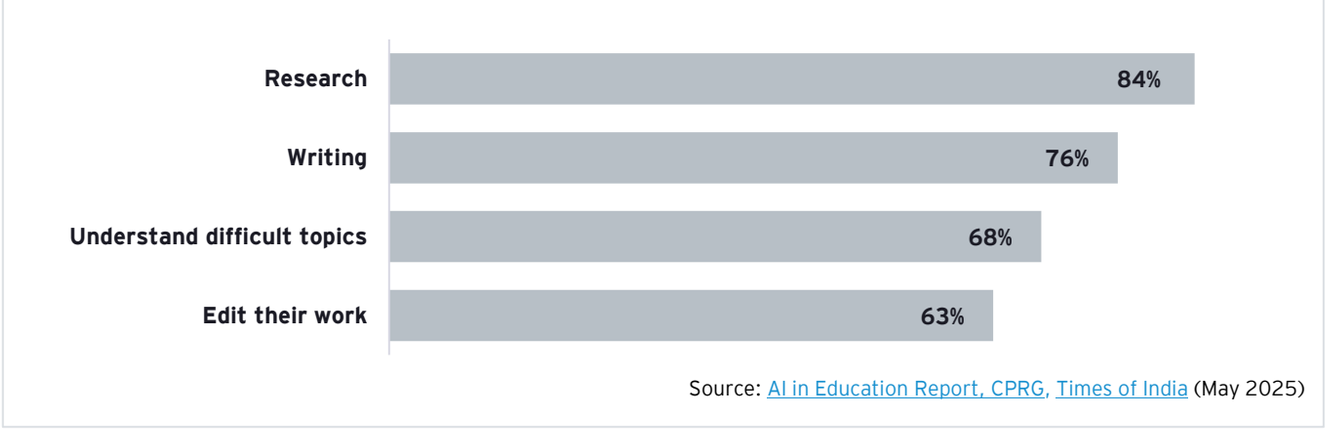
The growing use of artificial intelligence (AI) in higher education is changing how students study, learn and complete academic work. Most discussions tend to focus on two extremes: AI as a powerful productivity tool or as a means of engaging in unethical practices. While both perspectives have merit, neither fully captures how students are actually using AI. Much of this use falls within a gray area that many institutions, particularly in India, are still trying to understand and manage.

Responsible AI use requires students to combine independent thinking with the capabilities of AI tools. This, in turn, requires clear institutional rules, ethical guidelines and appropriate training to help students understand when and how AI can be used responsibly.

2.1.1 The spectrum of use: Mapping what students are actually doing

AI adoption among students in India is already widespread. Around 23% of students use AI every day, while 43% use it several times a week. Only 1% report never using it. Unlike the gradual transition from print-based resources to the internet, AI has quickly become a regular part of how students learn and complete academic tasks.[2]

Fig 1: How students use AI



Source: AI in Education Report, CPRG, Times of India (May 2025)

2. AI in Education Report, CPRG (May 2025), Times of India

Source: EY-Parthenon analysis

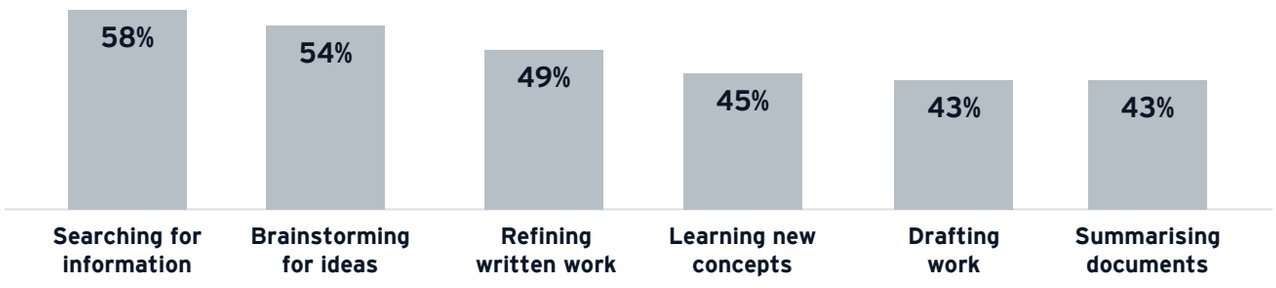
AI is becoming deeply embedded in everyday student learning, transforming how students research, write, practice and engage with academic work

Students often use AI to manage heavy academic workloads more efficiently. For common language tasks, AI helps them overcome initial writer's block, organize ideas and improve the quality of their writing. Many students draft their thoughts in simple language or bullet points and then use AI tools to turn them into well-structured, polished paragraphs in the required academic style. Some also rely on AI to draft professional emails to faculty, saving time and reducing effort.

Fig 2: How students use AI globally

Top 6 AI use cases by students

Question: What do you usually use AI for? (Select all that apply)



Source: Digital Education Council Global AI in Higher Education Global Survey 2026

Insights from interviews and focus group discussions suggest that student use of AI can be categorized into three broad bands.

Band 1: Productive use

Productive use involves treating AI as a support tool to structure thinking, refine original work, clarify complex concepts and navigate information more efficiently. In these cases, AI assists the learning process without replacing independent analysis or original effort.

Students are using AI productively in several ways:

- **Searching for information:** AI helps students quickly navigate large amounts of research that would take much longer to review manually.
- **Checking grammar:** AI tools are used to improve and refine writing, rather than to create it.
- **Summarizing and understanding:** Students use AI tools to simplify long documents and explain concepts that may not be clear in lectures.
- **Paraphrasing and structuring ideas:** Students use AI to reframe content and plan assignments while continuing to conduct the main analysis themselves.
- **Strategic sounding board:** Leveraging AI to brainstorm to expand the scope of thinking, surface different angles and stress test initial ideas.

> *"Generative AI has become an integral tool for exploring new ideas, enabling voice to text accessibility, validating thoughts and learning about unfamiliar topics. I use AI to refine and articulate my ideas more effectively, transforming abstract concepts into concrete outputs and breaking complex tasks into smaller, more manageable components. Additionally, I leverage AI to gain alternative perspectives and insights when approaching new or challenging assignments. Through consistent use of large language models (LLMs) over the past 10 months, the model has adapted to my preferred communication style and response format, resulting in more personalized and relevant interactions. As a result, my approach has evolved from a traditional "thinking-first" mindset to an "AI-first" approach, where AI serves as a key collaborator in ideation, problem-solving and decision-making. More broadly, I have observed that students utilize Generative AI in diverse ways, with some employing it to address open-ended and creative problems, while others use it to automate routine tasks and improve efficiency."* ***- BITS Student Focus Group Discussion***

Source: EY-Parthenon analysis

AI is bridging language gaps and academic uncertainty, while also reshaping how students access learning, career support and personalized guidance

AI as a language support tool

For many students who study in English as a second or third language, particularly those outside leading institutions in major cities, AI serves as an important language support tool. It helps students express their ideas more clearly in English without necessarily changing the depth or originality of their thinking. This use case is often overlooked and, in some instances, treated in the same way as academic misconduct. However, this may reflect a gap in how institutions understand AI-assisted learning rather than a problem with how students are using the technology.

In our focus group discussions (FGDs), it was revealed that students also use AI to write emails, particularly to help individuals whose first language is not English to communicate more effectively.

Source: EY-Parthenon analysis



Uncertainty around AI policies is leading to widespread but inconsistent student practices, blurring the line between support and misuse

Band 2: The gray zone

The gray zone is likely to be both the most consequential and potentially the largest category of student AI use. It includes practices that are neither clearly productive nor explicitly dishonest, but whose acceptability depends on the nature of the assignment, the extent of AI involvement and the expectations of the instructor.

Despite the growing and often thoughtful adoption of AI, many Indian universities have yet to establish clear institutional policies defining acceptable and unacceptable uses of the technology. As a result, both students and faculty operate amid considerable uncertainty. Some instructors strictly prohibit AI use, while others permit it informally or under specific circumstances. Many students are therefore left without a clear understanding of where legitimate assistance ends and academic misconduct begins.

This ambiguity represents a missed opportunity for institutions to guide students toward responsible and effective AI use.

Within this gray area, students commonly adopt mixed practices such as:

- Copying assignment questions into ChatGPT and rewriting the response in their own words.
- Using AI to develop an assignment structure before adding their own analysis and examples.
- Relying on AI to complete parts of assignments intended to assess their individual skills.
- Making decisions about AI use without clear guidance because institutional policies are absent, inconsistent or poorly suited to real academic situations.

Globally, 42% of students report that AI-generated content receives good grades.[3] This suggests that the challenge lies not only in student behavior but also in the design of assessments, which may not always distinguish between genuine understanding and polished AI-assisted output.

The ethical dilemma of AI-generated work

AI tools are increasingly becoming routine study collaborators. They can explain difficult topics, improve grammar, support writing and help students build confidence. When used appropriately, these tools can strengthen the learning process. Concerns arise, however, when students depend on AI to complete entire assignments or substitute for the intellectual work that the assignment is intended to assess.

Submitting fully AI-generated work limits students' opportunities to think critically, communicate clearly and solve problems independently. These capabilities are central not only to academic development but also to future employability.

While many students use AI to support their work, ethical concerns emerge when students:

- Submit AI generated essays as their own.
- Copy answers without understanding the underlying concepts.
- Use AI to avoid reading or conducting research.
- Bypass feedback and revision by relying on ready-made, polished outputs.

Some primary interviews revealed concerns about significant student dependence on AI. In a few cases, participants estimated that as many as 80% of students in a classroom rely heavily on AI tools, often without adequately reviewing or verifying the output. Faculty also reported instances where students treated AI-generated responses as authoritative, arguing that their answers were correct because they were produced by AI rather than engaging with feedback or examining potential flaws in the reasoning.

"AI has made me so lazy. I took an oath at the start of the semester not to use AI, but one week into the semester, I had already used it multiple times." ***- Student participant, Focus Group Discussion***

3. Student Generative AI Survey, Higher Education Policy Institute, 2026

Source: EY-Parthenon analysis

Highly visible cases of AI misuse are shaping institutional responses and public debate, but they capture only a narrow slice of a much broader and more complex reality

Band 3: Explicit misconduct

This category includes the clearest forms of AI misuse, in which established academic rules are directly violated. Although explicit misconduct often receives the greatest attention in public discourse, it represents only one part of the broader challenge posed by student AI use.

Direct submission of AI-generated work

One common form of misconduct involves submitting AI-generated content as original work without modification or disclosure. This may include:

- Using AI tools to generate complete essays, assignments or code and presenting the output as one's own.
- Skipping the research process and relying entirely on AI-generated answers.
- Using AI to answer questions during online examinations or other graded assessments.

"A student used Claude to complete a project that was later appreciated by the professor; however, the student acknowledged that they had not gained any meaningful learning from the process. Similarly, some students paste entire hackathon brochures into AI tools to generate answers, resulting in minimal engagement with the material and little to no learning." ***- BITS Student FGD***

Fabricated content and references

Students may also use AI to generate citations, data or sources that do not exist. Because AI tools can produce information that appears credible while being inaccurate or entirely fabricated, this form of misuse can be difficult to detect and particularly damaging to academic standards.

Use of AI during supervised assessments

Another significant concern is the use of AI during proctored examinations or quizzes where such tools are explicitly prohibited. Excessive reliance on AI can create a false sense of learning and make it difficult for instructors to assess what students genuinely understand.

Why it matters?

AI should support learning rather than replace it. It can be highly effective in generating ideas, improving writing and explaining difficult concepts. However, when students rely on AI to perform the entire task, they lose opportunities to develop essential academic and professional skills.

To date, many institutions have focused primarily on controlling explicit misuse through AI-detection tools and strict prohibitions. This approach, however, often overlooks the much larger gray zone in which a substantial share of student AI use appears to occur.

A policy gap

Existing academic-integrity frameworks are not fully equipped to address AI-assisted work. Traditional plagiarism rules, including University Grants Commission (UGC) guidelines, focus primarily on textual similarity and prescribe penalties based on the proportion of overlapping text. AI-generated content, however, may contain entirely original wording even when the ideas, structure and argument have been generated or substantially shaped by an AI system. This creates a significant conceptual gap. Students operating within these unclear rules are not necessarily attempting to cheat; many are navigating academic systems that have not yet adapted to the realities of AI-enabled learning.

Source: EY-Parthenon analysis

Outdated regulatory frameworks are struggling to keep pace with AI, creating ambiguity where existing rules fail to capture how AI-assisted work is actually produced

Illustrative example: AI integration in higher education

UGC

The UGC's existing regulations cap permissible textual similarity at 10%–20% for research submissions, with up to 10% similarity receiving no penalty, 10%–40% requiring revision and anything greater than 40% leading to escalated penalties.

These rules were designed for traditional forms of plagiarism and provide little conceptual clarity on how AI-assisted writing should be classified when the language is original but the ideas, structure, or argument have been generated by AI.

Students navigating this ambiguity are not being dishonest; they are navigating a regulatory framework that was built for a world that no longer exists.



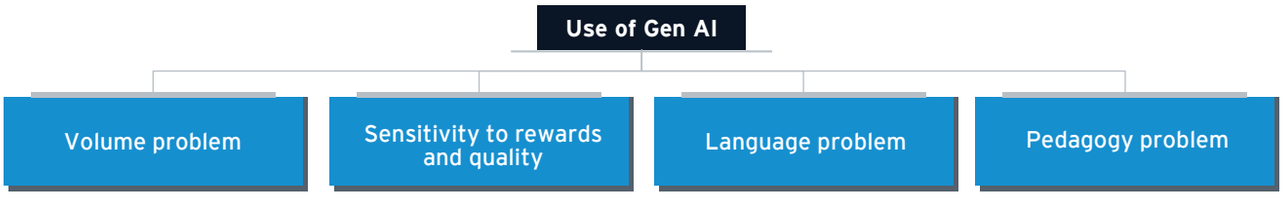
Sources: UGC, EY-Parthenon analysis

Student reliance on AI is shaped less by choice and more by system design, where heavy workloads and unclear learning goals make AI a rational response

2.1.2 Why students are reaching for AI: The structural logic

Student AI adoption is not primarily a moral phenomenon; it is a structural one. Within the current conditions of Indian higher education, using AI is often a rational and adaptive response to institutional shortcomings rather than evidence of a moral failure on the part of students.

Four structural drivers merit particular attention.



1 The volume problem

Colleges and professional programs often require students to complete a large volume of written work within compressed timelines. However, the purpose of this workload is not always clearly linked to specific learning objectives. To manage the resulting pressure, students turn to AI tools to accelerate tasks such as research, drafting and editing.

In such situations, AI does not simply provide a shortcut around learning. It becomes a way of navigating a system that may not be supporting meaningful learning in the first place.

When assignments are not designed intentionally, a high volume of coursework may shift the emphasis from learning to task completion. This may reveal deeper weaknesses in course design, including:

- **Pedagogical padding:** Syllabi are filled with repetitive tasks rather than fewer, more substantive projects that encourage critical thinking and deeper engagement.
- **Exhaustion instead of learning:** Continuous deadlines increase stress and burnout, potentially reducing how much students understand and retain.
- **Submission over understanding:** When every stage of an assignment is graded, students may prioritize completing the task over engaging meaningfully with the material, making AI an attractive means of accelerating submission.

In our primary interviews, we found that the students are using AI to cope with the increasing volume of work. This appears to be cyclical, as students have adapted to the availability of AI tools, faculty are proportionately assigning larger volumes of work because AI has accelerated the speed at which students can complete and submit assignments.

Source: EY-Parthenon analysis

Performance pressure, quality expectations and language gaps are driving students to rely on AI as both a support tool and a coping mechanism

2 Sensitivity to rewards and quality

Another important factor is how much students care about grades. Students who are highly focused on achieving good grades may be more likely to use AI tools because they see them as a way to improve performance. In some cases, this strong focus on rewards may lead students to misuse AI tools, such as relying on them too heavily or using them in ways that may cross academic integrity boundaries.[4]

> *"There is a pressure to use AI. If nine out of ten groups are using it, then you feel compelled to use it too, even if you wouldn't otherwise."* ***- Student Participant, Focus Group Discussion***

3 The language problem

India's higher education system enrolls students from highly diverse linguistic backgrounds. These include first generation English medium learners, students educated entirely in regional language schools and students for whom English is a tool of academic and professional navigation rather than their primary language of thought.

For many of these students, AI removes a barrier that the education system itself has created. English proficiency is often treated as a prerequisite for academic success, without adequate support to help students develop it.

Students are increasingly using AI-enabled speech translation, local language chatbots and dynamic transcript generation tools to bridge the gap between English medium instruction and vernacular learning. These tools allow students to understand complex lectures and study materials in their mother tongues, enabling language to support rather than constrain learning.

The growing use of AI in vernacular education is beginning to reshape classrooms across the country:

- **Speech to speech translation:** In October 2025, IIT Madras used AI to translate live English language classroom lectures into 11 regional languages in real time.[5]
- **Multilingual government platforms:** Bhashini, the Government of India's AI-enabled language platform, supports the translation of digital educational content and powers voice-based chatbots that can respond to academic questions in local languages.[6]

Some primary interviews suggested that students, particularly those from Tier 2 and Tier 3 universities, are able to spend more time on the thinking and problem-solving aspects of their work by using AI to generate polished English language outputs. Participants also noted that AI can help simplify complex concepts and make assignment requirements easier to understand, allowing students to focus more on the substance of their work.

4. Springer,
5. IIT Madras,
6. Bhashini

Source: EY-Parthenon analysis

The deeper challenge lies in pedagogy, where outdated teaching methods, unclear purpose, are pushing students toward AI as a substitute for meaningful learning

4 The pedagogy problem: The most fundamental and least discussed

Across much of Indian higher education, teaching continues to rely on traditional, lecture-based methods that emphasize memorization. However, contemporary academic and professional environments increasingly require practical, analytical and problem-solving capabilities.

This mismatch contributes to passive learning, weak student engagement and assessments that often measure recall or compliance rather than genuine understanding. Only a limited number of universities have meaningfully redesigned their pedagogy and assessment systems to reflect these changing requirements.
This broader concern is reinforced by MIT's 2026 Ad Hoc Committee on AI Use in Teaching, Learning and Research Training, which raises a more fundamental question about what universities understand learning to be. If learning is treated primarily as the delivery and reproduction of content, increasingly capable AI tools can perform much of that process. The distinctive role of higher education instead lies in developing capabilities that require active engagement: reasoning through unfamiliar problems, exercising judgment, testing ideas, receiving feedback and learning through interaction with others. AI therefore makes the shift from content transmission to capability development more important, not less.

- **AI has outgrown traditional assignments:** Many assessment formats, particularly open book essays and take-home reports, were designed in a pre-AI era when independently produced written work was a reliable indicator of individual understanding. Generative AI has weakened that assumption. These assignments are now inherently compatible with AI, not necessarily because students intend to misuse technology, but because the tasks themselves no longer require distinctly human cognitive effort.
- **Learning conditions have changed:** The effort and reasoning once required to complete such assignments can now be partly or entirely outsourced to AI. Although the assignment formats have remained largely unchanged, the environment in which students complete them has been fundamentally transformed.
- **This is a design problem, not only an ethical one:** The challenge lies primarily in the design of courses and assessments. Detection software cannot resolve a fundamental mismatch between existing assessment formats and current technological realities. When a task can be completed effectively through externalized reasoning, responsibility cannot rest solely with the student. The solution lies in redesigning assignments to create conditions for genuine thinking, rather than relying primarily on tools intended to detect or police AI use.

"Class participation is going down because students know AI can explain concepts easily. We are also seeing fewer thoughtful questions being asked in class." ***– Faculty interview***

"AI has made it easier to produce polished content, but also a lot more fluff. Whether it is papers or emails, the challenge is increasingly filtering through the volume and getting to the actual substance." ***- Faculty Interview***

Source: EY-Parthenon analysis

A parallel AI-driven learning ecosystem is emerging, operating alongside formal education and reshaping how students access support, guidance and peer learning

2.1.3 The shadow education system

India has long operated through a dual education ecosystem, a formal system of institutions and curricula, alongside a parallel and largely informal network of coaching centers, private tutors, peer-learning groups and examination-oriented communities.

Student reliance on AI is contributing to the emergence of a new, technology enabled layer within this parallel ecosystem, an AI driven **"shadow education system."** This comprises a rapidly expanding range of AI-powered tools, chatbots and personalized learning platforms that students use independently, often beyond formal institutional oversight. Together, these tools serve as an alternative support system for homework, essay drafting, examination preparation, concept clarification and, in some cases, emotional support.

In practice, this shadow system is evolving in several ways:

- Students describe AI as "the only teacher who explains things until I actually understand." The statement captures both AI's value as an endlessly available tutor and the limitations of formal pedagogy that make such support necessary.

> *"I turn to AI because it can break down complex concepts into simple terms. When I go to a professor, there's often an assumption of some baseline knowledge, whereas with AI, I can start from zero and build my understanding from the ground up."* ***- Student participant, Focus Group Discussion***

- Platforms such as Amity University Online have begun integrating virtual faculty (e.g., "AMI") into their learning management system (LMS), emphasizing 24/7 support. While framed as innovation, this also implicitly acknowledges the limitations of human faculty in providing continuous, on-demand academic assistance.[7]

> *"If my faculty member is available, I would prefer to ask them questions directly, as they can provide deeper insights and context. However, there are many times when faculty members are not accessible, such as when we are studying late at night or working from our dorms. In those situations, we often rely on AI for guidance and support."* ***– VIT student FGD***

- Peer-led YouTube channels are increasingly teaching students how to use AI tools effectively for discipline-specific tasks, particularly in engineering and management education.
- Discord communities have formed around academic specializations, where members collaboratively generate, critique and refine AI-assisted outputs.[8]

The resulting stratification problem

The benefits of this emerging system are not distributed evenly:

- Students who participate in peer-led AI communities may develop meaningful AI literacy through shared experimentation, feedback and iterative learning.
- Students who are socially or economically excluded from these networks often use the same tools in isolation, producing outputs that they may lack the knowledge or confidence to assess critically.

> *"My understanding of AI use cases is often shaped by my environment. If my peers start using a new tool or discover a new approach, I'm more likely to explore and adopt it myself."* ***- Student participant, Focus Group Discussion***

7. Amity University Online ,
8. Research Gate

Source: EY-Parthenon analysis

AI is not reducing inequality in higher education but amplifying it by rewarding those with existing advantages while leaving others further behind

2.1.4 The amplification problem: Who thrives, who does not

The risk is that AI may not function as an equalizer in Indian higher education. It may function as an amplifier and in a system already characterized by stark structural inequalities, amplification accelerates divergence rather than reducing it.

AI improves the student experience for many – but not all. Almost half (49%) of students believe AI has improved their student experience, particularly by saving time, improving understanding and providing instant support. However, a minority feel AI has worsened their experience, citing concerns about fairness, skill erosion, social isolation and future employment.[9]

Students who benefit most from AI typically share a cluster of enabling characteristics:

- Sufficient foundational knowledge to assess whether AI outputs are accurate, relevant and well reasoned (most cited characteristic in primary interviews).
- English-language proficiency that enables them to prompt effectively and edit outputs critically.
- Access to laptops or desktop computers rather than mobile-only access.
- Exposure to peer or institutional networks in which AI use is discussed critically, rather than used only as a means of completing tasks.

Institutional environments, typically within premier institutions or well-resourced private universities, that reward original thinking clearly enough to motivate effort beyond AI-assisted completion.

Students who may be less well positioned to benefit from AI often face the opposite conditions:

- Weak subject foundations that make it difficult to assess the quality or accuracy of AI outputs.
- Limited exposure to critical discussions of AI's limitations and failure modes.
- Mobile-only internet access, which restricts the range of tools and features they can use effectively.
- Social or economic circumstances that make AI their primary, or sometimes only, source of academic support.

While concerns about unequal access were raised in some interviews, most stakeholders suggested that these disparities are most pronounced when students first enter university. Over time, however, universities can act as a leveling field by providing students with similar learning environments, institutional resources and peer support, thereby reducing differences in access and capability.

However, when students use AI to mask gaps in basic understanding over the course of a three-year degree, they may graduate with qualifications that do not reflect their true abilities. They then enter a labor market, discussed further in Section 3, that increasingly values applied thinking, independent judgment and the ability to perform without constant AI assistance.

The students who can least afford the consequences of AI-dependent learning may therefore be those most likely to rely on it heavily. This reinforces the need for data-driven analysis that distinguishes between students at Tier 1, Tier 2, Tier 3 and state-affiliated institutions. National averages may otherwise obscure the widening gaps between these groups.

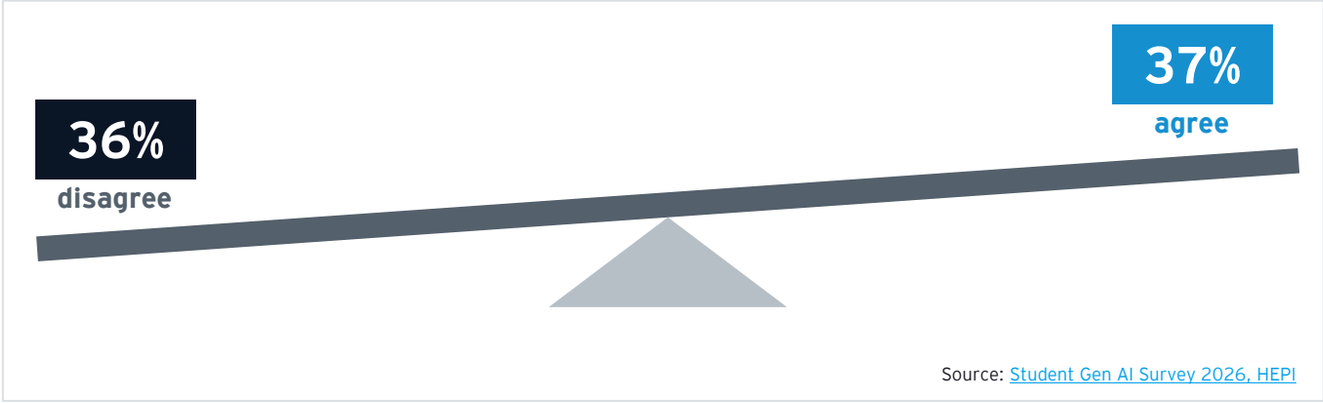
9. Student Generative AI Survey, Higher Education Policy Institute, 2026

Source: EY-Parthenon analysis

Students recognize AI skills as essential but inconsistent institutional support is creating gaps in how effectively these skills are developed

Students are almost evenly split on whether their institutions encourage AI use, with 37% agreeing and 36% disagreeing. While students clearly see AI skills as essential, with 68% believing they are necessary to succeed in today's world, institutional support has not kept pace. Fewer than half (48%) feel that teaching staff are helping them develop these skills for their future careers, with arts and humanities students especially likely to feel under supported.[10]

Fig 3: Do universities encourage AI use? (Student response)



Source: Student Gen AI Survey 2026, HEPI

In our primary interviews, we found that teachers are generally comfortable with students using AI as part of their coursework. However, students are expected to be transparent about how AI has been used, including its contribution to their work, the process followed and how they arrived at their final responses. Students should be able to explain and justify their reasoning independently, rather than relying solely on AI-generated outputs.

2.1.5 The mental health dimension: What institutions are not seeing

AI's relationship with student wellbeing is complex and often contradictory. It can contribute to isolation and dependency, while also providing companionship, advice and emotional support. Around 40% of students report that AI affects their experience of loneliness, with positive and negative effects reported in broadly equal measure. Approximately 15% of students use AI for companionship, advice or support in addressing loneliness.[11]

Recent research by scholars at Wharton highlights a phenomenon described as "cognitive surrender," in which individuals gradually defer their own critical reasoning to AI systems. In controlled experiments involving 1,372 participants, around 80% accepted AI-generated answers without independently verifying them, even when several of the answers were deliberately incorrect. This suggests a growing reliance on AI's perceived authority and a weakening of the cognitive instincts that ordinarily help protect against errors.[12]

10, 11. Student Generative AI Survey, Higher Education Policy Institute, 2026

12. Thinking—Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender, Wharton School Research Paper, 2026

Source: EY-Parthenon analysis

AI is both easing and intensifying student stress, creating new patterns of dependence, burnout and confidence erosion

The effects of AI on student wellbeing are not uniformly negative. Research published in the *International Journal of Indian Psychology* found that young Indian adults who used AI recorded lower average levels of stress and burnout than non-users.[13] AI may reduce stress by lowering the immediate cognitive demands of academic work, although this does not necessarily strengthen students' underlying capacity to manage those demands independently.[13]

A separate study found that the perceived usefulness of AI tools can encourage dependence. Through mechanisms associated with addictive use, this dependence may contribute to higher levels of learning burnout and reduced engagement.

> *"AI does not inherently create or reduce stress; its impact largely depends on how it is used. If prompts are unclear or ineffective, AI can be frustrating to use. However, when users understand how to interact with AI effectively, it can significantly reduce stress by improving productivity and simplifying tasks. The effects of AI also depend on the extent and manner of its use. While responsible use can enhance efficiency and learning, excessive reliance on AI may increase stress and reduce confidence in performing tasks independently."* ***- Deakin University GIFT City Campus, Student FGD***

AI-dependency anxiety

Some students have outsourced so much cognitive work to AI that they struggle to begin writing, analyzing or reasoning without consulting it. When access to these tools is limited, restricted or unreliable, the result may be genuine distress. This pattern is becoming increasingly visible in institutional counseling data where staff are equipped to ask students about their AI use.

In these cases, academic stress and declining self-confidence both contribute to and result from AI dependency, creating a self-reinforcing cycle that can be difficult to break.

Erosion of confidence

As students progressively outsource intellectual work to AI, they may begin to lose confidence in their own judgment. Some struggle to evaluate their writing without AI-generated feedback and experience acute difficulty in assessments where AI use is restricted. Heavy reliance on AI can therefore leave students less prepared for tasks that require independent reasoning than they might have been had they not used the tools at all. This occurs when AI replaces underlying capability rather than helping to develop it.

Burnout from volume without depth

AI enables students to produce work at a pace to which institutions have not yet adapted. This is creating a growing group of students who satisfy submission requirements on paper but experience the exhaustion of continuous production without corresponding intellectual engagement, understanding or learning.

13. Gurdatta, Gururani & Gupta, 2025

Source: EY-Parthenon analysis

Faculty capacity and choices are shaping how AI is used in classrooms, determining whether it supports real learning or reinforces surface level performance

2.2 Faculty reality check: Who's ready, who's resistant and why

Faculty play a pivotal role in shaping whether AI becomes a genuine learning aid or merely a tool for superficial performance. As the designers of assessments, they effectively set the terms of engagement deciding whether tasks encourage critical thinking, originality and application, or can be easily outsourced to AI. Beyond assessment design, faculty also shapes classroom culture, they influence whether AI is approached as a tool to be thoughtfully interrogated and integrated into learning, or as a threat that should be closely monitored and controlled. In doing so, they determine not just how students use AI, but how they learn to think alongside it.

2.2.1 A typology that resists caricature: The innovators, the pretenders, the resistors

The three-way typology proposed here is a diagnostic instrument, not a hierarchy of virtue. Its purpose is to name the full range of faculty responses accurately, so that institutional support can be designed for actual conditions rather than idealized ones.

Characteristics of the innovators, the pretenders and the resistors

Innovators	▪ Invest personal time, often ahead of any institutional mandate, in understanding how generative AI works and what it means for teaching and learning. ▪ Redesign assessments to require demonstrations of understanding that AI cannot easily simulate. These may include oral examinations, iterative drafts accompanied by documented reasoning and project-based evaluations that apply concepts to unfamiliar or context-specific problems. ▪ Engage with AI as a pedagogical tool rather than simply as a means of improving administrative efficiency or generating teaching materials.
Pretenders	▪ Attend institutionally organized workshops, receive certificates of completion, but make little substantive change to how they teach or assess students. ▪ Appear comfortable with AI at an operational level without understanding how it changes student behavior, assessment validity or the learning process.
Resistors	▪ Express concern about what students lose when they outsource foundational tasks such as drafting, synthesis, analysis and problem-solving. ▪ Question the reliability of AI in specialist disciplines, where hallucinations or subtle errors may be difficult to identify without deep subject knowledge. ▪ Worry about the long-term effects of habitual AI use on students' capacity for independent thought and judgment.

Primary interviews revealed examples across all three categories. Some faculty resist AI because they are concerned about the loss of essential intellectual capabilities when students outsource core tasks such as drafting and synthesizing ideas. Others remain skeptical of AI's reliability in specialized fields, where errors may appear plausible and remain undetected without significant subject knowledge.

Such resistance should not automatically be dismissed as conservatism. Productive resistance is accompanied by alternative teaching and assessment methods that still require students to demonstrate genuine understanding. The distinction is therefore not between faculty who accept AI and those who reject it, but between those who have thoughtfully adapted their pedagogy and those whose response, whether enthusiastic or resistant, remains superficial.

Source: EY-Parthenon analysis



Flagship AI partnerships are advancing structured adoption, but their impact remains limited to elite institutions

Illustrative example: OpenAI partnership with IIM Ahmedabad

The OpenAI partnership with IIM Ahmedabad and other institutions, announced around the AI Impact Summit in February 2026, promising structured AI exposure for 1,00,000 students alongside faculty training, is one of the first systematic attempts to institutionalize this at scale, but remains confined to the top of the institutional pyramid.

Structure of the institutional push

- **Enterprise access:** Institutions gain secure, enterprise-grade access to ChatGPT Edu, shifting AI adoption beyond mere productivity hacks into core curriculum integration.

Discipline specific use

- **IIM Ahmedabad:** Integrates AI fluency into management, strategy, finance and public policy, alongside offering structured OpenAI certifications.
- **Faculty training:** Structured onboarding and responsible-use frameworks are being deployed to train educators to adapt to rapid skill shifts.

Expanding beyond the apex

Acknowledging that elite institutions represent only the top of the pyramid, OpenAI and the broader Indian education ecosystem are taking steps to democratize AI training:

- **Edtech integrations:** Collaborations with platforms such as upGrad, to extend structured AI courses to Tier-2 or Tier-3 students and early-career professionals.
- **Multidisciplinary scaling:** Broader university networks such as Manipal Academic of Higher Education (MAHE) and UPES are also receiving structured capability tracks, helping push AI literacy into health sciences, design and regional colleges.

Sources: IIM Ahmedabad, Open AI, Business Wire, EY-Parthenon analysis

Faculty hesitation toward AI reflects deeper structural and psychological barriers, not just simple resistance to new technology

2.2.2 Why many faculty still avoid AI: Fear, pride, fatigue and lack of time

Faculty avoidance of AI is not simply resistance to technology; it reflects a deeper mix of professional, institutional and psychological barriers.

I have not integrated AI into my classroom teaching. The classroom is where students build strong fundamentals and develop deep subject knowledge. Those foundations ultimately enable them to use AI more effectively in the workplace."
- Dr. Alok Bhardwaj, Assistant Professor, IIT Roorkee

Reasons faculty avoid AI

Professional Barriers

Fear of academic integrity collapse

- Faculty worry that students will use AI to avoid reading, writing, coding, analysis, and independent thinking.
- They may also fear a loss of classroom authority as AI changes the traditional relationship between teachers and students.
- Students can arrive with alternative explanations, more sophisticated questions, or AI generated work that faculty may not know how to evaluate confidently.

Students are increasingly using AI tools to generate presentations from financial reports. However, excessive reliance on AI can limit their understanding of the underlying concepts, making it difficult to answer basic questions and potentially affecting their performance in interviews and placements. AI should serve as a facilitator that enhances learning and productivity for both students and faculty. ***- Dr Anju Vaswani, Vivekananda Business School, Mumbai***

Institutional Barriers

Lack of time for real experimentation

- Meaningful AI integration requires faculty to test tools, redesign assignments, develop clear rubrics, and determine how AI should be used within the context of a specific discipline. Most faculty are not given protected time to undertake this work. Without such support, AI adoption is likely to remain limited to low-effort applications that fit within existing teaching practices rather than prompting deeper pedagogical change.

Incentives mismatch

- Promotions and recognition still reward publications, seniority, administrative work, and compliance, not teaching innovation
- A faculty member who spends a full semester redesigning a course around AI is often doing work that brings little professional reward

Psychological Barriers

Professional pride and discomfort

Many faculty members built their reputation on deep subject knowledge and strong teaching. Admitting uncertainty about AI can feel risky and uncomfortable.

Not all resistance stems from reluctance to change. Hallucinations, superficial learning, student dependence, data privacy risks, and weakening core skills are legitimate concerns that warrant attention.

Source: EY-Parthenon analysis

Surface level AI adoption by faculty reflects structural barriers and no genuine integration by higher education institutions

2.2.3 The silent crisis: Faculty who pretend they understand AI but do not

A more difficult challenge for universities is not overt resistance to AI, but superficial adoption, faculty who appear AI ready yet are not equipped to meaningfully guide students. While many instructors demonstrate moderate to high familiarity with AI tools, actual use remains concentrated in low barrier, general purpose applications such as chatbots and translation services.

Policy familiarity without pedagogical clarity

Many faculty are familiar with institutional AI policies and can articulate broad principles around responsible use. However, translating these principles into classroom practice remains a challenge. Interviews suggest uncertainty around what students should be permitted to do within specific assignments, where the line should be drawn between support and overreliance and how AI use should be disclosed or evaluated. Their own use of AI also tends to be concentrated in general-purpose tools such as chatbots and translation platforms, while pedagogy-specific applications, including intelligent tutoring systems and adaptive or personalized learning platforms, see far more limited adoption. As a result, faculty may understand the policy framework but still lack clarity on how to apply it consistently in teaching, assessment and student guidance.[14]

Tool exposure mistaken for AI literacy

Attending a workshop, generating lecture notes, or experimenting with ChatGPT does not amount to a meaningful understanding of AI's implications for education. Operational familiarity does not necessarily equip faculty to assess how AI changes learning behavior, weakens or strengthens assessment validity, or affects students' judgment and intellectual development.

Ambiguous classroom guidance

Students are told "do not misuse AI" or "do not copy from ChatGPT," but are not given clear examples of acceptable, unacceptable and gray zone use.

Assessment designs remain unchanged

Faculty may acknowledge the presence of AI, yet continue assigning essays, reports and take home tasks that can be easily generated by AI. Although faculty recognize the limitations of current assessment and evaluation models, there is still no clear consensus on what should replace them.[15]

Students left unguided

In this environment, students are left without clear direction neither fully empowered to use AI responsibly nor provided with boundaries that prevent misuse. The result is an expanding gray zone of AI use, shaped not by intentional learning design but by ambiguity, uneven guidance and a lack of pedagogical coherence.

Why this matters

Superficial adoption may be more damaging than transparent non adoption because it creates the appearance of progress while leaving fundamental pedagogical weaknesses unresolved. Faculty may appear aligned with AI because they cite policies, attend training programs, or experiment with basic tools, yet still lack the capacity to provide students with practical guidance or redesign assessments appropriately.

---

14. Faculty Perceptions and Adoption of AI in Higher Education: Insights from Two Lebanese Universities ,
15. Digital Education Council's AI in Higher Education Latin America (LATAM) Survey 2026

Source: EY-Parthenon analysis

AI is reshaping disciplines in uneven ways, strengthening scientific methods while challenging core ideas in the humanities and social sciences

2.2.4 Where AI integration is failing: A disciplinary map

Universities are beginning to examine how AI may transform research, reshape disciplines and alter teaching and learning. Its impact is most visible in science and engineering, but it is also raising fundamental questions in the humanities and social sciences about interpretation, authorship, creativity and what it means to understand something.

AI does not affect all disciplines in the same way. In science and engineering, it typically strengthens established methods by improving measurement, analysis, modeling and prediction. In the humanities and social sciences, however, AI can generate text, interpret patterns and reproduce forms of argument, challenging conventional assumptions about meaning, originality and human knowledge.

As AI strengthens some disciplines while disrupting the foundations of others, universities should avoid adopting a uniform approach to integration. They need to build discipline-specific capabilities, invest in appropriate infrastructure and encourage collaboration across departments. The central challenge is to distribute AI's benefits across fields rather than allowing existing disciplinary gaps to widen.

Adoption and implementation are likely to vary across courses. Some of these differences may narrow as familiarity and experience increase, but a clear transition gap remains. It is particularly visible among faculty and students from non-technical backgrounds, where limited exposure to AI tools may constrain confidence, experimentation and meaningful integration into teaching and learning.

Alignment in the sciences

AI aligns relatively well with science and engineering because it builds on methods already central to these disciplines, including mathematical modeling, simulation, computation and data analysis. Rather than fundamentally changing how these fields define knowledge, AI often acts as an advanced support tool that strengthens existing research processes.

This is visible across several disciplines. In bioinformatics, AI supports the analysis of genetic data. In climate science, it improves simulations and forecasting. In chemistry, it helps identify potential drug candidates, while in materials science it predicts how different materials may behave. In astronomy, AI is used to process telescope data and in medicine it helps identify patterns and detect potential health conditions.

Its immediate value lies largely in automating repetitive work, processing large datasets and accelerating experimentation. This allows researchers to test more hypotheses, run more simulations and explore a broader range of possibilities. In emerging fields such as quantum computing and autonomous systems, AI can improve performance and identify solutions more efficiently, even where the underlying research questions remain unchanged.

Challenges in the humanities and social sciences

The humanities and social sciences face a different set of challenges. These disciplines rely heavily on interpretation and engage with texts, history, culture, institutions and human behavior. Although AI can identify patterns and process large bodies of material, its growing use raises questions about whether it can genuinely understand meaning, context, ambiguity and cultural nuance.

Yet the relevance of AI extends beyond disciplines traditionally associated with technology. In fields such as music, performing arts and fine arts, AI is emerging as a tool that can support creative and analytical processes, from identifying ragas and classifying musical patterns to assisting choreography and documenting artistic knowledge systems. These applications illustrate how AI can augment, rather than replace, human creativity and capability. Researchers are also using AI to analyze large collections of texts, explore historical archives, study misinformation and identify patterns in social behavior. AI is not replacing these disciplines, but it is beginning to change the scale, speed and methods through which research is conducted.

Limitations of AI adoption across Indian higher education institutions follow identifiable patterns shaped by disciplinary fit, faculty capability, infrastructure and the extent to which existing teaching and research methods can accommodate AI. The gaps are structural rather than random.

In our primary research, we identified that STEM disciplines appear to have a significantly higher level of AI integration compared to non-STEM departments. It is noteworthy that all the universities we engaged with directed us to their STEM departments when arranging interviews, further reflecting the stronger association of AI initiatives with STEM related programs and faculties.

Source: EY-Parthenon analysis

Strong AI ambitions are being undermined by weak execution, as regulatory ambiguity, procurement constraints and leadership gaps prevent intent from translating into action

2.3 Administration and leadership: Intent vs. action gap

Institutional leaders across Indian higher education have made increasingly visible commitments to AI. References to AI now appear across institutions of every type and tier, including in Vice-Chancellors' convocation addresses, AI readiness statements on university websites and press releases announcing new committees and task forces.

However, the operational changes required to translate these declarations into practice have not kept pace. Procurement systems remain poorly suited to emerging technologies, data governance frameworks are often incomplete, assessment redesign is limited, faculty development efforts are rarely sustained and few institutions have established governance structures capable of making timely and informed decisions.

2.3.1 The announcement trap: Why intent does not become action

Although AI governance has gained significant policy attention across higher education, evidence from stakeholder interviews suggests that implementation and enforcement remain at varying stages of maturity across institutions.

Three conditions repeatedly turn genuine institutional intent into delayed or incomplete implementation:

Regulatory opacity

- Institutions are being asked to make consequential decisions about AI governance in an environment that offers limited practical clarity on what responsible adoption requires.
- AICTE's declaration of 2025 as the "Year of AI," covering 14,000 colleges and 40 million students, was an important statement of ambition. However, it was not accompanied by a detailed compliance framework, implementation guidance or corresponding resource allocation.[16]
- Institutional leaders waiting for regulatory clarity are therefore waiting for guidance that has not yet arrived and may not emerge in the form they expect. At the same time, leaders choosing to act are being forced to make decisions within a genuine governance vacuum.

Procurement constraints

Government aided and state universities typically operate under procurement systems designed for conventional infrastructure acquisition. Requirements such as formal tendering, comparative bidding and multiple layers of administrative approval are often incompatible with the speed at which AI tools evolve.

For example, an institution with a six-month procurement cycle, may find it difficult to make meaningful decisions about AI tools. By the time a contract is finalized, the capabilities, pricing and competitive position of the selected product may have changed substantially.

This should not be understood simply as institutional inertia. It reflects a structural mismatch between procurement design and technological reality. The effects of this mismatch are also uneven; it disproportionately constrains the institutions where equitable and thoughtful AI integration may be most urgently required.

The leadership capability gap

- The capability gap within institutional leadership deserves direct examination rather than diplomatic silence.
- Vice Chancellors and Registrars, who are generally appointed for their academic or administrative experience, are now also expected to evaluate AI vendor proposals, interpret data-processing agreements under the DPDP Rules, 2025 and determine which AI tools are appropriate for their students.
- Institutions either adopt the tools promoted most aggressively by vendors without sufficient scrutiny or delay decisions indefinitely in the hope that clearer guidance will eventually appear.

"AI ownership and governance are critical. Whether the solution is managed by IT, a third party or a business function, clear accountability must be established. Organizations need a defined control owner who is responsible for oversight, risk management and responding to auditors. If outsourced, service providers must assume accountability through robust governance and contractual controls." ***- Global AI and innovation leader, education business unit, global IT services company***

16. AICTE AI Integration Initiative

Source: EY-Parthenon analysis

Critical data governance risks under the DPDP framework are being underestimated, leaving institutions exposed to legal, operational and student rights failures

2.3.2 The DPDP problem no one is taking seriously enough

The notification of the DPDP Rules 2025 in November 2025 represents one of the most significant regulatory developments affecting AI governance in Indian higher education institutions over the past two years. However, many institutions do not yet appear to be treating it as a central component of their AI adoption strategies.[17]

What the DPDP rules actually require of HEIs:

Educational institutions today carry a far greater responsibility than simply adopting AI tools for learning, they operate as data fiduciaries, accountable for every aspect of the student data they process. This includes behavioral data, performance metrics, interaction logs and the continuous streams of information generated whenever AI-powered tutoring systems, proctoring platforms or generative AI tools are used within an institutional ecosystem. Once an institution deploys such a tool under its license, the legal and ethical responsibility for that data shifts squarely onto the institution. It must confirm that data is collected only with meaningful consent, used strictly for the purposes that were communicated, safeguarded through robust security measures and never transferred to jurisdictions that fall outside approved regulatory frameworks.[18]

This responsibility becomes even more complex because many AI learning tools widely used across campuses are global products whose data flows often extend across borders. Institutions that adopt them cannot simply rely on vendor assurances, They must actively evaluate whether the data processing involved complies with the cross-border transfer provisions outlined in the DPDP Rules. In practice, this means understanding where student data is stored, how it is processed, whether it is shared with third-party providers and whether any part of this chain exposes students to privacy risks or regulatory non-compliance. As AI becomes more deeply embedded in learning environments, institutional stewardship of data is no longer merely a compliance exercise but a core component of safeguarding student rights and maintaining trust in educational systems.

The practical exposure institutions are not accounting for:

Organizations in India now face serious legal consequences for data breaches, with penalties reaching up to INR250 crore per violation under the DPDP Rules, 2025. Any institution that fails to report a breach within the mandatory 72-hour window or submits delayed or incomplete reports risks additional fines, compounding the financial and reputational damage. With full enforcement set to begin on 13 May 2027, institutions should start preparing today, though many have not yet even begun the simplest and most critical step: data mapping.[19, 20]

This regulatory landscape also introduces a powerful student rights imperative. Students are entitled to know what data is being collected about them, the reasons for its collection and the safeguards in place to protect it. As enforcement accelerates under the DPDP regime, institutions that fail to openly communicate about data practices and protections are not only risking non-compliance but are also perpetuating a breach of student rights and institutional accountability.

BITS Goa: *Whenever AI tools are deployed within a department, data privacy and security are treated as critical priorities. At BITS Goa, a dedicated IT manager oversees data governance, privacy and AI related compliance requirements. Any use of AI tools involving institutional data requires appropriate approvals from the IT team and relevant departments. Strict protocols are in place to ensure that sensitive or confidential information is not uploaded to external AI platforms. Additionally, non disclosure agreements (NDAs) are maintained with technology vendors to safeguard institutional data and ensure responsible data handling practices.*

17. PIB
18 EdTech Magazine
19. PIB
20. EY
Source: EY-Parthenon analysis

Institutional AI ambitions are outpacing infrastructure and financial readiness, exposing gaps in access, bandwidth and sustainable investment

2.3.3 The infrastructure reality beneath the announcements

If we move beyond the language of "digital transformation" and examine the actual conditions on the ground in most Indian higher education institutions, a clear gap between aspiration and infrastructure quickly becomes visible.

Bandwidth

AI tools such as generative AI platforms, AI-enabled proctoring systems and adaptive learning technologies, require high and stable internet connectivity. These tools are far more network intensive than the systems most institutions were originally designed to support. Existing IT infrastructure in many campuses was built primarily for administrative needs, such as student records, payroll, email and basic learning management systems. Scaling this infrastructure to support thousands of students using AI tools at the same time requires significant financial investment.

Device access

Discussions around AI integration often assume that students have access to personal laptops or desktops. However, this assumption may not hold across the full diversity of Indian higher education and is more likely to reflect the realities of certain segments than the sector as a whole. A large number of students, especially in state universities and affiliated colleges, rely mainly on mobile phones for internet access. While many AI tools are becoming mobile-friendly, their full capabilities are still best experienced on laptops or desktops.[21]

AI integration strategies built around students with reliable broadband and personal computers risk excluding those who need these tools the most. Without addressing these gaps in access, efforts to use AI for improving equity in education may end up reinforcing existing inequalities instead of reducing them.

2.3.4 Hidden and unbudgeted costs

Enterprise AI tool licenses introduce a fundamentally different kind of financial commitment for higher education institutions. These tools are typically priced on a per user or per query basis, meaning costs increase directly with usage. As adoption scales across an institution, these expenses can grow rapidly and become difficult to predict. In addition, as the AI market matures and consolidates, providers are likely to increase pricing power, making long-term cost trajectories even more uncertain. Unlike traditional capital investments such as buildings, laboratories or equipment, AI tools represent a recurring operational expense. This is a category of spending that most higher education financial systems have not historically planned for, either in budgeting structures or in long-term financial models.[22]

At the same time, meaningful integration of AI into teaching requires a parallel investment in faculty development one that goes far beyond short term workshops or introductory sessions. Effective adoption depends on sustained, discipline specific training that enables faculty to redesign assessments, rethink pedagogy and integrate AI into learning in ways that promote genuine understanding rather than superficial completion. This kind of deep, continuous professional development requires dedicated funding, time and institutional support. However, such investments have not always received the same level of attention within existing planning frameworks, which often place greater emphasis on research outputs, infrastructure or administrative expansion.

> *"For institution-wide AI deployments, such as Claude/ChatGPT/Grok licensing, a dedicated committee is constituted to assess implementation considerations. This includes evaluating licensing options, determining user access, estimating token consumption and establishing governance mechanisms."* ***- Symbiosis University***

21. Insights on India
22. Enterprise Generative AI Pricing Models and Cost Strategies

Source: EY-Parthenon analysis

AI adoption in higher education is outpacing institutional readiness, creating a gap between widespread use, limited capacity and weak governance

Conclusion

Closing the section: The question that remains

Students are already making use of AI tools in various aspects of their academic work, although approaches to usage often vary across individuals and contexts. In many cases, guidance and institutional frameworks are still evolving, leading students to develop their own practices and interpretations of responsible AI use.

Faculty members are also navigating the implications of AI in teaching and learning. While many are interested in adapting their approaches, doing so can be challenging alongside existing teaching, research and administrative responsibilities. Developing AI-enabled pedagogies, redesigning assessments and building familiarity with emerging tools often require additional time, resources and support.

At the institutional level, universities are increasingly signaling their commitment to AI through strategic initiatives, working groups and policy discussions. However, translating these ambitions into consistent practice remains an ongoing process. Areas such as governance, policy development, funding, technical support and implementation frameworks are still maturing in many institutions.

As a result, the key challenge for higher education is not whether AI will influence teaching and learning, but how effectively institutions can develop the capabilities, governance structures and supporting infrastructure needed to guide its adoption in ways that align with educational goals.

Source: EY-Parthenon analysis



What the industry actually wants

Chapter 3

While AI is becoming common in education, employers are redefining what it truly means to be "AI-ready" at work

3.1 The AI skills employers trust

From an employer's perspective, AI readiness extends far beyond the ability to operate AI tools. Prompting proficiency alone does not demonstrate the capacity to solve complex business problems, just as faster execution does not necessarily produce better outcomes. Employers increasingly value individuals who can apply AI with purpose, exercise sound judgment, critically evaluate its outputs and use it to create meaningful business value.

The central question is not whether a candidate can use AI, but whether they can use it to improve productivity, strengthen the quality of work and support better decisions while maintaining accuracy, managing risk and avoiding unnecessary complexity.[23]

Primary interviews reinforce this distinction. Employers associate AI readiness with candidates who combine strong subject fundamentals and critical thinking with the ability to use AI to improve the speed, productivity and quality of their work.

Employees with strong AI capabilities can automate routine tasks and devote more time to higher value activities such as strategic planning, creative problem solving, innovation and relationship-building. AI proficiency is therefore becoming an important recruitment consideration, with organizations increasingly assessing whether candidates can work effectively alongside AI rather than simply operate individual tools.[24]

3.1.1 Baseline AI literacy: A work readiness skill, not a tech specialist skill

AI literacy is no longer confined to technical roles. It is becoming a foundational workplace capability across functions and professions.

An estimated 33% of entry-level tasks globally are now performed by AI, rising to 37% in India. As routine work becomes increasingly automated, the effects are becoming visible across industries. AI-powered chatbots can perform tasks previously distributed across several junior employees, predictive analytics dashboards can assume responsibilities traditionally handled by supervisors, and AI-enabled procurement platforms can automate activities that once required larger operational teams.[25]

Employers increasingly expect workers in both technical and non-technical roles to operate effectively in this environment. In India, 91% of organizations consider AI skills important even for non-technical positions, while 61% report difficulty finding talent with the capabilities required to succeed in an AI-enabled workplace. The gap is therefore not simply one of access to AI tools, but of knowing how to use them appropriately in a professional context.[25]

"Students cannot afford to be AI shy. The expectation today is not just to be AI aware, but to be AI conversant. You cannot simply use AI as a tool to reduce your workload, do the heavy lifting, or cut corners. As recruiters, we expect candidates to demonstrate maturity, sound judgment and discretion in how they use AI. The key is understanding what to leverage AI for, what not to rely on it for and how to apply it responsibly and effectively." **- Avantika Tomar, Partner, EY-Parthenon**

"AI has not fundamentally changed what we look for in candidates. We still assess core problem-solving, critical thinking and technical fundamentals, with AI not permitted during case or technical rounds. At the same time, we value candidates who can use AI effectively to improve productivity while applying their own judgment to evaluate and build on its outputs." **- Campus Recruiter, Established e-commerce company**

23. LinkedIn
24. Fischer Phillips
25. *The AI Workforce Pulse: The Adaptability Imperative*, Cognizant and Pearson, 2026

Source: EY-Parthenon analysis

As AI reshapes entry-level roles, employers are prioritizing problem framing, judgment and human skills over basic task execution

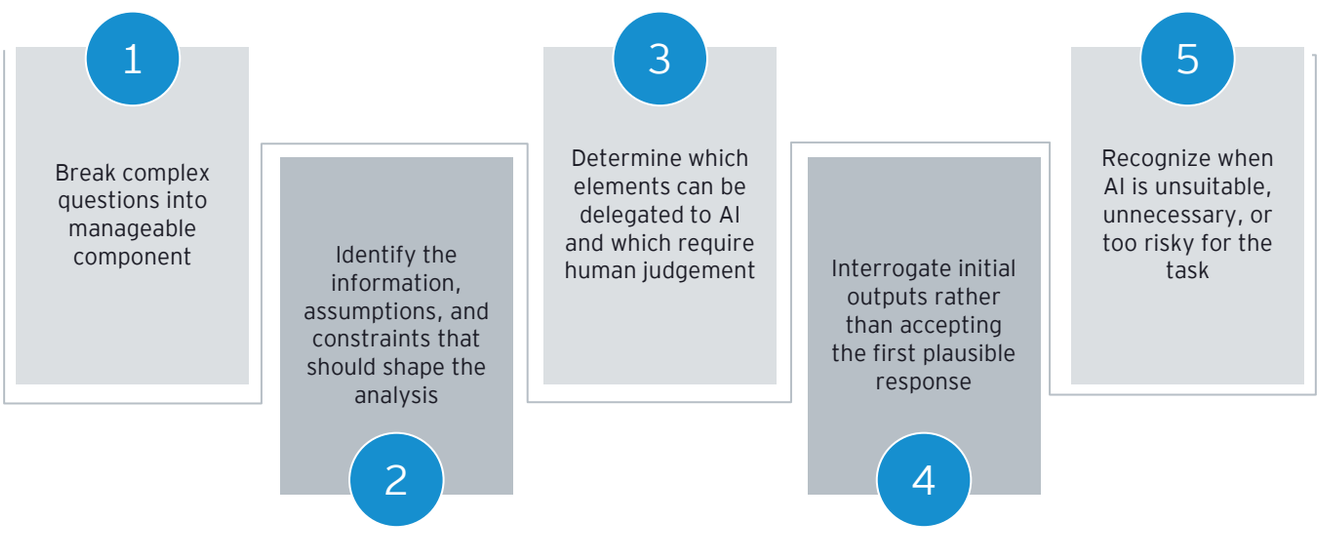
3.1.2 The skill employers trust most: Problem framing before prompting

The most valuable AI skill is not the ability to write sophisticated prompts. It is the ability to identify which problems are worth solving, determine whether AI is appropriate for the task and define the question clearly enough to generate a useful outcome.

Prompting and problem framing may appear similar, but they require different levels of thinking. A prompt is an instruction given to a tool. Problem framing requires the user to understand the objective, context, constraints, evidence and decision involved before approaching the tool.

As one interviewee observed, *"With traditional search engines, you still had to understand the problem and figure out what to search for. Now, many students can paste the entire question into AI without first thinking through what is actually being asked." – IIT Mandi, Student FGD*

Employers increasingly value people who can:



Employers therefore place a premium on candidates who can distinguish useful evidence from indiscriminate or low-value research, decide what should be delegated to AI and what requires human judgment and use AI to test and refine their thinking rather than replace it. The ability to generate more information is of little value if a candidate cannot identify which information is relevant to the question at hand.[26]

Strong problem framing requires intellectual discipline. The best questions often emerge only after the user has resisted the urge to seek an immediate answer and first clarified what they are actually trying to understand or decide.

The biggest shift is that students need to ensure the content is valid and relevant to what is being asked. The thinking cannot be outsourced, even if the research can be. AI can assist with gathering information, but critical thinking, judgment and contextualization remain your responsibility. ***- Avantika Tomar, Partner, EY-Parthenon***

26. WEF Job and Future of Work 2026

Source: EY-Parthenon analysis

As AI masks performance, employers are shifting focus from polished outputs to verification, integrity and real understanding

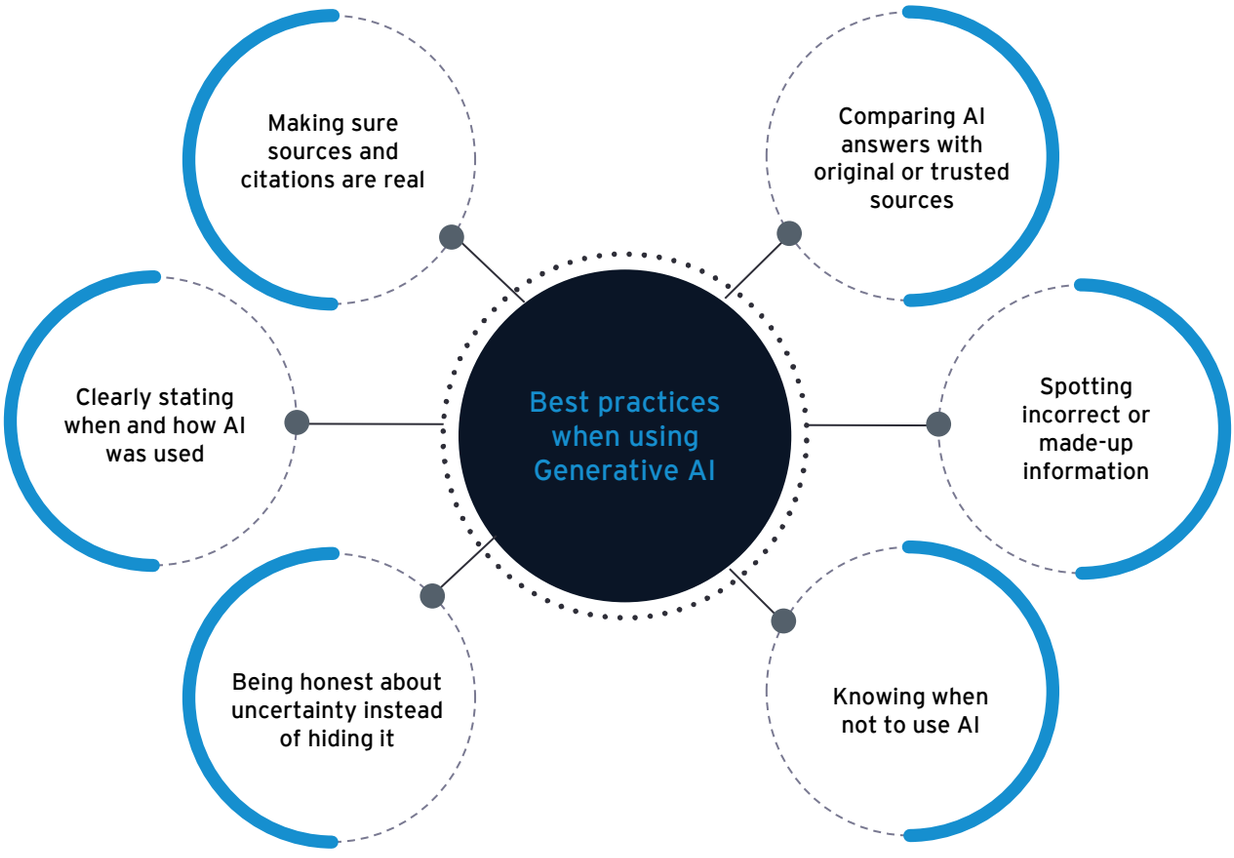
3.1.3 Verification, source discipline and intellectual honesty

Before the widespread use of generative AI, weak capability was often visible through poor writing, disorganized structure, or shallow analysis. AI can now produce polished and coherent work that conceals these weaknesses, making surface quality a less reliable indicator of genuine understanding.[27]

Verification has therefore become a core professional skill. This involves more than checking whether an AI- generated claim is factually correct. AI can produce large volumes of plausible information with little regard to what is most relevant or useful. Employers need individuals who can trace important claims to credible sources, identify unsupported assumptions, distinguish useful evidence from generic research and separate confidence of presentation from quality of evidence.

Critical thinking, adaptability, communication, problem-solving and decision-making remain essential because they allow employees to interpret AI outputs rather than merely reproduce them.

"Leadership, communication, critical thinking, problem-solving and decision-making will continue to be among the most valuable workplace skills, even in an AI-enabled world. MBA is likely to remain highly relevant, as MBA programs are designed to develop these distinctly human capabilities alongside business acumen, strategic thinking and leadership effectiveness" **– Dr Raman, Vice Chancellor, Symbiosis International University**



Global frameworks similarly emphasize that responsible AI use requires ethical awareness, critical thinking and an understanding of how AI systems work, rather than the ability to operate them alone. [28]

27. Bristow Holland, 2026
28. UNESCO

Source: EY-Parthenon analysis

As AI handles routine tasks, communication and human judgment are becoming the key skills that set graduates apart

3.1.4 Communication: The human skill AI has made more important, not less

AI can process information and execute routine tasks at speed, but it remains limited in situations requiring contextual judgment, interpretation, empathy and accountability. As a result, communication is becoming more important rather than less.

Active listening and clear communication allow individuals to interpret meaning that may not be explicit in the available data, make better decisions and build stronger professional relationships.[29]

Employers increasingly value graduates who can explain AI-assisted work in clear and accessible language. Employees should be able to:

- Describe how AI was used and what contribution it made.
- Explain the assumptions, evidence and reasoning behind a recommendation.
- Justify key decisions and tradeoffs to clients, managers and users.
- Translate AI-generated information into practical and decision-relevant insights.
- Communicate uncertainty and limitations without hiding behind jargon.
- Adapt the message to the knowledge and expectations of different audiences.

As AI becomes more capable of generating polished answers, the differentiating capability becomes metacognition: the ability to understand, question and explain one's own reasoning. This makes the communication of thought processes and decisions particularly valuable in collaborative and leadership settings.[30]

Research published in the Journal of Intelligence similarly argues that critical thinking is becoming more important because AI systems remain capable of producing errors and therefore require human oversight. Employees who can evaluate evidence, challenge assumptions and communicate their reasoning would remain valuable even as automation expands.

The World Economic Forum's Future of Jobs Report 2025 reinforces this view. It identifies analytical thinking, creativity, resilience, flexibility, curiosity and continuous learning as capabilities that will remain critical alongside technical skills in AI and data.

29. Havard Business School
30. Forbes, 2025 , Journal of Intelligence and World Economic Forum's Future of Jobs Report 2025
Source: EY-Parthenon analysis

As AI reshapes assessments, CGPA is losing importance while real skills and proof of ability gain greater value

3.1.5 Why Cumulative Grade Point Average (CGPA) is weakening, but not disappearing

The growing emphasis on skills-based hiring does not mean that CGPA has ceased to matter. Its role is better understood as changing across stages of recruitment.

NACE's *Job Outlook 2026* found that 42% of employers continued to use GPA as a screening criterion, while nearly two-thirds reported using skills-based hiring practices, particularly during screening and interviews. Grades therefore remain an efficient initial filter, particularly in high-volume recruitment, even as employers look for stronger evidence of capability at later stages.[31]

Primary interviews reinforce the continuing importance of CGPA in this role.

> *"CGPA is still very relevant because when we receive 400 CVs and need to shortlist 40 profiles, it serves as an effective first filter. As a professor, I may not completely agree with relying on it, but as a Partner, I believe it remains highly relevant as one of the inputs. Today, I also look for projects or internships that demonstrate the use of AI, prompt engineering, or effective query management. I am interested in seeing evidence of genuine and meaningful AI application. However, AI has become such a buzzword that almost everyone now mentions it on their CV. The real differentiator is whether they can demonstrate practical understanding and responsible use of AI, rather than simply listing it as a skill."* **- Avantika Tomar, Partner, EY-Parthenon**

The distinction is important. CGPA remains useful because it allows recruiters to process large applicant pools quickly and provides a common academic reference point. What is weakening is its usefulness as a **standalone proxy for capability**. Once candidates move beyond the initial shortlist, employers increasingly assess technical knowledge, practical experience, communication, attitude, problem-solving and the ability to learn.

> *"Earlier, we used to look primarily at CGPA and the courses a candidate had taken. While those factors are still important, we now also assess how candidates perform on practical tasks such as vibe coding. The focus has shifted from what you know to how effectively you can apply that knowledge in practice."* **- Mr Kalyan Kolachala, India Head, Symphony AI**

This becomes particularly relevant as AI makes assignments, project submissions, reports and other academic outputs easier to produce. Where employers have less confidence that these outputs reliably represent individual capability, interviews, live tasks, internships, portfolios and work samples become more useful complements to grades.

CGPA is therefore not disappearing. It remains a valuable screening mechanism, particularly where applicant volumes are large. What is changing is the amount of work employers expect it to do: grades may help determine who enters the assessment process, but they are less likely to determine who ultimately succeeds within it.

31. NACE Job Outlook 2026

Source: EY-Parthenon analysis

As industry voices grow clearer, the gap between what universities produce and what employers need is becoming harder to ignore in the AI era

3.2 The truth from CHROs, CTOs and hiring panels

A significant gap persists between higher education and industry requirements in India. Recent research suggests that nearly 75% of higher education institutions are not fully equipped to meet employers' needs.[32] One of the principal reasons is the limited alignment between academic curricula and rapidly changing workplace requirements.

This challenge is particularly acute in fields such as artificial intelligence and cybersecurity, where technologies and professional practices evolve faster than many institutions can update their programs. Outdated or weakly relevant curricula therefore remain a major constraint on graduate employability.[33]

Recruitment practices are changing in response. Consulting firms are increasingly asking applicants how they would use AI chatbots during the recruitment process, and some have introduced AI-assisted interviews as part of final-round assessments. Candidates may be presented with a business problem similar to real consulting work and expected to use AI to gather information, structure their thinking, test hypotheses and refine insights. The focus is not on letting AI provide the answer, but on demonstrating sound judgment in how it is used to strengthen problem-solving and decision making.[34]

3.2.1 What is missing in Indian graduates

India faces a persistent graduate employability challenge. Millions of students enter the workforce each year, yet the *India Graduate Skill Report 2025* estimates that only 42.6% meet prevailing industry standards for job readiness.

A recent HirePro report highlights the scale of the perception gap: only 9% of corporates believe fresh graduates are job-ready on day one, while 68% of students believe they are. This disconnect reinforces the need for clearer employer feedback loops, stronger work-integrated learning, and more credible evidence of practical capability.

Employers consistently identify gaps in the practical application of knowledge, industry exposure, communication, presentation, problem-solving and critical thinking. Deficiencies in digital capability and adaptability further constrain graduates' ability to perform in technology-enabled workplaces.

Primary interviews similarly suggest that the challenge is not simply whether graduates have used AI. Employers value candidates who can apply knowledge in practical settings, experiment with tools, distinguish useful information from noise and solve real problems. Graduates need to be able to use AI to build, test, verify and improve outcomes without allowing the tool to substitute for their own understanding.

NASSCOM and Indeed's 2026 report similarly indicates that AI is reshaping hiring expectations in India: 86% of employers report some impact of AI on job roles and responsibilities, while 40% prefer demonstrable AI skills or certifications over degrees. This does not imply that formal qualifications are becoming irrelevant. Rather, employers increasingly expect them to be accompanied by evidence that candidates can apply their knowledge in practice.[35]

32. From Degree Factories to Employability Hub, Team Lease EdTect 2026
33. Fortune
34. The Guardian
35. Nasscom Indeed India's AI Talent Inflection Point 2026 , India Graduate Skill Report 2025, HirePro

Source: EY-Parthenon analysis

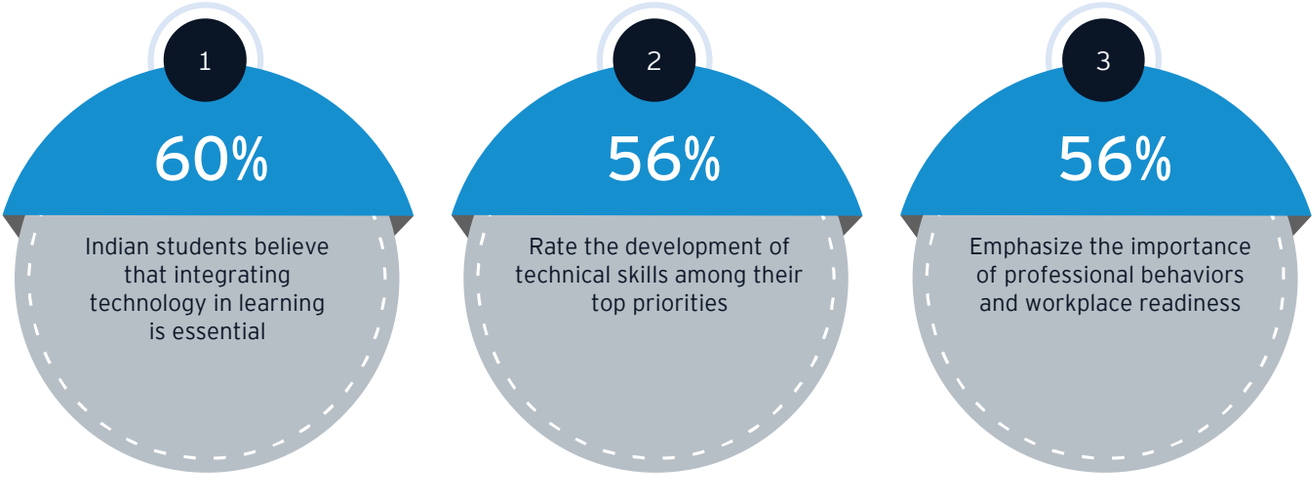
While gaps in institutional readiness persist, Indian students' adaptability and ambition could become a key advantage in the AI-driven global workforce

3.2.2 Where Indian students may outperform global peers

India's advantage does not lie in the readiness of its universities for the AI era. It lies in the willingness of its students and workforce to experiment, learn quickly and adopt new technologies.

This adaptability represents a meaningful strength, but it is not self executing. Without adequate institutional and employer support, the benefits of AI adoption may remain concentrated among students with stronger networks, better infrastructure and greater access to high-quality learning opportunities. The challenge is therefore to convert individual enthusiasm into broad-based, structured and sustainable capability.

Student expectations are already changing. Learners increasingly want higher education to offer more than lectures and academic knowledge. They expect practical learning, stronger employability outcomes and clearer pathways to global careers. A survey of 3,000 students and parents found that Indian respondents place greater emphasis on hands-on learning and employment preparation than students in many other countries.[36]
In our FGDs, students from the Vellore Institute of Technology emphasized that their university experience extends beyond academic learning. They highlighted the importance of universities in providing opportunities for networking, skill development, extracurricular engagement and placements, noting that while AI can support learning, it cannot replace the broader value and experiences offered by higher education institutions.



Source: Economic Times

India also begins from a position of scale. NITI Aayog's AI jobs roadmap notes that the country already has more than nine million professionals working in technology and customer-experience roles. It argues that India could emerge as a global center for AI talent and innovation by 2035 if government, industry and educational institutions act in coordination.[37]

Microsoft and LinkedIn's India Work Trend Index similarly indicates that Indian knowledge workers use AI at rates above the global average. This is encouraging, but adoption should not be confused with proficiency. Frequent use of AI does not necessarily imply that users can evaluate outputs, apply the technology responsibly, or generate better professional outcomes.

India's comparative advantage therefore depends on whether it can translate high adoption and ambition into stronger reasoning, verification, judgment, and domain-specific application.

36. Economic Times
37. Press Information Bureau
Source: EY-Parthenon analysis

As AI reshapes work, the real gap is shifting from generating answers to reasoning, verifying and exercising judgment

3.2.3 AI specific gaps, reasoning, verification, judgment

The traditional employability gap centered on whether graduates could apply academic knowledge in practice, the transition from "knowing" to "doing." In the AI era, the gap is evolving. When plausible outputs can be generated almost instantly, employers increasingly distinguish between candidates who can produce an answer and those who can determine whether the answer is sound.

Three gaps require particular attention.

- **Reasoning gap:** A candidate may be able to produce an answer but struggle to explain how they arrived at it. They may find it difficult to distinguish correlation from causation, summary from analysis or an answer from a well-reasoned argument.
- **Verification gap:** A candidate may trust AI outputs too readily without checking sources, numbers, assumptions or facts. They may also struggle to distinguish relevant evidence from generic research, leaving errors or weak assumptions unchallenged.
- **Judgment gap:** A candidate may not recognize when AI should not be used. They may overlook risks related to privacy, confidentiality, bias, regulation or reputation, treating AI as a universal source of answers rather than a tool with context-dependent limitations.

The problem is, therefore, not simply overuse. It is the failure to retain responsibility for the reasoning, evidence and final judgment.

UNESCO's AI competency frameworks[38] reflect this broader conception of readiness. They emphasize a human-centered approach built on critical thinking, ethical awareness and responsible use rather than technical operation alone.

UNESCO also recommends that AI competency frameworks for students and teachers form part of a broader institutional strategy for capacity building. This includes universal digital access, ethical design principles, responsible deployment and greater attention to the environmental effects of AI systems.



38. UNESCO's AI competency frameworks

Source: EY-Parthenon analysis

As AI removes routine tasks, entry-level roles are shifting from learning through doing to exercising judgment from day one

3.2.4 The entry-level work paradox

AI is not only automating entry-level tasks but also reducing the volume of low-risk work through which junior employees traditionally developed professional capability.

In the past, new employees learned by completing relatively simple tasks, receiving feedback and gradually assuming greater responsibility. These activities provided a structured environment in which they could build judgment, understand workflows and learn from manageable mistakes.

As AI takes over a larger share of routine work, graduates may enter roles in which part of this traditional learning layer has already been removed. They may be expected to handle exceptions, interpret outputs, engage with clients and make decisions earlier in their careers, even though they have had fewer opportunities to develop these capabilities.

According to research by Cognizant and Pearson:[39]

- 96% of HR leaders expect entry-level roles to shift towards supervising, managing, or working alongside AI systems.
- 94% expect AI to create new categories of entry-level employment.
- 97% believe that communication, critical thinking, teamwork, problem-solving and judgment will become more important as AI adoption increases.

In practice, fresh graduates are likely to spend less time completing repetitive tasks and more time reviewing AI outputs, interpreting results, managing exceptions and determining when human intervention is required.

The role of the entry-level employee is therefore beginning to resemble that of an "air traffic controller", coordinating AI-supported workflows, validating outputs, identifying risks and stepping in where contextual or professional judgment is required.

39. Cognizant and Pearson The AI Workforce Pulse: The adaptability imperative 2026

Source: EY-Parthenon analysis

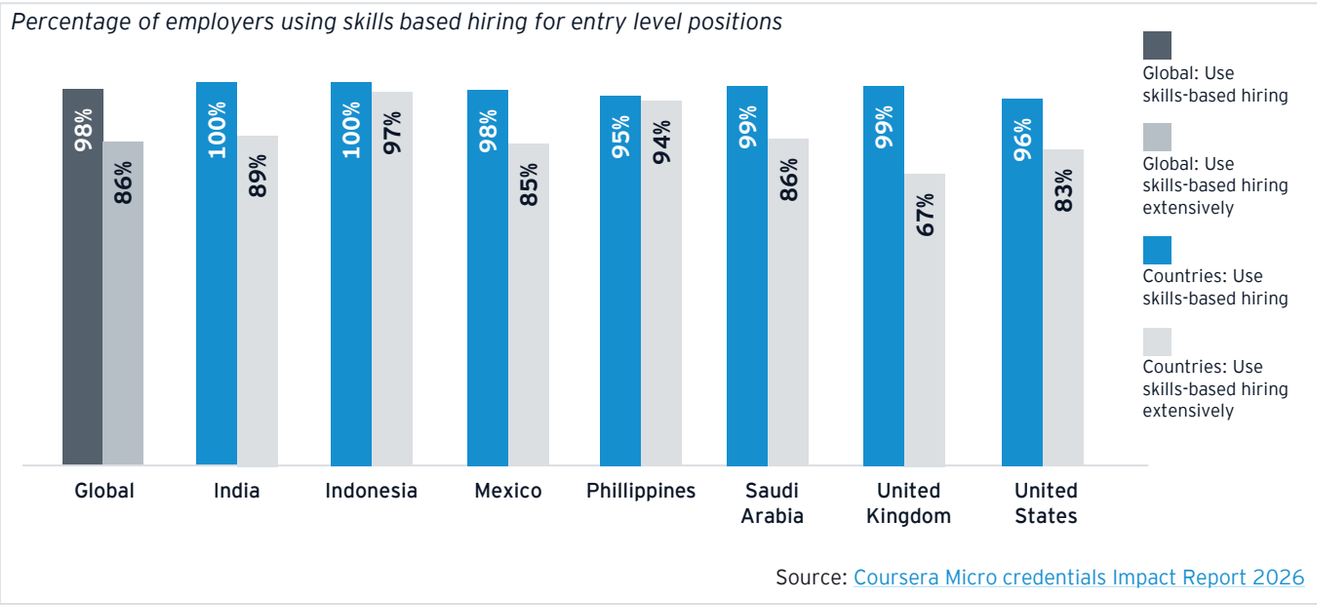
As skills-based hiring reshapes entry-level roles, employer trust is shifting from degrees and certificates to capability and proven work

3.3 Micro-credentials, certifications or internships – What has value?

Skills-based hiring is becoming increasingly important for entry-level roles. By 2030, employers expect nearly one-third of core job skills to change, while 86% already rely extensively on skills-based hiring for entry-level recruitment.[39] Graduates therefore need credible evidence of job-relevant capability that employers can assess quickly.

As hiring shifts towards skills-based evaluation, the signals available to employers can be grouped into three broad categories.

Fig 4: Employers turn to skills-firts hiring



Source: Coursera Micro credentials Impact Report 2026

1 Signals employers increasingly trust

Internships, portfolios, work samples and applied projects are among the strongest signals because they provide tangible evidence of how candidates approach problems and apply their skills in practice. Well-designed micro-credentials can also differentiate candidates when they incorporate rigorous assessment and evidence of real-world application.[40]

According to Coursera, 97% of employers say entry level hires with micro credentials perform better in the first year.[40]

2 Signals employers consider useful but limited

CGPA, short courses and introductory certifications can support screening or indicate exposure to a subject, but they provide limited evidence of whether a candidate can apply the underlying knowledge in practical settings.

3 Signals that are becoming resume decoration

Generic certifications and disconnected badges are becoming less differentiating as they proliferate. Listing a credential or familiarity with a tool without demonstrating how it was applied is increasingly a weak signal in a skills-based hiring environment.

39. NACE Job Outlook 2026
40. Coursera Micro credentials Impact Report 2026

Source: EY-Parthenon analysis

As credential overload grows, employer trust is shifting from the number of certificates to the depth, relevance, and proof of real skills

3.3.1 The credential inflation problem

Students today accumulate a wide range of credentials and titles, including Coursera certificates, LinkedIn Learning badges, Google, IBM and Microsoft certifications, short AI boot camps, prompt engineering courses, lightly supervised internship letters, hackathon participation certificates, and campus leadership positions.

Indian graduates may be the most credentialed generation yet. Their résumés contain more courses, specializations, and digital certifications than ever before, but many remain underprepared for the realities of professional work. This contradiction lies at the center of India's employability paradox: students are accumulating qualifications faster than their readiness for employment is improving.[41]

Part of the explanation lies in credential inflation. The growth of online learning platforms and short-term certification programs has democratized access to specialized knowledge. In principle, this is a positive development. The problem arises when credentials accumulate more rapidly than the underlying knowledge is understood, practiced, or retained. In our FGDs, students also highlighted that some peers overstate or exaggerate the scope and impact of their projects on professional platforms such as LinkedIn.

As a result, the certificate itself is increasingly difficult to interpret. Employers need to know what it actually represents: whether it involved meaningful assessment, applied project work, mentorship, external validation, or genuine engagement with the subject.

A hiring manager described a campus recruitment experience involving a student whose resume included certifications in data analytics, digital marketing, AI tools, and academic competitions. Although the candidate appeared exceptional on paper, they struggled to frame an ambiguous real-world business problem that did not have a clearly defined answer.[42]

Primary interviews reinforce this concern.

> *Credentials are often difficult to distinguish between certifications that are truly valuable and those that are not. Certifications carry limited weight on a résumé alone. Their real value becomes apparent during interviews, where you can assess whether a candidate genuinely understands the subject matter and can apply the knowledge in practice. As a result, live projects and hands-on experience are generally preferred over certifications, as they provide stronger evidence of real-world capability and problem-solving skills.* ***- Avantika Tomar, Partner, EY-Parthenon***
>
> *"We place limited weight on certifications alone, as they are not always a reliable indicator of capability. What matters more is evidence of applied learning, such as research projects, papers, or hands-on work in areas like LLMs and AI (for tech roles). Relevant certifications may add context, but they are generally a secondary consideration."*
> ***- Campus Recruiter, Established e-commerce company***

The implication is not that credentials have no value. It is that they should be treated as evidence of **possible capability rather than proof of capability**. Their value rises considerably when the candidate can demonstrate the learning through an interview, project, work sample, or other applied setting.

41,42. Hindu

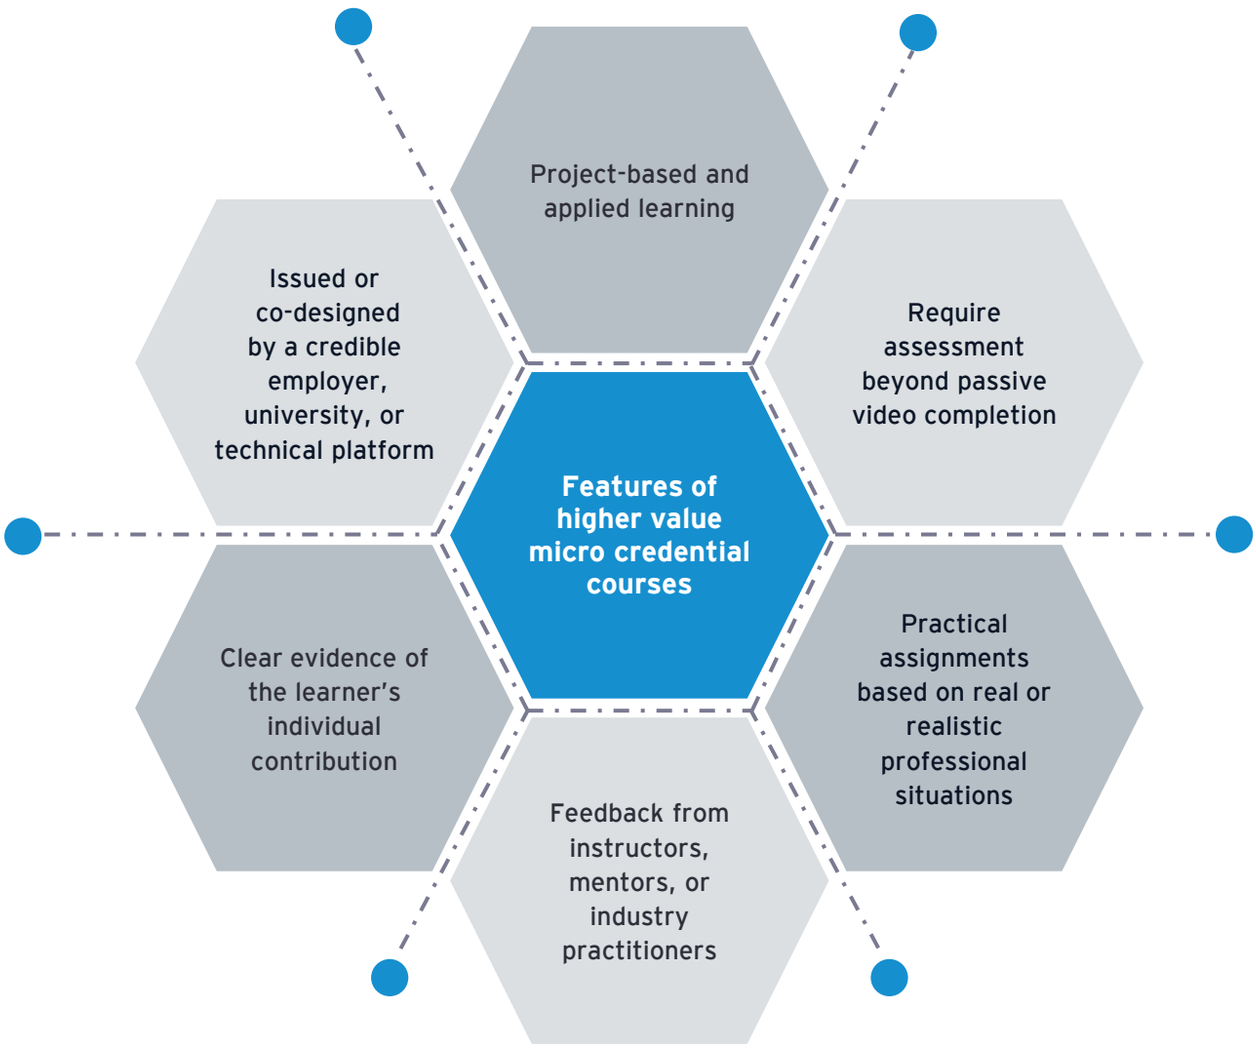
Source: EY-Parthenon analysis

As workplace complexity rises, employer respect is shifting from basic knowledge to problem-solving, adaptability and real-world judgment

3.3.2 Which micro credentials employers may respect

Most graduates already possess a basic understanding of common digital tools and management concepts. Another introductory certificate is, therefore, unlikely to distinguish them meaningfully in the labor market.

Higher-value micro-credentials are more likely to be closely connected to a role, require meaningful assessment and applied work, and make the learner's individual contribution visible.



Source: EY-Parthenon analysis

As learning shifts toward hands-on execution, micro-credentials are moving from theory-based exams to real-world projects and job-ready skills

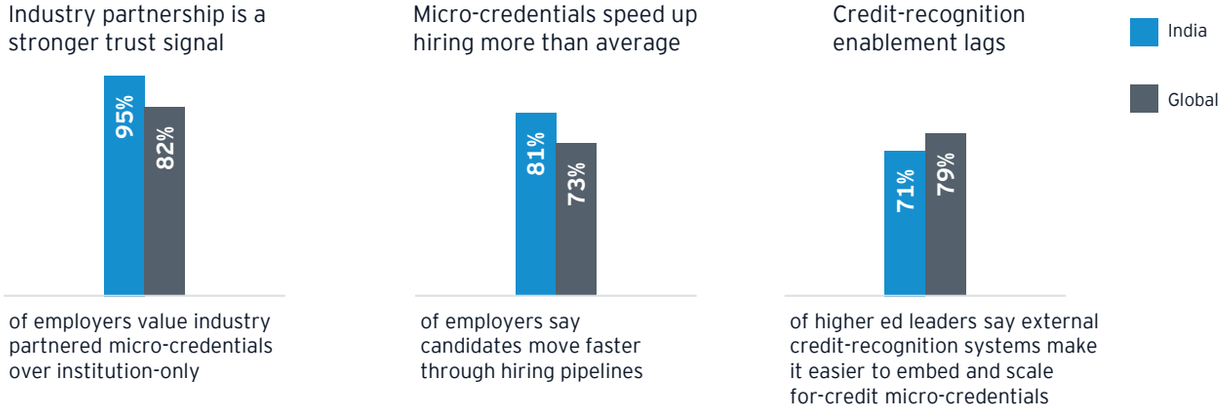
Professional certificates offered by organizations such as Google, IBM and Microsoft increasingly incorporate case studies, practical assignments, hands-on laboratories, or workflow-based assessment. These features make the credential more informative because employers can infer more about what the learner was required to do, rather than simply whether the course was completed.

Universities are also beginning to integrate such credentials directly into academic programs rather than treating them solely as optional additions. This can help institutions respond more quickly to changing workforce requirements. Formal curriculum revision may take several years, whereas well-designed micro-credentials can be updated more frequently as tools, practices and employer expectations evolve.

Fig 5: India stands out for the strength of its labor market demand for micro-credentials.

Employers prioritize skills-based hiring and see value in candidates who can demonstrate job ready skills. For universities, the opportunity is not to prove the case for micro-credentials—it is to scale them faster through credit-aligned, industry-validated pathways that are easier to embed into degrees.

GLOBAL: Key takeaways



Source: Coursera Micro credentials Impact Report 2026

The *Coursera 2026 Micro-credentials Impact Report* indicates that 92% of employers believe entry-level employees with micro-credentials perform better during their first year, while 96% have hired at least three candidates with micro-credentials during the previous year.[43]

These findings point to strong employer interest, but primary interviews add an important qualification: recognized credentials are most persuasive when candidates can explain what they learned and where they applied it. Micro-credentials can strengthen the signal; they do not remove the need to demonstrate capability.

43. Coursera Micro credentials Impact Report 2026

Source: EY-Parthenon analysis

As internships evolve from passive participation to active contribution, employer focus is shifting from certificates of attendance to evidence of real capability

3.3.3 Internships: from attendance certificate to capability evidence

Internships can provide an important bridge between academic learning and professional employment by exposing students to workplace expectations, real-world problems, feedback and accountability.

Strada Education Foundation's 2025 Work-Based Learning Report finds that experiences aligned with a student's field of study can provide an effective bridge between education and employment, although access remains limited. Students also primarily value these experiences as a means of developing skills and experience for a planned career rather than simply exploring possible occupations.[44]

However, the term internship now covers experiences of very different quality. Employers are therefore likely to distinguish among four broad types.

1 Genuine internships with real tasks, feedback and deliverables

High-quality internships require students to contribute to meaningful assignments in actual work environments. Engineering students may contribute to product development, while management students may work on market research, business analysis or campaign execution.

The defining characteristics are not the employer's name or the duration of the internship, but the quality of the work, the level of supervision and feedback and the student's ability to explain what they contributed and learned.

2 Shadow internships where the student observes and does little

In shadow internships, students primarily observe professionals rather than undertake substantial work. They may attend meetings, follow a mentor and observe how tasks are performed.

These experiences can build awareness of workplace structures and professional routines, but provide limited opportunities to make decisions, solve problems or develop confidence through execution.

3 Certificate internships where the student receives documentation but no meaningful work

In certificate internships, the principal outcome is a completion letter rather than professional learning. Students may complete basic, repetitive or low-impact tasks, with limited supervision and little connection to meaningful business outcomes.

> *"Students need to undertake more real-time projects that incorporate AI exposure to become truly industry-ready."*
> ***- RB Mouli, Chief Manager of Placement Unit, BITS Goa***

Such internships may add a line to a resume but provide employers with little credible evidence of professional capability.

44. Strada Education Foundation's 2025 work-based learning report

Source: EY-Parthenon analysis

As AI reshapes internships, learning is shifting from using tools to demonstrating judgment, execution and real-world impact

4 AI-era internships where the student uses AI to accelerate output but still demonstrates judgment, communication and execution

AI era internships should allow students to use AI to improve productivity while continuing to demonstrate judgment, communication and execution.

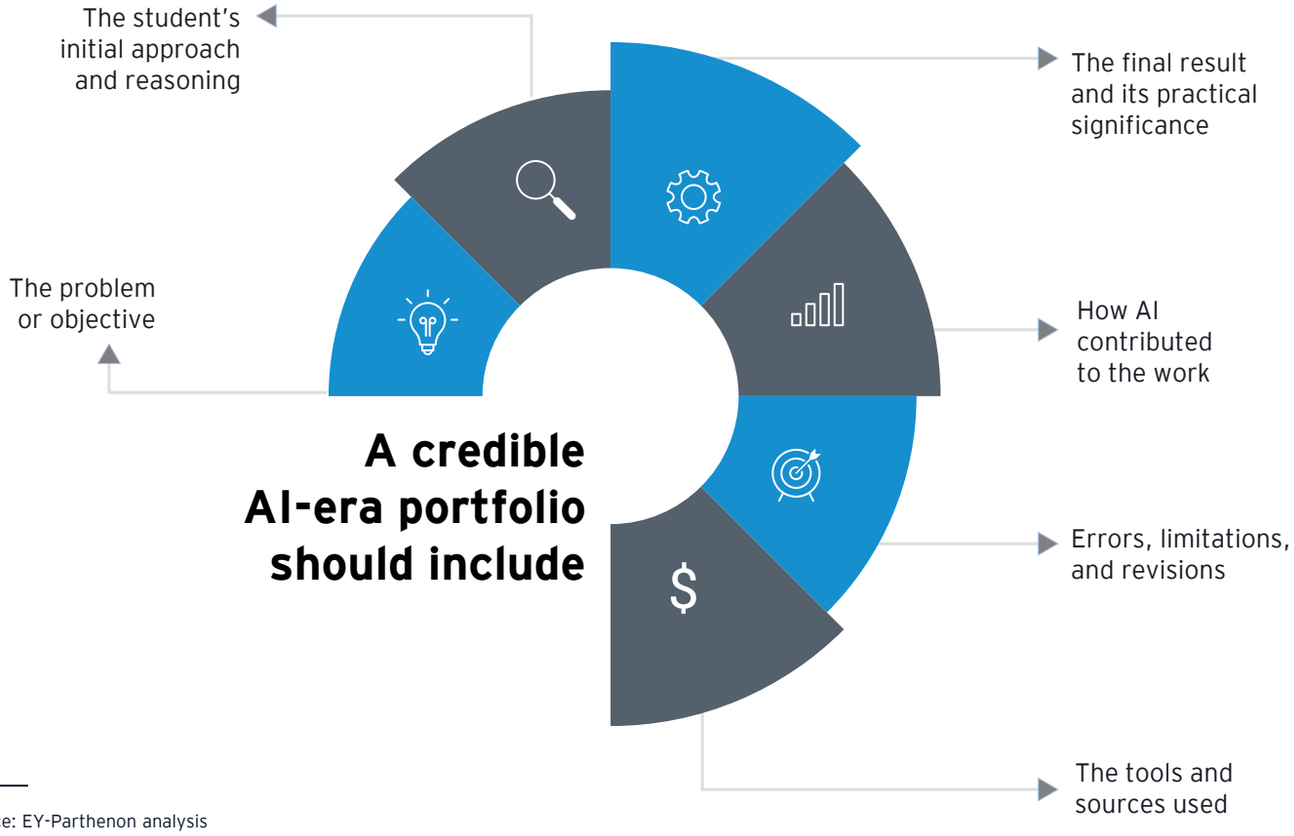
Students may use AI to accelerate research, drafting, analysis, or routine workflows. They must nevertheless remain responsible for:

- Assessing the accuracy and relevance of outputs
- Identifying and correcting errors
- Explaining their reasoning and decisions
- Communicating effectively with colleagues and stakeholders
- Managing confidentiality and data risks
- Delivering meaningful outputs and measurable results

The objective is not to restrict AI use, but to verify that internships continue to develop judgment, problem-solving, communication and accountability in workplaces where human and machine capabilities are increasingly combined.

3.3.4 Portfolios as the new middle layer between degree and job

A portfolio should not simply be a folder of polished outputs. It should provide a record of how a student thinks, works, revises, verifies and applies tools. This becomes particularly important in the AI-era, when a final output alone may reveal little about how much of the work was produced by the student and where AI contributed.



Source: EY-Parthenon analysis

As skills first hiring expands talent pools, employers are shifting from job titles to capabilities, diversity and real-world skill relevance

Portfolios can therefore provide a useful bridge between formal qualifications and employer trust. Work samples and documented projects give employers evidence of how candidates approach unfamiliar problems, use tools, respond to errors and feedback and distinguish their own contribution from AI-assisted work.

This is consistent with the wider movement towards skills-first hiring. LinkedIn data suggests that hiring based on demonstrated skills rather than degrees, certificates or previous job titles can help employers respond more effectively to changing role requirements and widen candidate pools.[45]

Their role should not, however, be overstated. In India, portfolios are unlikely to replace degrees or CGPA in high-volume campus recruitment. Their value lies primarily in later stages of evaluation, where employers need richer evidence of what a candidate can actually do.

A degree and CGPA may therefore continue to establish baseline eligibility; **the portfolio can provide the next layer of evidence about capability.**



45. LinkedIn Skills First: Reimagining the Labor Market and Breaking Down Barriers 2023

Source: EY-Parthenon analysis

As AI reshapes talent development, industry-led CoE models are emerging as operational engines to replace ad hoc partnerships with deep, scalable collaboration

3.4 Industry co-optation of the Centre of Excellence model

A deeper partnership between higher education institutions and industry requires moving beyond advisory relationships and positioning institutions as strategic co-creators of AI era capability. Localized Centres of Excellence (CoEs) can provide a structured platform for this collaboration by enabling industry and academia to shape curricula, co develop technologies, build faculty capability and accelerate the translation of research into practical deployment.

Realizing this potential would require a shift from symbolic partnerships to execution-oriented models. Higher education institutions can draw lessons from approaches in which research, product development, teaching and industry engagement take place within closely connected environments, as illustrated by the IIT Madras case study on the following page.

3.4.1 Why the current partnership model falls short

Traditional partnerships between higher education institutions and industry are not sufficient for AI era talent development. The challenge is partly one of speed: AI technologies and enterprise workflows can change within months or even weeks, while academic systems typically take much longer to adapt. But slow curriculum change is only one part of the problem.

Several structural factors make deeper collaboration difficult.

- Incentives are often weak or misaligned. Industry teams may see limited reason to commit sustained time unless a partnership creates a clear talent, research, innovation, or commercial benefit. Faculty and academic departments, meanwhile, may receive limited recognition for the time-intensive work required to build meaningful industry engagement.
- The value proposition is not always communicated clearly. Institutions and employers may agree that collaboration is desirable without having a shared understanding of why it matters, what each side is expected to contribute, or what success should look like. Without greater awareness of the benefits for each party, engagement can remain superficial.
- University leadership faces competing priorities. Rankings, research output, admissions, placements, accreditation and institutional growth all command senior-management attention. These priorities matter, but they can crowd out sustained focus on the underlying quality and relevance of teaching, curriculum and student capability.
- Finally, India still has too few widely visible success stories demonstrating what sustained industry–academia collaboration can achieve. Strong examples can show the benefits not only for institutions, but also for faculty, students and recruiters. Without such cases, it is harder to build internal buy-in or justify the effort required on either side.

These constraints help explain why most higher education institutes (HEI)-industry collaborations continue to revolve around memoranda of understanding, guest lectures, short workshops, internships, campus placements, event sponsorships and limited curriculum advisory roles. These interventions create useful connections, improve institutional visibility and support recruitment, but are often episodic and transactional rather than part of a sustained model of capability development.

Source: EY-Parthenon analysis

Beyond AI CoEs: IIT Madras Research Park highlights the opportunity to build collaborative, industry-driven innovation and talent ecosystems

The problem is not that these activities lack value. It is that they rarely change the core of how teaching and learning are designed. Students have limited access to authentic industry problems, datasets and workflows; faculty engagement is often confined to occasional workshops; and hiring outcomes are rarely fed back into curriculum or pedagogy in a systematic way.

There are, however, examples of institutions responding more quickly when the internal conditions allow it. R.B. Mouli of the BITS Goa Placement Cell observed:
Faculty at BITS Goa are very prompt in adapting to changes suggested by employers where there are no financial implication. If a change in recruiter expectations is communicated back to the institution, the turnaround time for incorporating it into courses or assessment structures can be as short as one semester.

This example suggests that the constraint is not necessarily academic unwillingness to adapt. Responsiveness can be rapid, where employer expectations are communicated clearly and implementation does not require substantial funding, new infrastructure, or formal administrative approvals.

Moving beyond basic networking would therefore require a model in which both sides see clear value and collaboration extends into curriculum, faculty development, applied learning, research and recruitment. When designed well, the CoE model can provide an institutional structure for bringing these activities together.

- **Illustrative example:** IIT Madras Research Park - The park acts as an independent corporate hub co-located on campus, providing 200+ startups and major enterprises full access to academic faculties and shared R&D equipment.
- **IIT Madras Research Park**

1.2+	250+	17+	16+	6,000+
Million sq ft of workspace	R&D companies and labs	Deep-tech sectors	Centres of excellence	Technologist

- The AI ecosystem at the IITM Research Park operates on several key fronts:
  - **Research centers and infrastructure:** It is home to the **School of AI and Data Science (WSAI)** and the Centre for Responsible AI (CeRAI), which focuses on building ethical, accountable and explainable AI systems for the Indian market.
  - **Startup incubation:** Startups at the IITM Incubation Cell are heavily leveraging AI and deep tech. Entities like Vakyam AI and Predixion AI operate directly from the park, with over 100 deep-tech startups pushing innovations in sectors from healthcare to autonomous systems.
  - **Skill development and workshops:** Through IITM Pravartak (a Technology Innovation Hub), the park offers advanced training, certification programs and boot camps in AI, Machine Learning and Generative AI for students and professionals

Sources: IIT Madras Research Park, EY-Parthenon analysis

As India builds a national AI CoE foundation, the real opportunity lies in using it as a scalable platform for industry led collaboration and distributed talent development

3.4.2 The Government CoE model as a platform for industry participation

India already has a policy foundation on which AI focused Centres of Excellence can be built. Under the IndiaAI Mission and the National Strategy for Artificial Intelligence, the Government has begun developing the infrastructure, shared resources and institutional mechanisms required to scale AI innovation.

In October 2024, the Ministry of Education announced the establishment of AI Centres of Excellence focused on healthcare, agriculture and sustainable cities. AIIMS, IIT Delhi, IIT Ropar and IIT Kanpur were designated as lead institutions. The Government approved an outlay of INR 990 crore for FY2024 to FY2028 to support the initiative.

These CoEs are intended to operate as collaborative ecosystems led by premier academic institutions in partnership with industry and startups. Their mandate includes conducting interdisciplinary research, developing AI-enabled applications and creating scalable solutions to critical sectoral challenges. [46]

The IndiaAI Mission provides a broader enabling ecosystem. With an approved outlay of INR10,300 crore, it seeks to strengthen India's AI capabilities through public compute infrastructure comprising 38,000 GPUs, access to high-quality datasets, development of indigenous AI models, startup financing and initiatives focused on safe and trusted AI. Collaboration between academia and industry is a central feature of this agenda.[47]

The objective should not be to replicate resource intensive, elite CoEs across every institution. A more practical opportunity lies in building a model that assigns different roles to institutions based on their capabilities.

Under such a model:

- Leading institutions would serve as hubs for advanced research, model development, evaluation and faculty development.
- Regional universities would function as applied domain laboratories, translating research into sector and context specific use cases.
- Tier-2 and Tier-3 institutions would act as centers for implementation, testing, data generation and talent development.
- Industry partners would participate as problem owners, data providers, mentors, funders, evaluators and eventually employers.

This structure would allow institutions to participate at different levels without requiring each of them to build the same capabilities independently.

46. PIB
47. PIB

Source: EY-Parthenon analysis

As CoE models evolve, their impact may depend on integrating curriculum, research, faculty and hiring into a continuous, outcome driven talent ecosystem

3.4.3 What an industry-co-opted CoE should actually do

A credible AI CoE should perform four interlocking functions:

Curriculum co-design

Industry and academia should jointly define the capabilities required for AI-enabled roles.

Industry partners can identify role-specific workflows, tools, tasks, risks and decision requirements based on current professional practice.

Faculty can then translate these inputs into coherent modules, assignments and assessments that fit within academic programs.

This would allow institutions to move beyond generic courses on the fundamentals of AI towards domain-specific applications such as AI in healthcare, finance, supply chains, law, manufacturing or education.

The aim should not be to teach every emerging tool, but to connect AI capability with the knowledge, judgment and workflows required in particular professions.

Faculty capability-building

Faculty development should be continuous and immersive rather than limited to occasional workshops.

Structured industry exposure can help faculty understand how AI is used in professional environments, how outputs are verified, where systems fail and which human capabilities remain essential. Industry practitioners can complement this process by co-teaching applied components of academic programs.

This two-way exchange would help faculty remain current while ensuring that teaching continues to be grounded in pedagogy rather than driven solely by the latest technology.

Applied research and problem banks

Industry partners could contribute authentic problem statements that students and faculty could address through coursework, capstone projects, or research.

These challenges may be supported by anonymized, synthetic, or appropriately governed data to protect privacy, intellectual property and commercial confidentiality.

Well-designed problem banks would expose students to the ambiguity, constraints and tradeoffs of real professional work.

The success of these projects may be evaluated not only through academic publication but also through practical relevance, quality of execution, scalability and potential for implementation.

Talent pipelines

The CoE could create a structured pathway connecting coursework, applied projects, internships and employment.

This way, employers could observe student capability over time rather than relying on a single interview or resume.

They could also assess how students frame problems, respond to feedback, work in teams, apply AI, verify outputs and improve across multiple stages.

This would create a more credible hiring process based on demonstrated performance, consistency and growth.

Source: EY-Parthenon analysis

As CoE models expand, their true impact may depend on whether they broaden access or reinforce existing inequalities in talent and opportunity

3.4.4 The equity risk: Preventing CoEs from becoming another elite campus advantage

The CoE model can either reduce inequality in higher education or deepen it, depending on how it is designed.

If AI CoEs remain concentrated within a small group of elite institutions, they may reinforce existing advantages. Students at these institutions already benefit from stronger faculty networks, better infrastructure, greater industry exposure and more established recruitment relationships. Additional investment may advance research and innovation but have limited impact on talent development across the broader higher education system.

State universities and affiliated colleges, where a large share of Indian students are enrolled, may receive limited benefit. Employer networks could become more concentrated, potentially narrowing recruitment funnels and widening the gap between institutions.

A different outcome may be possible if CoEs are designed not as institution-specific assets, but as regional or networked hubs that extend their capabilities to a wider group of universities and colleges. An anchor institution could provide shared access to infrastructure, programs and industry partnerships, while participating institutions connect to the same platform rather than building these capabilities independently. Tier 2 and Tier 3 institutions could then gain access to curated industry problem statements, modern AI tools and shared infrastructure, faculty training and mentoring, standardized assessment rubrics, externally validated projects and structured links to internships and employment.

Shared standards may also make it easier for employers to identify capable students beyond a narrow set of campuses. Over time, CoEs could support regional talent pipelines aligned with local industries and economic clusters.

This vision reflects the idea of an "AI Highway", a common backbone of infrastructure, knowledge, standards and opportunity that enables institutions at different levels to participate in the AI economy.

It is also consistent with UNESCO's guidance on AI in education[48], which emphasizes inclusion and equity as central principles of technological adoption. The success of the CoE model may therefore be assessed not only through the quality of research produced at leading institutions, but also through the extent to which it expands access to capability, opportunity and employment across the wider higher education system.

48. UNESCO Guidance on AI in Education

Source: EY-Parthenon analysis



India vs. the world: Where India stands in AI-driven higher education

Chapter 4

Leading universities are operationalizing AI governance, while many institutions are still building policy frameworks

4.1 What are the top global universities doing differently and where do Indian HEIs stand? A comparative diagnostic

AI is becoming increasingly embedded in higher education. Students and faculty may already be using AI extensively, but widespread usage does not necessarily translate into institutional readiness. A number of leading global universities are beginning to move beyond experimentation toward more institution-wide approaches.

Rather than treating AI as a collection of isolated tools or standalone courses, they are embedding it into academic and institutional systems. These emerging practices provide useful benchmarks for how AI can be integrated in a more structured, responsible and scalable manner.

Indian institutions are also progressing, with several leading universities developing formal policies, institutionally licensed tools and more deliberate approaches to AI integration. However, implementation across the broader sector remains uneven. The comparative question is therefore not whether Indian HEIs are adopting AI, but how far they have progressed in building the institutional capabilities required to support that adoption effectively.

4.1.1 Operating model, governance and policy maturity: Decision rights, policy to practice, change management

AI is increasingly moving from experimental adoption to embedded practice within higher education. Leading institutions are shifting their focus from offering standalone AI courses to integrating AI capabilities across programs, disciplines and learning experiences, thereby positioning AI fluency as a core graduate attribute rather than a specialized technical skill.[49]

Global benchmark

The global benchmark is not simply the existence of an institution-wide AI policy. It is the ability to translate that policy into day-to-day academic and administrative practice through clear decision rights, course-level guidance, student expectations, tool approval processes, procurement reviews, data protection, risk-based oversight and regular policy updates.

Northeastern University has adopted a university-wide approach to embedding AI fluency throughout the curriculum. Rather than treating AI as a separate subject area, the university seeks to develop AI-related competencies across disciplines through classroom instruction, project-based learning and experiential education, including industry placements.

To support this effort, Northeastern drew on the Digital Education Council (DEC) AI Literacy Framework as a foundational reference and adapted it to its own institutional context and academic priorities. Implementation is faculty-led and bottom-up, with representatives from different colleges and disciplines defining what AI readiness should entail within their respective fields. The model is also designed to evolve continuously. Feedback from employers, industry collaborators and applied learning environments informs curriculum development, helping keep AI-related learning aligned with changing professional practices.[50]

The European Union provides a complementary regulatory benchmark. For higher-risk educational AI applications, including areas such as admissions, grading and proctoring, the emerging framework places greater emphasis on governance, risk management, accountability, transparency and human oversight. The underlying principle is that the level of institutional control should increase with the potential consequences of AI-supported decisions.[51]

Weaving AI readiness into the Northeastern curriculum Northeastern has moved beyond AI policy and standalone AI courses to establish a faculty-led framework that embeds AI readiness across disciplines, aligns learning outcomes with industry needs and continuously updates curricula through employer and experiential-learning feedback.

49,50. Digital Education Plan, 2026

51. European Commission, Digital Education Council

Source: EY-Parthenon analysis

As AI policy adoption grows across Indian universities, the challenge is shifting from policy creation to institutional implementation

Where India stands

According to the 2025 FICCI EY-Parthenon survey, approximately 57% of Indian HEIs had established an AI policy, while a further 40% were in the process of developing one. However, implementation remained uneven, with adoption often occurring through fragmented initiatives rather than coordinated institutional strategies.

Symbiosis International University provides an example of a more developed approach. Its formal AI policy, introduced in 2024 and revised in May 2026, extends AI-related learning across academic programs rather than restricting it to technical disciplines.

> *"Institution-wide technology decisions are evaluated through dedicated committees that consider licensing, access protocols, token allocation and governance requirements before recommendations are submitted to the Executive Board. Program directors retain responsibility for implementation and budgeting within their units, creating a balance between central oversight and decentralized execution."* ***- Dr Raman, Vice-Chancellor, Symbiosis International University***

Primary interviews, however, suggest that this level of institutional maturity is not yet representative of the wider sector. In many universities, AI use continues to depend on individual faculty members, departments or students using personal subscriptions or freely available tools. Even where formal policies exist, course-level guidance, assessment rules and implementation mechanisms may remain inconsistent. Some interviewees reported limited clarity on which tools can be used, for what purposes and under what conditions, while others highlighted the need for practical, use case-based training and workshops, particularly for non-technical departments, to support more consistent adoption.

The comparative gap is therefore increasingly one of **operationalization rather than policy intent**. Leading global institutions are moving from broad AI principles to implementation within courses, assessments and institutional processes. Many Indian HEIs have begun developing the policy layer, but the mechanisms required to translate that policy into consistent practice are still evolving.

Sources: 2025 FICCI EY Parthenon survey , EY-Parthenon analysis

As AI adoption matures, universities are moving beyond tools to architecture, governance and risk management

4.1.2 Technology and data: Reference architecture in practice (SSO, LLMs, data storage units, procurement and risk controls, bias tests, etc.)

Global benchmark

Leading universities increasingly treat AI as an institutional technology capability rather than a collection of individual tools. The benchmark is a secure and integrated technology environment comprising:

- Single Sign On (SSO) and identity management
- Learning Management System (LMS) integration
- Approved and governed AI tools
- Secure data storage and data governance frameworks
- Cloud and high-performance computing access
- Procurement and vendor risk controls
- Usage monitoring and audit trails
- Technical support and platform administration
- Unified student and institutional data platforms

Serious AI adoption therefore requires a reference architecture. At the institutional level, this means that access, identity, integration, data storage, procurement, monitoring and security are designed as parts of the same system rather than managed independently.

The global benchmark is not that every student has access to a large language model. It is that the institution understands which AI tools are being used, under what data conditions, for what academic purposes and with what safeguards.

The London School of Economics: Institutional access with data safeguards

LSE's Claude for Education model provides a useful benchmark because it converts access to generative AI from an individual purchase into an institutional resource available to students, academic staff, professional staff and researchers.

Its guidance also explicitly addresses data protection, noting that student and staff data would not be used to train Anthropic's models. The significance lies not only in providing access but in providing access within defined institutional conditions.[52]

IU International University of Applied Sciences: Integration across the learning journey

IU International University in Germany has collaborated with Microsoft to provide institution-wide AI access and integrate AI capabilities into the student learning journey. Through the IU Copilot School, students use Microsoft Copilot alongside Syntea, its AI-powered teaching assistant built on Azure OpenAI. AI-supported learning is integrated across disciplines through personalized support, guided study and critical reflection exercises rather than being confined to technical programs.[53]

University of Sydney: Institutionally governed AI infrastructure

The University of Sydney developed Cogniti, an AI assistant hosted on its secure Azure environment, in response to both pedagogical opportunities and concerns around data privacy, security and intellectual property.

The platform allows educators to create customized AI assistants for specific learning contexts while keeping prompts and responses within a governed institutional environment and preventing their use in training external models.

> *"[Faculty aren't] being replaced by technology; their expertise is reflected in the way that it works. Cogniti provides the framework a teacher needs... so that they can strengthen their relationships with students. We want Cogniti to be community developed: built by educators for educators."* ***– Adam Bridgman Pro Vice Chancellor of Education Innovation, University of Sydney***

52. LSE
53. IU International University of Applied Sciences

Source: EY-Parthenon analysis

Leading institutions are investing in the capabilities and support systems required for effective AI adoption

Where India stands

The 2025 FICCI–EY-Parthenon report identifies digital infrastructure, LMS adoption, cloud capability, data systems, high compute technologies, AI tools in IT services and unified data platforms as important dimensions of institutional readiness. However, across much of the Indian higher education sector, AI adoption remains more application-led than architecture-led.

There are notable examples of more structured approaches:

- At **BITS Goa,** open-source AI models have been deployed internally to support selected administrative and query-based use cases. The Placement Department funds and manages access to Claude for specific operational and student-support requirements, while BITS provides institution-wide access to LLM for students, faculty and staff.
- At **Symbiosis International University,** budgets for AI initiatives are generally managed at the program level. For institution-wide deployments, a dedicated committee evaluates licensing options, user access, expected token consumption and governance requirements before submitting recommendations to the Executive Board. Once approved, implementation is managed by the relevant program or functional directors.

These examples show that institutionally governed AI environments are beginning to emerge in India. Primary interviews, however, suggest that comparable arrangements remain concentrated among a relatively small number of better-resourced institutions.

The comparative gap is therefore not primarily one of **access to AI tools**. Students at Indian and global universities may often be using the same underlying models. The difference lies in the institutional layer around those tools: identity management, integration, secure data environments, procurement controls, technical support and governance.

Leading global universities are increasingly building these capabilities as part of a coherent institutional architecture. Across much of Indian higher education, access to AI tools has advanced faster than the systems required to integrate and govern them consistently. The national infrastructure and funding implications of this gap are addressed separately in Chapter 6.

Sources: 2025 FICCI EY Parthenon survey , EY-Parthenon analysis

As AI adoption grows, leading universities are increasingly moving beyond standalone faculty training programs

4.1.3 Faculty capability: Upskilling programs, templates, rubrics

Global benchmark

The global benchmark is not simply faculty training. It is the establishment of a broader support system that enables educators to translate AI awareness into changes in teaching, course design and assessment. Leading universities increasingly provide faculty with practical resources such as course-policy templates, assignment redesign examples, assessment rubrics, student disclosure guidelines, examples of acceptable and unacceptable AI use, AI- resilient assessment models and ongoing instructional design support. While training can help faculty understand what AI tools are and how they work, it does not necessarily equip them to decide how a particular assignment should change, what constitutes appropriate AI use within a discipline or how learning should be assessed when students have access to generative AI.

Stanford Teaching Commons	Stanford's AI teaching resources and course-policy guidance provide faculty with practical support for developing course-specific approaches to AI rather than relying solely on broad institutional rules. Stanford's Center for Teaching and Learning provides instructional guidance that helps faculty translate institution-level principles into decisions about syllabi, assignments, classroom practices and assessment. The value of this model is therefore not simply that faculty are trained in AI. It is that instructors are given resources that reduce the burden of translating general institutional guidance into individual course design.
MIT Sloan AI Hub	MIT Sloan's AI Hub frames faculty support around how generative AI can augment teaching and enrich learning, with resources that are updated as the field evolves. The Teaching Commons resource hub offers faculty a range of practical resources for AI-enhanced teaching and learning. These include guidance on the use of custom GPTs, AI-supported instructional practices and the design of AI-resilient learning experiences.

At a system level, China is also developing national teacher AI literacy standards, tiered training pathways, assessments and integration of AI knowledge into teacher certification processes.

Sources: MIT Sloan AI Hub, China Government , EY-Parthenon analysis

Faculty capability remains the focus, but institutional support is the missing link in AI adoption

Where India stands

Faculty training has emerged as one of the most common institutional responses to AI adoption in India. The 2025 FICCI–EY-Parthenon report repeatedly identifies faculty upskilling, AI literacy and capability building as important priorities for HEIs.

Indian universities are responding through faculty development programs and workshops. Symbiosis International University, for example, conducts five to six AI-focused faculty-development programs each month through the Symbiosis Teaching Learning Resource Centre. Vice-Chancellor Dr Raman also leads workshops on emerging AI applications, including workflow automation using tools such as Claude.

The majority of primary interviews, however, indicate that AI-related workshops are often designed for a broad institutional audience and delivered uniformly across disciplines. As a result, faculty members often receive general introductions to AI tools and concepts rather than training tailored to their specific teaching contexts, disciplinary requirements or assessment challenges.

While such institution-wide programs help raise awareness and build baseline familiarity with AI, they may be insufficient to support meaningful pedagogical transformation. Effective AI integration requires faculty to understand not only how AI tools function but also how they can be applied within their respective disciplines, learning outcomes and assessment frameworks.

What AI guidelines in universities commonly focus on:

- How faculty should use AI tools
- Academic integrity
- Classroom policies
- Acceptable use of ChatGPT and AI writing tools
- Student assessment practices



What AI guidelines in universities should focus on:

- Course level AI policy templates
- Discipline specific examples and use cases
- Assignment redesign frameworks
- Sample rubrics and disclosure formats
- AI resilient assessment models
- Ongoing instructional design support

The comparative gap is therefore not simply one of **faculty AI literacy**. India is already investing in faculty awareness and training. The more important difference lies in the extent to which institutions provide sustained, discipline-specific support that helps faculty convert that awareness into changes in teaching and assessment.

For Tier 2 and Tier 3 institutions, replicating dedicated AI teaching centers or specialist support teams may not be practical. The relevant global lesson is therefore not that every institution needs to reproduce the same organizational structures, but that faculty require access to the underlying capability. Shared templates, discipline-specific toolkits, common assessment resources and central or regional instructional design support can provide similar benefits without requiring every institution to build these systems independently. The policy mechanisms required to extend such support at scale are addressed in Chapter 6.

Sources: 2025 FICCI EY Parthenon survey ,EY-Parthenon analysis

As AI expands research output, leading universities are shifting focus from publications to real-world impact

4.1.4 Research and innovation capacity

Global benchmark

Universities are increasingly measuring AI research success not by publication volume alone, but by their ability to connect research with real-world impact. Globally, AI research is becoming more organized around applied challenges in areas such as healthcare, education, climate, finance, manufacturing and public policy. Universities are creating structures that bring together computer scientists, subject matter specialists, industry partners, entrepreneurs, and students to accelerate research, application and commercialization.

The emerging benchmark is therefore not simply the number of AI research papers produced, but the institution's ability to:

- Conduct interdisciplinary AI research
- Secure industry-funded research partnerships
- Translate research into products, services and deployable solutions
- Support AI start-ups and entrepreneurship
- Build student research and innovation pipelines
- Establish domain-specific AI centers and laboratories
- Generate measurable societal and economic impact from AI research

University of Oxford

Oxford has increasingly positioned its AI research around interdisciplinary collaboration and real-world applications. Newer partnerships, including the Oxford–UBS Centre for Applied AI, focus on the application of AI across business, society, governance, sustainability and economic systems. Oxford also supports AI centers and doctoral programs that combine academic research with industry-linked projects and pathways toward commercialization. The significance of this model lies in the connection between research capability and translation. Research is not treated as an isolated academic output, but as part of a broader ecosystem linking disciplines, industry, entrepreneurship and implementation.[54]

Where India stands

India has made significant progress in research output over the past decade, and several leading institutions have developed substantial AI research capabilities. Premier universities and institutes have established specialized AI centers, industry collaborations and advanced research programs, including the IIT Madras covered in the previous section.

Other institutions are also building specialized capacity. Thapar Institute of Engineering and Technology, for example, has collaborated with a leading global technology company to establish the Thapar School of Advanced AI and Data Science. The initiative combines specialized AI and machine-learning programs, AI-literacy micro-credentials for students across disciplines and advanced computing infrastructure with more than 227 petaflops of capacity, planned to expand beyond 500 petaflops. [55]

A different model is emerging at FLAME University, where the Centre for Inter-disciplinary Artificial Intelligence explicitly supports faculty from outside computer science to undertake AI-enabled research. The Centre invites projects across areas such as humanities, social sciences, Indian languages, music and art and supports selected proposals with computational resources, cloud and API access, software developers and interns. The model is notable because it lowers the technical and infrastructure barriers for domain experts to participate in AI research, rather than concentrating capability within computing departments alone.

The challenge is therefore not the absence of advanced AI research in India; it is its concentration. Deep research infrastructure, interdisciplinary centers, industry-linked programs and commercialization pathways remain concentrated within a relatively small group of institutions, while capabilities across the wider higher education ecosystem are considerably more uneven.

The comparative gap is consequently one of **research translation and diffusion rather than research activity alone**. Leading global universities increasingly connect AI research to industry, entrepreneurship, student innovation and deployment. India has strong pockets of comparable capability, but these models have not yet become widespread across the higher education system. The relevant benchmark for Indian HEIs should therefore extend beyond publication counts to the extent to which research generates meaningful academic, economic and societal outcomes.

54. University of Oxford
55. Thapar Institute of Engineering and Technology

Source: EY-Parthenon analysis

Leading institutions are redesigning programs around how AI changes professional practice

4.1.5 Curriculum, pedagogy and assessment comparison

Global benchmark

The emerging global benchmark is **discipline-specific AI integration combined with assessment that remains credible in an AI-enabled learning environment.**

Leading institutions are increasingly moving beyond standalone AI-literacy modules and asking how AI changes the knowledge, workflows and capabilities required within individual disciplines. Curriculum design and assessment are then adapted accordingly.

- **Wharton's AI for Business major** where AI is taught in the context of business problems and managerial decision-making, illustrates how AI can be embedded within professional education rather than treated as an isolated technical subject.
- MIT Sloan's guidance on resilient learning similarly reflects a shift away from treating AI primarily as a tool to be prohibited. Instead, it focuses on designing learning experiences and assessments that remain meaningful in an environment where students already have access to generative AI.
- Singapore's guided approaches to AI-enabled learning reflect a similar principle: AI use is structured around pedagogical objectives and guardrails rather than treated as unrestricted access to technology.[56]

Across these examples, the benchmark is not simply whether AI appears within the curriculum. It is whether institutions have reconsidered **what students need to learn, how AI changes disciplinary practice and how that learning can still be assessed credibly.**

Where India stands

Primary interviews indicate a significant variation in both the depth and reach of AI integration across Indian higher education. AI is becoming increasingly visible within engineering, management and several leading private institutions, but integration remains less developed across many affiliated colleges, non-STEM disciplines and resource-constrained institutions.

Some institutions are already beginning to move beyond generic AI adoption towards program- and course-level integration. UPES, for example, describes an "AI-first" approach in which AI is embedded across all seven schools and applied differently by discipline: law students test AI-generated briefs against statutes and case law, business students critique AI-generated strategies using real data, health-sciences students evaluate AI-supported diagnoses against clinical evidence and design students use generative AI while retaining responsibility for authorship. Its approach also extends into assessment, with greater emphasis on interrogating AI outputs and demonstrating reasoning, judgment and verification. At the course level, business schools such as IMI (International Management Institute) have introduced requirements for faculty to identify where AI-enabled learning is relevant within individual subjects, expose students to function-specific AI tools, incorporate cases on responsible use and address issues such as bias, limitations and appropriate application. Together, these examples illustrate how some leading institutions are beginning to translate AI from a generic institutional capability into discipline- and course-specific learning.

However, such approaches remain far from universal. Evidence from primary interviews points to uneven exposure to AI-enabled learning. For example, a painting professor at an institution reported limited familiarity with relevant AI applications and consequently little integration of AI into classroom teaching. The example illustrates a broader disciplinary issue: the relevance of AI may be recognized at an institutional level without yet being translated into what AI means for teaching and learning within individual fields.

56. Government of Singapore

Source: EY-Parthenon analysis

As AI usage becomes widespread, the focus is shifting from AI access to AI literacy

Indian universities are also beginning to adapt assessment practices, but implementation remains inconsistent. Students may encounter different rules across courses, while faculty vary in how they interpret AI use within assignments. The issue here is not simply the policy-to-practice gap discussed in Section 4.1.1, but its pedagogical consequence: the same institution may lack a coherent approach to determining how AI should alter learning outcomes and assessment design within different disciplines.

The comparative gap is therefore one of **generic AI integration versus program-level redesign**. Leading global institutions are increasingly asking how AI changes professional and disciplinary practice and redesigning curricula and assessments around that question. Across much of Indian higher education, AI integration remains more commonly layered onto existing programs than used as a basis for systematically rethinking them.

The next stage of maturity would therefore require program-level responses across disciplines such as law, finance, medicine, education, social sciences, design, architecture, public policy, management and others, rather than a uniform model of AI awareness across the institution.

4.1.6 Student preparedness

Global benchmark

The emerging global benchmark is not whether students use AI, but whether they can use it competently and responsibly. Student preparedness increasingly encompasses the ability to understand the capabilities and limitations of AI systems, apply AI appropriately, evaluate outputs critically, identify errors and bias, disclose its use transparently and defend AI-assisted work through independent reasoning and judgment.

Yale University

Yale has developed a Generative AI Literacy Framework that defines capabilities students should acquire when working with AI. These include understanding how AI systems function, recognizing their strengths and weaknesses, evaluating reliability and data quality and making informed judgments about AI-generated information. Illustrative practices include AI literacy learning outcomes, evaluation of AI reliability and credibility, understanding model limitations and biases and critical review of AI-generated content.

The framework therefore treats AI literacy as a cognitive and judgment capability rather than simply proficiency with a particular tool.[57]

Where India stands

Primary interviews indicate that Indian students are already active users of generative AI. Students described using AI as a regular collaborator for learning, problem solving, content creation, concept clarification, idea generation and rapid exploration of unfamiliar topics.

There are also signs of more sophisticated usage. Some students described using AI to validate their thinking, explore alternative perspectives, break complex problems into manageable components and accelerate work under tight timelines. Several emphasized the importance of checking outputs and retaining independent judgment.

However, usage patterns remain highly uneven. Some students engage with AI as a thought partner and critically interrogate its responses, while others use it primarily to automate routine tasks or generate finished outputs. As discussed in Chapter 2, heavy adoption can therefore coexist with uncertainty over verification, appropriate use and the boundary between assistance and substitution.

The comparative gap is consequently not one of **AI adoption**, where Indian students may already compare favorably with global peers. It is the extent to which institutions systematically convert that adoption into **demonstrable AI literacy**.

The benchmark should therefore be what students can **demonstrate**, rather than simply whether they have access to or frequently use AI tools.

---

57. Generative AI Literacy Framework

Source: EY-Parthenon analysis

As AI reshapes work, industry partnerships are shifting from placement support to capability co-production

4.1.7 Industry alignment

Global benchmark

The emerging global benchmark is not employer engagement at the point of recruitment, but **capability co-production throughout the student learning journey**. Leading universities increasingly involve industry partners in the development of curricula, AI-focused courses, applied projects, innovation labs, research, internships, capstone experiences and other forms of work-integrated learning. The objective is not to allow employers to dictate academic programs, but to help keep students exposed to contemporary professional workflows and the capabilities developed within universities remain relevant to changing workplace requirements.

Imperial College London

Imperial's AI programs incorporate industry engagement through group projects, applied innovation modules, entrepreneurship-focused coursework and opportunities for industrial internships. Several programs require students to develop AI products and accompanying business cases, often with industrial input, reflecting a deliberate effort to align learning outcomes with employer and innovation ecosystem needs.[58]

Where India stands

The 2025 FICCI-EY-Parthenon AI Adoption Survey indicates growing recognition of the importance of industry collaboration in AI education. Institutions are increasingly introducing AI through degree programs, minors, interdisciplinary modules and industry-linked learning experiences. The report also highlights the need for stronger ecosystems connecting academia, industry and government to support AI adoption and workforce readiness.

India therefore does not lack industry engagement. The more important question is the **depth and continuity** of that engagement.

Several Indian institutions are beginning to move beyond traditional placement-oriented industry engagement toward deeper collaboration in AI capability development. The Robert Bosch Centre for Data Science and Artificial Intelligence (RBCDSAI) at IIT Madras integrates industry supported research, student projects and innovation activities across multiple domains. IIT Delhi's Corporate Relations and industry partnership initiatives have similarly expanded opportunities for industry sponsored projects, internships and collaborative research in emerging technologies. Institutions such as BITS Pilani have strengthened industry engagement through practice schools, innovation centers, startup ecosystems and partnerships with technology companies that expose students to real- world problem solving. Across management and engineering institutions, industry participation is increasingly visible through AI focused labs, hackathons, capstone projects, challenge-based learning and curriculum inputs from employers.[59] [60] [61]

Taken together, these examples suggest that India is making meaningful progress toward stronger industry academia collaboration in AI. However, such initiatives remain concentrated within leading institutions and centers of excellence. Across much of the higher education system, employer engagement continues to be focused on internships, guest lectures, placements, advisory participation and selected projects rather than systematic co-development of capabilities throughout the student learning journey.

As discussed in Chapter 3, traditional HEI industry relationships frequently remain concentrated around internships, guest lectures, placements, advisory participation and selected projects. These activities remain valuable, but they do not necessarily mean that industry is systematically contributing to how capabilities are defined, developed and validated across the student learning journey.

The comparative gap is therefore one of **engagement versus co-production**. Leading global models increasingly integrate industry into applied learning and capability development, while much of Indian higher education still engages employers more intensively at the later stages of the student journey.

The benchmark should not simply be whether employers recruit from an institution or participate in its activities. It should be whether industry engagement improves the relevance and authenticity of students' learning experiences and helps validate the capabilities required in an AI-enabled workplace.

58. Imperial College London,
59. RBCDSAI) at IIT Madras
60. IIT Delhi's Corporate Relations ,
61. BITS Pilani

Source: EY-Parthenon analysis

Global AI models provide direction, but India need to adapt them to its own realities

4.2 Why global models do not "copy-paste" into India and what can

Global benchmarks provide valuable direction for Indian higher education, but they cannot be transferred without accounting for the institutional conditions in which they were developed. Many leading global universities operate with conditions that make institution wide AI integration relatively easier, established teaching and learning centers, flexible curriculum processes, dedicated instructional design teams, strong digital infrastructure, substantial technology budgets, smaller class sizes, robust student support systems and greater institutional autonomy. These capabilities allow institutions to move more quickly from experimentation to coordinated implementation and governance.[62]

India operates under materially different conditions. The system encompasses institutions with widely varying levels of funding, autonomy, infrastructure, faculty capacity and digital maturity. The relevant question is therefore not whether India should replicate global models, but **which underlying principles are transferable and how they need to be adapted to Indian conditions**.

4.2.1 Cultural, financial and structural differences

Several structural features shape the feasibility of AI adoption in Indian higher education:

- **Large class sizes and mass enrolment** make high touch models of individualized support, feedback and assessment redesign more difficult to implement at scale.
- **Affiliated college structures** can slow curriculum reform and create significant variation between policy decisions made centrally and implementation at the college level.
- **Uneven digital infrastructure** remains an important constraint, particularly across state universities and rural institutions, where access to connectivity, devices, technical support and computing resources varies considerably.
- **Language diversity and differences in digital fluency** require AI systems and learning models that can operate across regional languages and varying levels of technological familiarity.
- **Resource and procurement constraints** limit the ability of many institutions to acquire premium tools, advanced computing infrastructure or specialist technical and instructional-design support.

These differences mean that some practices visible at elite global universities are unlikely to be directly scalable across India.

Universal access to premium AI platforms, for example, may be feasible within institutions such as LSE or Stanford but considerably more difficult across a system containing thousands of resource constrained colleges. Similarly, small-class teaching models that rely on frequent personalized feedback or intensive assessment redesign may be difficult to reproduce where faculty-to-student ratios are substantially higher.

Institutional autonomy also matters. Faculty at universities such as Stanford may have considerable discretion to establish course-specific AI policies and redesign assessments. In affiliated college systems, curricula and examination structures are often determined at the university or regulatory level, reducing the scope for individual instructors to make comparable changes independently.

Advanced compute infrastructure presents a similar challenge. The principle of providing secure, institutionally governed access to AI is transferable; expecting every HEI to build a sophisticated enterprise AI environment is not.

The relevant lesson from global benchmarks is therefore to separate the **principle from the institutional form through which it is delivered**.

62. Times Higher Education

Source: EY-Parthenon analysis

Global AI benchmarks are transferable, but their implementation need to reflect Indian realities

Several principles can transfer effectively:

- **Program level curriculum redesign:** AI can be integrated within disciplines such as law, finance, healthcare, education, public policy and management rather than introduced only through generic AI literacy courses.
- **Course level AI guidance:** Students can receive clear expectations on permitted AI use within individual courses and assessments without requiring significant new infrastructure.
- **Practical faculty support:** Templates, assessment frameworks, sample rubrics and discipline-specific resources can complement faculty training without requiring every institution to maintain a dedicated teaching innovation center.
- **Core governance principles:** Data protection, risk-based oversight, procurement discipline, human accountability and periodic review remain relevant irrespective of institutional size.
- **Shared infrastructure:** Regional platforms, common cloud resources, shared AI services and centrally developed support systems can substitute for institution-by-institution infrastructure investment.

The objective should therefore be **adaptation rather than imitation**: preserving the principles that make global models effective while redesigning their delivery for Indian institutional realities.

4.2.2 India's scale and diversity problem

Scale creates a second challenge that is distinct from the institutional constraints faced by any individual university.

Many global benchmarks are drawn from a relatively small number of autonomous and well-resourced institutions. Even where those models are highly effective, they do not demonstrate how the same approach would operate across tens of thousands of institutions with sharply different starting conditions.

India's higher education system includes globally competitive research universities, large private institutions, state universities, autonomous colleges, affiliated colleges, specialist institutions, rural campuses and institutions serving predominantly first-generation or regional-language learners. A single model of AI readiness is therefore unlikely to be appropriate across the entire system.

A Tier 1 research university may need access to high-performance compute, advanced model-development capabilities, secure institutional AI environments and specialized research infrastructure. A Tier 2 or Tier 3 college may derive substantially greater immediate value from reliable connectivity, secure access to approved AI tools, faculty support resources, multilingual applications and clear assessment guidance.

The implication is not that different tiers should face different standards of responsible use. Core expectations around student protection, academic integrity, data governance and human accountability should remain consistent. What should differ is the **pathway and level of capability through which those standards are achieved**.

This distinction is important because uniform requirements can themselves create inequity. A standard designed around the capabilities of elite institutions may become impossible for smaller colleges to implement, while a standard designed around the lowest common denominator could constrain institutions capable of more advanced innovation.

India therefore requires a differentiated model of AI readiness: a common minimum foundation across all recognized HEIs, with progressively more advanced capabilities for institutions with greater resources, research mandates and technical capacity.

This also changes how global benchmarking should be interpreted. The relevant comparison is not whether every Indian HEI looks like Stanford, MIT or LSE. It is whether the Indian system can make the **capabilities underlying those models accessible in forms appropriate to different institutional contexts**.

Source: EY-Parthenon analysis

The goal is not to match global leaders, but to build AI systems that work across India

4.2.3 Regulatory and accreditation constraints

Global experience also illustrates a movement toward more structured AI governance, particularly for applications that can materially affect student rights, outcomes or opportunities. International frameworks increasingly emphasize proportional oversight, transparency, accountability, human review and stronger safeguards as the consequences of an AI supported decision increase.

Stanford established an AI at Stanford Advisory Committee to evaluate AI use across education, research and administration, identify policy gaps and recommend university wide principles. The framework emphasizes human oversight, responsible experimentation, ethical use, transparency and continuous policy review. Governance responsibility is distributed across relevant university bodies, creating an ongoing mechanism for adapting policies as AI evolves.[63]

India faces an additional challenge because higher education institutions operate within a multilayered regulatory environment. Depending on their programs and institutional form, HEIs may be subject to requirements from the UGC, AICTE and other professional councils, state governments, affiliating universities, procurement authorities, accreditation bodies and emerging data protection requirements.

The difficulty is therefore not simply the absence of AI regulation. It is that institutions need to integrate AI governance into an already complex system of overlapping academic, professional, administrative and legal responsibilities.

This has two implications:

- First, global governance frameworks cannot simply be imported as an additional layer of compliance. If AI requirements are developed independently by multiple regulators or professional bodies, institutions may face conflicting definitions, duplicated reporting obligations and inconsistent expectations across programs.
- Second, the appropriate balance between national consistency and institutional discretion need to reflect the structure of Indian higher education. Certain principles, including data protection, transparency, student rights, human accountability and governance of consequential AI applications, require consistency. By contrast, decisions about how AI should be used within a particular discipline, course or assessment often require greater pedagogical flexibility.

The global lesson that transfers is therefore the **principle of proportional and accountable governance**, rather than any single regulatory model.

For India, the challenge is to create sufficient consistency across regulators and institutional types without imposing detailed rules that become obsolete as technologies evolve or prevent legitimate disciplinary variation.

The detailed design of risk classification, institutional accountability, data protection and national policy architecture is addressed in Chapters 5 and 6. The comparative lesson here is narrower: global models work because governance responsibilities are increasingly being made explicit, India need to achieve the same clarity within a substantially more complex regulatory and institutional system.

63. Stanford University

Source: EY-Parthenon analysis

Healthcare shows that AI scales through infrastructure and governance, not technology alone

4.3 What education AI can learn from healthcare AI in India?

India's healthcare sector provides a useful domestic comparison for higher education because it demonstrates how digital transformation can be built at a national scale across a large and heterogeneous system. The most relevant lesson is not the adoption of individual AI applications, but the development of common digital infrastructure, standards, governance mechanisms and implementation capacity that allow technology to scale across institutions.

Higher education already has several comparable digital building blocks. The question is whether these can evolve from largely separate systems into a coordinated architecture capable of supporting AI-enabled learning, student services, research and institutional decision-making.

4.3.1 Comparative adoption levels

Healthcare benchmark

India's healthcare sector has developed one of the country's strongest examples of digital public infrastructure through the Ayushman Bharat Digital Mission (ABDM). Rather than building a collection of standalone applications, ABDM established a common architecture based on Ayushman Bharat Health Account (ABHA) digital identities, health facility and healthcare professional registries, consent-based data exchange and interoperable digital health records.[64]

The scale of adoption is significant. In May 2026, the Ministry of Health and Family Welfare reported that more than 100 crore health records had been linked to ABHA accounts, supported by over 450 integrated public and private health technology solutions.[65]

The relevance of ABDM for higher education lies in its architecture. Different applications and providers can operate within a shared set of identity, interoperability, consent and data-exchange mechanisms rather than each building these foundations independently.

Education comparison

Higher education already possesses several important digital building blocks, including APAAR IDs, Academic Bank of Credits (ABC), DigiLocker, the National Academic Depository, learning management systems, ERP platforms, examination systems, placement platforms and digital credential infrastructure.

However, these components do not yet operate as a comparable, integrated architecture for AI-enabled education. Student records, learning data, assessment systems, advising platforms, employability systems and institutional data repositories often remain distributed across separate platforms and institutions.

The distinction is therefore not that healthcare has digital infrastructure while education does not. Both sectors have substantial digital assets. The difference lies in the extent to which those assets operate through common standards for identity, interoperability, consent and data exchange.

This becomes increasingly important as AI applications move beyond generic chatbots towards systems that draw on institutional and student data to support learning, advising, assessment, or administrative decisions. Such applications require reliable data flows and clear rules governing how information can be accessed, shared and used.

The healthcare comparison therefore suggests that education's next stage of digital maturity lies in connecting its existing building blocks into a more interoperable and governable architecture, rather than continuing to add standalone applications.

---

64. Ayushman Bharat Digital Mission
65. PIB APAAR IDs, Academic Bank of Credits (ABC)

Source: EY-Parthenon analysis

Policy momentum is growing, but education still lacks the architecture needed to scale AI

4.3.2 Drivers of adoption: policy, funding, regulation and infrastructure

Healthcare benchmark

India's digital health experience demonstrates that large-scale technology adoption is not driven by an interest in technology alone. ABDM combined national policy direction with common technical standards, digital identity, consent mechanisms, registries, interoperability requirements, governance structures and supporting infrastructure.

This created an implementation layer through which public and private providers could participate within a common ecosystem. The sequencing matters. Shared infrastructure and standards created the conditions through which applications could subsequently scale, rather than requiring every hospital, provider, or technology company to solve the same foundational problems independently.

Education comparison

Higher education is experiencing comparable policy momentum. The IndiaAI Mission, AI Centres of Excellence, AI-focused curriculum initiatives, faculty-development programs, institutional pilots and other national and institutional efforts have created significant momentum across the sector.

The 2025 FICCI–EY-Parthenon AI Adoption Survey similarly indicates that HEIs increasingly view AI as a strategic priority and are experimenting with applications across teaching, learning, research and administration.

The difference lies less in ambition than in implementation architecture.

Higher education does not yet have the same degree of coordination around the institutional processes required to scale adoption across a highly fragmented system. Questions around common infrastructure, approved tools, data exchange, procurement, institutional responsibilities, recurring funding and implementation support remain less consistently resolved.

This helps explain why a large number of pilots or institutional initiatives do not necessarily translate into system-level adoption. Without common enabling mechanisms, institutions are required to solve many of the same implementation problems independently, producing a significant variation in capability and quality.

The healthcare comparison therefore highlights a broader principle: **national missions create momentum, but implementation architecture converts momentum into scale**.

The detailed policy, funding and governance mechanisms required for higher education are addressed in Chapters 5 and 6. The comparative lesson here is that these elements must operate as a coordinated system rather than as parallel initiatives.

4.3.3 Equity: Access is necessary but not sufficient

Healthcare benchmark

India's digital-health experience also demonstrates that access to digital infrastructure does not automatically produce equitable usage or outcomes.

ABDM has created a common infrastructure through digital identities, interoperable records, consent-based exchange and connected platforms. Yet, the ability of individuals and institutions to benefit from these systems still depends on factors such as awareness, digital literacy, trust, language accessibility, institutional readiness and last-mile implementation capacity. Digital infrastructure can therefore remove one category of constraint while leaving others intact.

Sources: 2025 FICCI EY Parthenon survey , EY-Parthenon analysis

AI equity depends on the learning ecosystem, not just tool availability

Education comparison

The same distinction applies to AI in higher education. Access to an AI tool is not equivalent to the ability to use it productively. Students' ability to benefit from AI depends on a wider learning environment that includes device access, connectivity, language proficiency, digital literacy, faculty guidance, institutional policies and the quality of available tools. Two students with nominal access to generative AI may therefore operate within very different learning environments. A student with a personal laptop, reliable connectivity, institutionally licensed tools, trained faculty and strong English proficiency has access to a substantially richer AI ecosystem than a mobile-first student at an under-resourced affiliated college who relies on free tools, has limited connectivity, receives little formal guidance and encounters language barriers.

The equity challenge is therefore not adequately captured by measuring whether a tool has been made available.

This parallels the healthcare experience: **infrastructure creates access, but complementary capabilities determine whether access produces meaningful outcomes**.
For higher education, these complementary conditions include multilingual access, affordability, device and connectivity support, faculty capability, accessible interfaces and shared institutional support for resource-constrained HEIs. This is particularly important because AI may amplify differences in existing capability. Students with stronger academic foundations and better institutional support may be better positioned to extract value from the same tools than students who require greater guidance. The benchmark for equitable adoption should therefore be meaningful usage and learning outcomes across student groups, rather than access alone.



Successful AI adoption may likely depend less on tools and more on the systems that support them

4.3.4 Lessons for 2026–2030: Towards an Indian Model of AI Readiness

India's healthcare experience offers a clear lesson for higher education: technology scales when the enabling system develops alongside the applications. The Ayushman Bharat Digital Mission was built on shared digital identities, interoperable records, consent mechanisms, governance standards and common infrastructure. Higher education faces a similar challenge. Sustainable AI adoption would depend less on the number of tools introduced than on whether institutions have the underlying systems required to use them effectively and responsibly.

The global benchmarks examined in this chapter reinforce the same conclusion. Leading universities are moving from experimentation towards institutionalization. AI is increasingly embedded within governed technology environments, discipline-specific curricula, assessment practices, faculty support systems, research ecosystems and student capability frameworks. The important lesson for India is not to replicate these institutional models wholesale, but to identify the principles that make them effective.

Four implications stand out:

- **Build common foundations before scaling applications:**
  - Higher education already possesses important digital assets through APAAR, the Academic Bank of Credits, DigiLocker, LMSs, ERP platforms and other systems. The next stage is to make these foundations more interoperable, secure and governable so that AI applications do not develop as another layer of disconnected technology.
  - The healthcare analogy is therefore about sequencing: establish the foundations that allow innovation to scale safely, rather than expecting every institution to solve the same infrastructure and governance problems independently.
- **Measure capability, not adoption:**
  - High AI usage does not necessarily imply AI readiness. The same applies at the institutional level: the existence of an AI policy, faculty workshop, license, laboratory, or Centre of Excellence does not by itself demonstrate meaningful integration.
  - The stronger benchmark is whether AI improves the quality of learning, teaching, research, judgment and institutional decision-making. This shifts the focus from inputs and activity towards demonstrable capability and outcomes.
- **Adapt global principles to Indian institutional diversity:**
  - The conditions under which Stanford, MIT, Oxford, LSE, Wharton, or Imperial operate are not representative of the wider Indian higher education system. India therefore requires differentiated implementation.
  - A research-intensive institution may need advanced compute, model-development capacity and specialized research infrastructure. A resource-constrained college may first require reliable connectivity, governed access to AI tools, multilingual support, faculty resources and clear academic guidance.
  - The objective should be common standards of quality, safety and accountability, without requiring identical institutional models.
- **Design for scale and equity together:**
  - India's scale makes equity a system-design issue. Access to the same AI tool can produce very different outcomes depending on connectivity, devices, language proficiency, academic foundations, faculty support and institutional capacity.
  - The benchmark should therefore not be whether AI has been made available, but whether students and institutions across different contexts can use it meaningfully. National-scale adoption would succeed only if models are designed for the median institution, not only for the frontier.

The broader implication for India

Taken together, the global comparisons and the healthcare experience point towards a distinctly Indian model of AI readiness. Global universities demonstrate what mature institutional capability can look like. India's healthcare transformation demonstrates how common infrastructure, standards and governance can enable technology to scale across a large and heterogeneous system. India's opportunity between 2026 and 2030 is to combine these two lessons: build common foundations that support the wider higher education system while allowing institutions with greater capability to innovate beyond them.

The success of India's AI transition in higher education will therefore not be determined by whether its best universities can imitate global leaders. It would depend on whether the country can create a model that works for the **median Indian student, the median Indian faculty member and the median Indian institution**, while still enabling its leading universities to operate at the global frontier.



Designing the ethical, effective, AI-enabled Indian campus

Chapter 5

While AI is becoming common in classrooms, leading institutions are redefining academic integrity around intellectual ownership and accountability

5.1 What does ethical AI use look like in the classroom?

The central question should not be whether technological tools, including AI, contributed to the production of a student's work. Rather, the key consideration is whether the student remains the accountable intellectual author of the submission. The focus should therefore be on the extent to which the student has exercised independent judgment, critical thinking, and substantive control over the ideas, arguments, analysis and conclusions presented. Where the student retains ownership of the intellectual process and assumes responsibility for the final content, the use of supporting tools does not necessarily diminish authorship. Academic evaluation should therefore prioritize evidence of intellectual responsibility, originality of thought and accountability for the work, rather than the mere presence or absence of technological assistance.

5.1.1 Academic integrity re-imagined

In an educational environment where AI is increasingly integrated into teaching and learning, academic integrity should be grounded in the principle of intellectual ownership rather than technological purity. A student may legitimately use AI while retaining authorship, provided that they define the problem, exercise independent judgment in selecting and interpreting evidence, critically evaluate AI generated outputs, verify factual claims, and remain capable of explaining, defending and taking responsibility for the final submission. In this framework, authorship is determined by the student's control over the intellectual process rather than simply by whether AI was used.

This necessarily means that acceptable AI use will vary across disciplines, courses and assessment types. AI-assisted language editing or grammar correction may be entirely consistent with the objectives of an economics assignment, where the primary focus is economic reasoning. The same level of assistance may be inappropriate in a language or writing course where sentence construction, linguistic proficiency and stylistic development are themselves being assessed. Ethical evaluation may therefore remain sensitive to the purpose of the learning activity rather than rely on universal permissions or prohibitions.

To support this distinction, institutional and classroom guidance should move beyond broad statements asking students to use AI responsibly or warning them not to misuse generative AI. Educators should adopt activity-based policies that specify which forms of AI use are permitted, restricted, or prohibited within a particular course or assessment. Faculty should indicate whether students may use AI for activities such as brainstorming, translation, outlining, source discovery, draft generation, coding, data analysis, editing, or final submission preparation. Stanford's Teaching Commons reflects this approach by emphasizing that appropriate AI use should be defined within specific pedagogical contexts rather than through institutional rules alone.[66]

Transparency should similarly be treated as a matter of scholarly disclosure rather than as an implicit admission of misconduct. Where AI has made a substantial contribution to the structure, analysis, coding, or substantive content of an assignment, students should provide a concise explanation of the tool used, the nature of its contribution, the elements accepted or rejected and the steps taken to verify the output. Disclosure requirements should remain proportionate, routine grammar correction may require no disclosure or only a standard declaration, whereas more substantive assistance may justify a brief process statement.

66. Stanford Teaching Commons

Source: EY-Parthenon analysis

While AI is becoming embedded across higher education, leading institutions are building frameworks for safe, transparent and accountable AI adoption

Institutions should also exercise caution in relying on AI detection software as the primary mechanism for enforcing academic integrity. Such tools generate probabilistic assessments and cannot yet reliably determine the extent of a student's intellectual contribution. Most universities interviewed reported following a benchmark that permits up to 20% AI assisted content in student submissions. However, detection systems cannot establish how ideas were developed, evaluated, or refined throughout the learning process. An excessive focus on detection risks diverting attention from the more important challenge of designing assessments that provide credible evidence of authentic learning, critical thinking and student accountability.

Ethical policy must also recognize differences in students' circumstances. A student who uses AI to translate an independently developed argument from a regional language into English occupies a fundamentally different position from one who uses AI to generate the underlying argument itself. Similarly, students with disabilities may rely on AI-enabled transcription, speech interfaces, text simplification, or image description tools as essential forms of accessibility support.

Responsible AI governance should therefore preserve human agency, inclusion, equity and educational purpose. The future of academic integrity lies neither in prohibiting AI nor in accepting its use without limits, but in developing frameworks that distinguish clearly between **assistance and authorship**.

5.2 What does safe AI use look like institution wide?

A safe institutional model should proceed in a clear sequence, first understand the data being processed, then establish meaningful notice and consent, classify the risk associated with each use case, and test systems for bias and other forms of failure. Procurement, security review, human oversight, incident response and appeal mechanisms should surround this sequence as continuing institutional responsibilities rather than one time compliance exercises.

5.2.1 Data mapping

Safe and responsible AI deployment begins with a clear understanding of how information moves through institutional systems. Before implementing AI-enabled solutions, institutions need visibility into the data they collect, how it is processed, where it is transferred, who can access it and the purposes for which it is ultimately used.

This requires systematic mapping of data flows across major institutional functions, including admissions, attendance, academic records, assessment submissions, LMS activity, financial aid, placements, research, disability support, counseling and other student services. Without a clear understanding of these flows, supported by a framework for classifying information according to sensitivity and risk, institutions cannot effectively assess the obligations associated with AI deployment.

Different categories of information require different levels of protection. Publicly available course materials or routine administrative information may, for example, be suitable for use within an institution specific AI assistant. Identifiable counseling records, unpublished research, disciplinary proceedings, health information or disability-related data require substantially stronger safeguards and access controls.

Source: EY-Parthenon analysis

As AI becomes embedded in higher education, leading institutions are placing data governance, transparency and student autonomy at the core of AI readiness

Institutions must also distinguish between data explicitly provided by students and data generated or inferred by AI systems. Students may knowingly submit information through applications, assessments, surveys, or learning platforms. AI systems, however, can create additional information through predictive and analytical processes, including assessments of academic risk, inferred learning preferences, or predictions of future performance.

The importance of data mapping is reinforced by the DPDP Rules, 2025. Compliance with requirements relating to notice, security safeguards, retention, breach management, and data-principal rights depends on institutions understanding how personal data is collected, processed, stored and shared across their systems.[67]

5.2.2 Consent flows

Consent for AI use cannot be reduced to a single privacy acknowledgement obtained when a student enrolls. AI systems may introduce new processing purposes, generate additional inferences, or involve third-party providers that were not contemplated when data was originally collected. Institutions should therefore provide clear, purpose-specific notices explaining what data will be used, why it is required, whether participation is mandatory or optional, who may receive the information, how long it will be retained and how students can exercise their rights.

Meaningful consent requires a genuine choice. Students cannot reasonably be considered having freely consented to an AI system if refusing to use it prevents them from completing a compulsory academic requirement or accessing an essential institutional service. Where feasible, institutions should therefore provide non-AI alternatives, particularly in areas such as advising, proctoring, career guidance and wellbeing support.

Transparency is equally important; information about AI systems should be communicated in clear and accessible language rather than in technical, legal or vendor-specific terminology, enabling students and faculty to understand what participation involves.

Institutions must also recognize the possibility that a formally optional tool becomes effectively compulsory when faculty or administrators assume universal adoption. Consent should therefore be understood not only as a legal requirement but also as a question of institutional power, student autonomy and academic fairness.

67. Ministry of Electronics and Information Technology

Source: EY-Parthenon analysis

As institutions advance AI integration, risk-based governance is becoming a cornerstone of responsible adoption

5.2.3 Risk categorization of AI tools

AI risk should be classified according to the consequences of the use case rather than the sophistication or popularity of the underlying tool. The same generative model may pose limited risk when used to produce practice questions and substantial risk when used to influence admissions, grading or disciplinary decisions.

Low risk

Low-risk applications generally include voluntary and advisory uses such as brainstorming, grammar assistance, translation, lesson planning and general study support.

These applications should still be subject to basic privacy, security and accuracy safeguards, but typically do not require the same level of oversight as systems that influence consequential academic or administrative outcomes.

Medium risk

Medium-risk applications may include tutoring systems, formative feedback tools, academic advising platforms, career guidance services, learning analytics and administrative triage. Although these systems may not independently determine outcomes, they can materially influence student behavior, decision making and access to institutional support.

They should therefore be subject to documented human oversight, local testing, clearly defined limitations and accessible mechanisms for correcting inaccurate or misleading outputs.

High risk

High risk applications include admissions decisions, scholarship allocation, grading, proctoring, disciplinary investigations, placement shortlisting, dropout risk prediction, biometric identification and mental health screening.

Because these systems can materially affect student rights, academic progression, access to opportunities, reputation and wellbeing, they require stronger governance. Their use should be supported by clear evidence of necessity, proportionality, accuracy, fairness and security, together with meaningful human review and accessible mechanisms for appeal and redress.

As AI becomes embedded in higher education, leading institutions are building trustworthy AI systems and AI-ready classrooms

5.2.4 Bias audits

AI systems deployed in Indian higher education should be evaluated within the environments in which they are intended to operate. Performance on international benchmarks alone does not establish reliability across India's linguistic, cultural, socioeconomic and institutional diversity. A system that performs effectively for English speaking students at a well-resourced university may behave differently for regional language learners or students in other educational settings.

Bias audits should therefore go beyond overall accuracy and examine factual reliability, hallucination rates, citation accuracy, consistency of outputs, performance across Indian languages, accessibility for students with disabilities and potential disparities across relevant student groups. The criteria should also reflect the application, admissions systems should be tested for differential selection outcomes, career guidance tools for stereotypical recommendations and early warning systems for disproportionate false positive rates among particular groups.

Audit findings should be documented transparently rather than reduced to a single fairness score. Institutions should record the groups tested, methodologies applied, performance thresholds used, limitations identified, mitigation measures adopted and the basis on which any residual risk is considered acceptable. Incident registers should similarly capture failures such as inaccurate misconduct flags, misleading guidance, unauthorized data use, or inappropriate automated recommendations, while human oversight should involve genuine independent judgment rather than a routine validation of an automated output.[68]

5.3 The AI-ready classroom starter blueprint

An AI-ready classroom is not defined by the pervasive use of AI across every aspect of teaching and learning. It is characterized by clearly articulated learning outcomes, transparent expectations regarding appropriate AI use, and a culture in which students are encouraged to engage with AI critically and reflectively. Within such an environment, students learn to evaluate, question and verify AI-generated outputs while strengthening the critical thinking and judgment that the technology is intended to support.

Assessment must continue to provide credible evidence of students' knowledge and capabilities. AI should therefore serve educational objectives rather than begin to determine them.

Emerging experimental evidence shows why the design of AI-enabled learning matters. In a randomized field experiment involving nearly 1,000 mathematics students, access to a standard GPT-4-based tutor improved performance on AI-assisted practice problems by 48%. However, once access was removed, these students performed 17% worse than students who had never received AI support. A redesigned tutor that provided structured, teacher-informed hints rather than direct answers largely mitigated this effect.

Conversely, a 2025 randomized controlled trial involving 194 undergraduate physics students found that a deliberately scaffolded AI tutor generated significantly higher learning gains than the comparison active-learning classroom, while requiring less student time. The tutor was designed around active engagement, scaffolding, targeted feedback, cognitive-load management and self-paced learning.

The implication is that AI itself is neither inherently beneficial nor harmful to learning. What matters is whether its use extends students' thinking or substitutes for it. Institutions should therefore design and evaluate AI-enabled learning around cognitive engagement, capability development and learning outcomes, rather than adoption or usage alone.

> *"Assignments were highly valuable in the past, but we are now relying more on viva voce examinations and group discussions. In a course on AI and International Relations, a significant portion of the assessment weightage was allocated to group discussions. I also incorporated surprise assignments and GDs as part of the evaluation process."* **- *Faculty Member, Private University in a tier-2 city***

68. NIST; MIT Sloan Teaching & Learning Technologies

Source: EY-Parthenon analysis

Reimagining assessment, rubrics and academic integrity to balance AI-enabled learning with independent student capability

5.3.1 Templates, rubrics and question frameworks

Each course should include a concise, course-specific statement that identifies permitted AI use, activities requiring disclosure or prior approval, prohibited uses, and, where relevant, approved or discouraged tools. Activity-based examples are particularly useful because students encounter AI through specific tasks. A statement such as, *"You may use AI to generate practice questions, but not to complete answers for the graded problem set,"* offers considerably more guidance than a general instruction to use AI responsibly.

The design of major assessments should similarly begin with the learning objective. Faculty should identify the knowledge, skill, or competency being assessed, consider which elements of the task AI can already perform, determine whether its use would support or compromise the intended learning outcome and establish what evidence students must provide to demonstrate independent understanding.

This supports a four-part classification of assessments: **AI free**, where unaided performance is integral to the capability being measured, **AI assisted**, where limited support such as brainstorming, language improvement, or formative feedback is permitted while the core intellectual work remains the student's, **AI integrated**, where students are expected to use AI but also critically evaluate its outputs and limitations and **AI directed**, where the task assesses the ability to design, manage, supervise, or evaluate AI-enabled processes.

Rubrics need to evolve accordingly. Traditional assessment criteria often reward fluency, polished presentation and broad comprehensiveness, all of which AI can now produce with relative ease. Greater weight should therefore be placed on problem framing, quality of evidence, critical interpretation, recognition of uncertainty, verification, contextual application, evaluation of tradeoffs and the ability to justify and defend conclusions.

5.3.2 AI-friendly and AI-resistant assessment designs

AI friendly assessments should incorporate AI where its use strengthens learning or reflects contemporary professional practice. Students may be asked to compare AI-generated responses with primary evidence, identify inaccuracies or hallucinations, refine weak outputs, test alternative hypotheses, debug generated code, adapt outputs to particular contexts, or explain why an AI-generated recommendation should be rejected. The objective is not to reward prompting proficiency, but to develop the ability to supervise, evaluate and apply AI using disciplinary knowledge and professional judgment.

This shift was strongly reflected in primary interviews. Several faculty members described moving assessment away from content generation towards the correction, validation and improvement of AI-generated outputs, asking students to identify where AI is wrong, verify information, detect errors and explain how an answer should be improved. In technical disciplines, some faculty members have modified coding assessments by withholding test cases or introducing corner and edge cases, thereby testing whether students understand the underlying logic rather than merely submitting generated code. Others use scenario-based questions that require students to explain *who, why and what*, placing greater emphasis on judgment than on reproduction.

Some institutions have also introduced greater transparency into AI-enabled assignments. In open-book or take-home assessments, students may be required to submit prompts or relevant chat histories, allowing faculty to examine the reasoning and verification process alongside the final output. Across these approaches, the principle is consistent: AI may assist the work, but students remain responsible for understanding the logic, validating the output and defending the decisions reflected in the final submission.

Source: EY-Parthenon analysis

As AI use increases, universities are shifting from output-based assessment to evidence of understanding

IISc provides an instructive example of how AI rules can differ across tasks and courses rather than being applied uniformly. In a 2026 *Introduction to Cloud Computing* course at the Department of Computational and Data Sciences, students are expected to work with designated AI agents during a capstone project and submit the resulting interactions, while distinguishing between contributions made by team members and by AI. Quizzes and examinations, however, are conducted independently in closed-book, proctored settings and students must also defend their project work through a viva.

A second course, *Scalable Systems for Data Science*, draws the boundary differently. Students are encouraged to use chatbots as part of the learning process, but generative AI is prohibited from producing assessed project code, reports or presentations. Taken together, the two courses show that assessment design can distinguish between contexts where AI use is part of the skill being developed and contexts where students are expected to demonstrate their own command of the subject.

Similar experimentation is emerging at Ashoka University, where faculty has explored both AI-integrated and AI-resistant approaches to assessment. Examples discussed in a university-wide faculty workshop included requiring students to submit prompts, compare AI-assisted and unaided work, or critically evaluate AI-generated responses, alongside greater use of handwritten assignments, oral examinations, in-class writing and presentations where independent performance needs to be demonstrated. The emerging principle is not to apply a uniform rule to AI, but to align its use with the capability that an assessment is intended to measure.

AI-resistant assessment should not be understood as an attempt to design tasks that AI can never complete, as capabilities improve, that objective will become increasingly difficult to sustain. Its purpose is instead to generate credible evidence of a student's independent understanding.

This principle is also visible in recent institutional practices. The University of Chicago has moved towards AI-free learning environments in parts of its foundational curriculum, including its core social science sequence, where the objective is to confirm that students first develop the ability to read, write and think independently. Its Law School has adopted a similar sequencing approach in legal research and writing, with students initially building foundational skills without AI before using AI-assisted tools later. The underlying logic is not to reject AI, but to preserve spaces in which core capabilities are developed and demonstrated independently before technology is used to augment them.

Primary interviews indicate growing use of oral examinations, vivas, presentations, group discussions, laboratory-based assessments, project demonstrations, supervised written assessments and practical evaluations for precisely this reason. Several interviewees reported allowing substantial AI use in coursework but validating student understanding through an oral examination or presentation afterwards. Others noted that take home assignments have become less reliable indicators of learning when students can submit sophisticated work that they cannot independently explain. The strongest assessment systems therefore combine AI-enabled work that reflects contemporary practice with bounded oral, written, laboratory, or practical assessments that verify the underlying knowledge and capability.

This reflects a broader change in assessment design: faculty are moving from evaluating final outputs alone towards creating systems that make understanding, judgment, creativity and individual contribution visible. Institutions can support this transition through sample syllabus statements, disclosure templates, assessment rubrics, discipline-specific examples, privacy guidance and instructional design support, with faculty development programs judged increasingly by the quality of course redesign rather than workshop participation alone.

Source: EY-Parthenon analysis

Few leading private universities are moving from standalone AI initiatives to institution-wide capabilities across disciplines and stakeholders

Case study: MAHE – Building institution-wide AI capability

Leading private universities such as MAHE are beginning to move beyond standalone AI courses towards broader institutional integration. MAHE's approach is notable for the breadth of functions and disciplines it is attempting to bring into its AI strategy.

AI is being positioned as a university-wide capability

- In 2026, MAHE announced a partnership with OpenAI to integrate AI across teaching, research and academic administration, with AI literacy intended to extend across health sciences, technology, management, law, humanities and social sciences. The university has also announced plans for structured AI-governance guidelines around responsible and transparent use.
- Through collaboration with Microsoft, the institution has established five AI-focused labs (three on the main campus and two in Bengaluru). The partnership also includes certification programs designed to equip students with the knowledge and skills required for the ethical and effective use of AI. AI initiatives are centrally governed, enabling continuous monitoring of implementation and day-to-day activities.

The emphasis is increasingly cross-disciplinary rather than confined to engineering

- Kasturba Medical College established a dedicated Department of AI in Healthcare, bringing together medicine, engineering, data science and public health, while Manipal Law School introduced coursework on AI-enabled legal practice, prompting, contracts, governance and professional responsibility. MAHE has also partnered with Microsoft to establish AI-focused labs across management, commerce, economics and law schools.

Capability building is extending beyond students

- Faculty-development initiatives have introduced educators in disciplines such as nursing to AI-supported teaching, research and evidence synthesis, with emphasis on verification and human oversight. At the institutional level, MAHE is also experimenting with AI in academic and administrative workflows through MAGIC, its agentic AI platform, with initial applications across timetabling, attendance, examination planning and performance insights and further use planned across admissions, finance, HR, procurement and research support.

AI governance and AI adoption

- The institution maintains a central dashboard to monitor AI tool usage across the ecosystem. Usage metrics, digitalization progress, vendor changes and related initiatives are reviewed monthly as part of a structured governance process.
- Currently, three AI agents have been deployed for students, faculty and administrative staff. The institution aims to achieve 80% to 90% adoption of these agents and expand the ecosystem to 30 to 40 interconnected agents over the next four to five months.
- **Budgeting and Resource planning:** Token budgeting and AI-related expenditure are managed at both the institutional and school levels, with budgets planned and reviewed through formal governance processes.

MAHE therefore illustrates how some leading Indian private universities are beginning to assemble a broader **AI capability stack** spanning literacy, faculty enablement, domain-specific applications, research, governance, industry partnerships and institutional operations. The next opportunity is to deepen these efforts further through more systematic integration into pedagogy, assessment and measurable student outcomes.

Source: EY-Parthenon analysis

Building institutional operating system for safe, scalable and accountable AI adoption

5.4 The "AI-ready institution" operating system

Classroom reform cannot be scaled effectively if individual faculty members and departments act in isolation. Institutional AI readiness should instead be understood as an integrated operating framework comprising four connected layers, a **data layer** governing what information may be used and under what conditions, an **identity layer** determining who may access systems, data, and functionality, an **application layer** governing which AI services are deployed and for what purposes and a **human governance layer** establishing decision-making authority and accountability.

5.4.1 Data layer

The data layer should define the information an institution possesses, the conditions under which it may be used within AI systems and the safeguards that govern its management. This includes questions of ownership, classification, storage, retention, access, sharing, anonymization, correction, deletion and restrictions on model training.

Institutions should distinguish among public information, teaching materials, student academic records, assessment submissions, administrative data, research data and highly sensitive information such as health, disability, counseling and disciplinary records, with each category subject to appropriate access and processing controls. A course-specific AI assistant should not automatically receive access to all institutional databases simply because it sits within the university's technology environment.

While the DPDP framework provides the legal foundation for personal-data governance, higher education institutions may require additional safeguards to protect academic freedom, research confidentiality, assessment integrity and student trust.[69]

5.4.2 Identity layer

Where feasible, institutionally supported AI systems should operate through institutional identities rather than unmanaged personal accounts. Single Sign-On and role-based access controls can differentiate permissions for students, faculty, researchers, administrators and external collaborators while ensuring that access is granted and withdrawn through institutional processes.

Identity management should govern not only access to AI tools but also the data and functionalities available to different users. Faculty developing course assistants, researchers handling confidential information and administrators preparing routine communications require different permissions and safeguards. Indian institutions are already beginning to move towards institutionally managed access: BITS provides institutional access to an LLM, while Symbiosis allocates AI access or tokens through institutional or departmental arrangements. Ashoka similarly provides institutional access to tools including Copilot and Claude to strengthen security and data privacy. Pearl Academy has also partnered with OpenAI to provide students and faculty across its campuses with institutionally governed access to ChatGPT Edu, extending managed AI access into creative disciplines such as fashion, design and communication.

Usage data can help institutions assess adoption, identify areas requiring additional support and evaluate the effectiveness of licensed services, but identity management should not become a mechanism for intrusive surveillance. Individual prompts and interactions should not routinely be monitored by police academic inquiry or intellectual exploration except where clearly defined security, compliance, or misconduct procedures justify such oversight.

5.4.3 Application layer

The application layer should translate institutional AI objectives into a governed portfolio of approved services, which may include tutoring and learning support, language and accessibility tools, research discovery platforms, coding and analytical assistants, faculty course design support, administrative workflow solutions, library search services and institution specific assistants grounded in verified institutional information.

Each application should have a defined purpose, target user group, accountable owner, risk classification, data boundary, performance metrics and review timeline. This is particularly important because institutions should avoid developing multiple uncoordinated AI tools serving overlapping purposes or allowing applications to expand beyond the use cases for which they were initially approved.

---

69. DPDP framework

Source: EY-Parthenon analysis

From human governance to shared infrastructure: Building the foundations for universal AI readiness in higher education

Institution specific assistants should also distinguish clearly between authoritative information and generative interpretation. Systems that provide guidance on scholarships, attendance requirements, academic policies, or institutional regulations should be grounded in approved documents and direct users to the relevant source material for verification.

5.4.4 Human governance layer

Human governance should establish clear decision-making authority rather than operate solely through an advisory committee. A cross-functional governance body should bring together academic leaders, faculty representatives from multiple disciplines, IT and cybersecurity professionals, legal and privacy expertise, research administrators, accessibility and inclusion specialists, administrative staff and student representatives.

Common safeguards, policies and risk thresholds should be established centrally, while departments and faculty retain responsibility for discipline specific and course level decisions. The central governance function should oversee high risk applications, institutional policy, vendor evaluation, incidents, audits, faculty support and communication of material policy changes; within these parameters, faculty needs to determine assignment level AI use, where students should have formal channels for raising concerns about unclear, impractical, or inequitable requirements. For example, Ashoka University is constituting a committee to develop clearer guidelines governing student use of AI in assignments, assessments and academic work.

Governance should also leave room for responsible innovation. Low risk pilots can operate within clearly defined limits on data, users and duration without unnecessarily burdensome approval processes, while higher risk applications should undergo more rigorous review and carry explicit senior-level accountability.

The NIST[70] **Govern–Map–Measure–Manage** framework provides a useful reference because it treats AI risk management as an ongoing institutional responsibility across the AI lifecycle. It connects the establishment of governance and accountability with an understanding of system context and risk, systematic testing and monitoring and continuous mitigation and review, providing a more useful model than treating responsible AI as a one-time approval exercise.[71]

5.5 The AI Highway, shared infrastructure for universal readiness

Most Indian HEIs do not have the resources to independently secure enterprise AI licenses, maintain specialist privacy and cybersecurity teams, conduct model evaluations, develop multilingual AI services and build the full range of technical and pedagogical capabilities required for responsible adoption. Expecting every institution to do so independently would likely produce duplication among well resourced universities, fragmented or potentially unsafe implementation among moderately resourced institutions and limited access for resource-constrained colleges.

The **AI Highway** should therefore be conceived as a federated shared infrastructure model through which institutions with different levels of digital maturity can access common computing resources, AI models, safeguards, training and implementation support while retaining autonomy over academic and institutional decision making. It should not be understood as a single national chatbot or a centralized repository for student data, but as an enabling ecosystem that allows institutions to draw on capabilities that would be inefficient or unaffordable to build individually.

The AI Highway would not require India to build an entirely new technology stack. Several upstream components are already emerging. IndiaAI's common compute infrastructure had expanded to more than 45,000 GPUs by June 2026, while AIKosh hosted more than 14,000 datasets and 331 AI models by July. The national AI Curriculum Taskforce has also proposed shared infrastructure for educational institutions.

The remaining gap is a higher-education-specific access layer that enables institutions to use these capabilities safely and effectively, including approved tools, institutional access controls, multilingual support, privacy and procurement standards, model evaluation and implementation support. Initiatives such as IIT Madras' GenAI4Edu and the emerging Bharat EduAI Stack provide early indications of how this layer could develop.

The opportunity is, therefore, to connect existing national AI infrastructure to institutions that would struggle to build these capabilities independently.

---

70. NIST: National Institute of Standards and Technology
71. NIST Govern–Map–Measure–Manage framework

Source: EY-Parthenon analysis

From compute access to campus readiness: Building equitable foundations for AI adoption across higher education

Examples of foreign governments coming up with AI policy for higher education

International developments reinforce this broader direction. China's **AI + Education Action Plan** establishes a national AI literacy framework spanning school education, higher education and lifelong learning while encouraging discipline-specific curricula and interdisciplinary AI programs. The **EU AI Act**, although a cross-sector regulatory framework rather than a dedicated education policy, similarly combines an emphasis on AI literacy with stronger safeguards for consequential applications such as admissions, assessment, learning-pathway decisions and AI-based monitoring. Together, these developments point towards a wider shift from encouraging AI adoption to building the common capabilities and safeguards required to support it responsibly.[72]

5.5.1 Shared infrastructure across institutional tiers

The AI Highway should provide pooled access to computing infrastructure, approved AI models and institutionally negotiated services through subsidized or differentiated pricing arrangements. A secure model gateway could authenticate users, route requests to approved providers, enforce appropriate data-protection controls and usage limits and allow institutions to adopt or switch models without rebuilding existing integrations.

Shared support should extend beyond the technology itself. Curated datasets, evaluation frameworks, multilingual and accessibility tools, governance policies, consent templates, procurement guidance, risk-classification frameworks, course-policy templates, sample rubrics, assessment-design resources and specialist expertise in privacy, cybersecurity, bias auditing and accessibility could all be made available through common infrastructure, reducing the need for institutions to recreate the same capabilities independently.

National platforms can provide the underlying compute, models and datasets, while the AI Highway converts these assets into institution-ready services. The objective is not to make institutions technologically uniform, but to establish a baseline of safe, equitable and effective access while allowing institutions with greater capacity to innovate beyond that common foundation.

5.5.2 Extending the institutional operating system through equity of access mapping

The effectiveness of shared infrastructure ultimately depends on whether institutions and students can use it meaningfully. Access should therefore be assessed not simply by whether an AI service is technically available, but by whether students and faculty have the connectivity, devices, language support, accessibility features, institutional guidance and capability required to benefit from it.

This matters because the burden of implementing responsible AI is not distributed evenly. Resource constrained institutions are less able to develop specialist legal, technical, cybersecurity and pedagogical capabilities independently, even though they may serve students for whom equitable access to AI-enabled learning is particularly consequential. The AI Highway should therefore complement institution-level operating systems by reducing these capability gaps through common infrastructure and shared support.

The true measure of an AI-enabled campus is not the number of AI tools deployed or the frequency of their use, but whether the institution can explain **how and why AI is being used, what data and risks are involved, where human responsibility remains, how errors can be challenged or corrected and whether students have an equitable opportunity to benefit from its adoption**.

Higher education institutions face mounting pressure from constrained budgets, legacy technology environments and geopolitical uncertainty. The challenge is not whether to adopt AI, but how to do so without replacing existing systems. Our focus is on modernizing the current technology estate, either by enhancing existing platforms or deploying an agentic AI layer on top of them. ***- Global AI and innovation leader, education business unit, global IT services company***

72. China's AI + Education Action Plan , European Commission , IndiaAI's compute platform , AIKosh

Source: EY-Parthenon analysis



India's AI imperatives for higher education, 2026-2030

Chapter 6

While India has established the foundations for AI adoption, a coordinated national readiness framework is needed to drive consistent implementation across higher education institutions

India already has important components of a national AI ecosystem, including shared compute, datasets, skilling initiatives and responsible AI programs. The challenge now is to build the policy and implementation architecture through which these capabilities can reach higher education institutions consistently and responsibly.

The principal gap is therefore not a lack of ambition or isolated initiatives, but the absence of a coherent framework that defines minimum institutional readiness, clarifies national and institutional responsibilities, establishes common standards where necessary, supports under-resourced institutions and provides a shared basis for measuring progress and outcomes.

Given the pace at which AI capabilities and educational applications are evolving, highly prescriptive rules are unlikely to remain relevant for long. A more durable approach would follow the "light but tight" principle articulated in the National Education Policy 2020, clear expectations and accountability at the system level, combined with sufficient institutional flexibility in implementation. The NEP's emphasis on critical thinking, conceptual understanding, multilingualism, equity, faculty development and appropriate technology integration remains directly relevant to the AI transition.[73]

A national framework should distinguish between baseline AI readiness and advanced institutional capability. Every recognized higher education institution should be expected to meet a minimum threshold in areas such as student guidance, assessment integrity, approval of AI tools, data protection, faculty orientation and grievance redressal. Institutions with stronger resources and research capacity should be encouraged to move beyond this baseline by investing in AI model development, advanced research, institutional AI platforms and industry linked innovation.

The period from 2026 to 2030 should be treated as a phased implementation horizon rather than a distant policy aspiration. The immediate priority in 2026 should be to establish common standards, governance principles, funding mechanisms and a credible measurement framework. The following two years should focus on strengthening faculty capability, expanding shared infrastructure and supporting institution level implementation. By 2029-2030, progress should be judged not by the number of AI tools deployed or committees established, but by demonstrable improvements in learning outcomes, equity, research productivity and graduate employability.

73. PIB

Source: EY-Parthenon analysis

Building the foundations for AI readiness through infrastructure, AI literacy and faculty capability development

6.1 National priorities that cannot be deferred any longer

1 Extend national AI infrastructure into an HEI-ready access layer

- IndiaAI Compute already provides important upstream infrastructure. Higher education institutions now require a broader enabling layer that includes secure access to approved AI tools, multilingual capabilities, data safeguards, procurement support and technical assistance.[74]
- Minimum infrastructure standards should be tiered. All institutions should have access to reliable connectivity and essential AI services, while those with stronger capabilities should be able to use more advanced cloud, data, compute and research infrastructure.
- Shared national or regional services could reduce duplication and help smaller institutions access capabilities that would be prohibitively expensive to build independently.

2 Position AI literacy as a core graduate capability across disciplines

- AI literacy should be treated as a foundational graduate capability rather than a specialized computer science subject.
- All students should understand, at a minimum, how AI systems function, where their limitations lie, how outputs should be verified, what data can be shared safely and when AI use should be disclosed.
- Curriculum reform should be embedded within disciplines rather than delivered through a single generic course. Law students should learn to verify AI-generated legal references and address confidentiality and professional responsibility. Commerce students should evaluate AI-generated financial analysis critically, while humanities students should engage with questions of authorship, interpretation and meaning.
- UNESCO's AI Competency Framework for Students offers a useful reference by framing AI literacy around human centered thinking, ethics, knowledge and application. India can adapt this model to its own institutional and disciplinary context rather than impose a uniform course across all programs.[75]

3 Establish a national faculty capability framework focused on teaching and assessment

- Faculty readiness should not be measured by attendance at introductory training programs. The more meaningful test is whether educators can redesign courses, assessments and classroom practices for an AI-enabled learning environment.
- A national faculty development strategy should combine a common foundation with discipline specific implementation support.
- The foundational component should cover AI limitations, academic integrity, data protection, disclosure, bias, verification and responsible use.
- Discipline specific support should help faculty redefine learning outcomes, redesign assignments and adopt assessment methods suited to the pedagogical requirements of their fields.
- Faculty would also require instructional design support, communities of practice and protected time to translate training into changes in teaching.

74. India AI Compute

75. UNESCO's AI Competency Framework for Students

Source: EY-Parthenon analysis

AI readiness requires assessment reform, institutional accountability and equitable implementation across higher education

4 Reform assessment for an AI-enabled learning environment

- National guidance should help institutions preserve academic integrity while allowing legitimate and discipline appropriate uses of AI.
- It should define broad categories of acceptable, restricted and prohibited use, supported by discipline specific examples rather than abstract principles alone.
- The objective should be to generate credible evidence of learning through methods such as supervised assessments, oral examinations, practical demonstrations, staged submissions, process documentation and carefully designed AI-enabled tasks.
- Assessment reform should focus less on detecting AI-generated language and more on making student reasoning, judgment and individual contribution visible.

5 Embed AI readiness within quality assurance and accreditation frameworks

- Accreditation should extend beyond assessing the existence of AI policies, laboratories or Centers of Excellence. Instead, it should evaluate whether
- Students understand institutional AI policies
- Assessment practices have been adapted for AI-enabled learning
- Faculty receive adequate support
- AI tools are governed effectively
- Data management processes are clearly defined
- Students have access to mechanisms for challenging harmful or unfair AI-related decisions

6 Make equity a design requirement

- National policy need to recognize that, without targeted intervention AI adoption is likely to reinforce existing institutional inequalities.
- Well-resourced institutions will be better positioned to procure advanced tools, train faculty and develop internal capabilities, while less resourced institutions may rely on limited public tools and informal adoption.
- Achieving equity therefore requires more than equal access to platforms; solutions need to be affordable, multilingual, accessible to students with disabilities and compatible with low-bandwidth environments. Progress should be measured through meaningful usage and educational outcomes rather than tool availability alone.

7 Establish clear institutional accountability

- Every HEI should have clear accountability for AI strategy, tool approval and oversight of high-risk use cases.
- National policy should define these core governance functions while allowing institutions to adapt responsibilities to their scale and capacity. Smaller institutions may benefit from shared regional support for privacy, security and model evaluation.

Source: EY-Parthenon analysis

As AI adoption accelerates, India needs clear and consistent governance frameworks to guide institutional decision-making

6.2 Policy gaps that need to be fixed in 2026

Higher education institutions are making consequential AI-related decisions within a fragmented policy environment shaped by data protection regulation, national AI initiatives, existing academic integrity frameworks, vendor agreements and institution specific policies. While these mechanisms provide broad direction, they do not yet amount to a coherent operational framework for AI governance in higher education.

In the absence of common standards, divergent practices are becoming increasingly entrenched. Assessment rules vary across institutions and even across faculty within the same institution. Vendor contracts provide inconsistent safeguards, while students often develop their own norms for AI use in response to unclear guidance.

The pace and maturity of adoption are also uneven. Many private universities are progressing more rapidly than publicly funded institutions because of differences in resources, infrastructure and institutional capacity. Without timely intervention, high-risk applications may move from experimentation to routine deployment before their educational, ethical and operational implications have been adequately assessed.

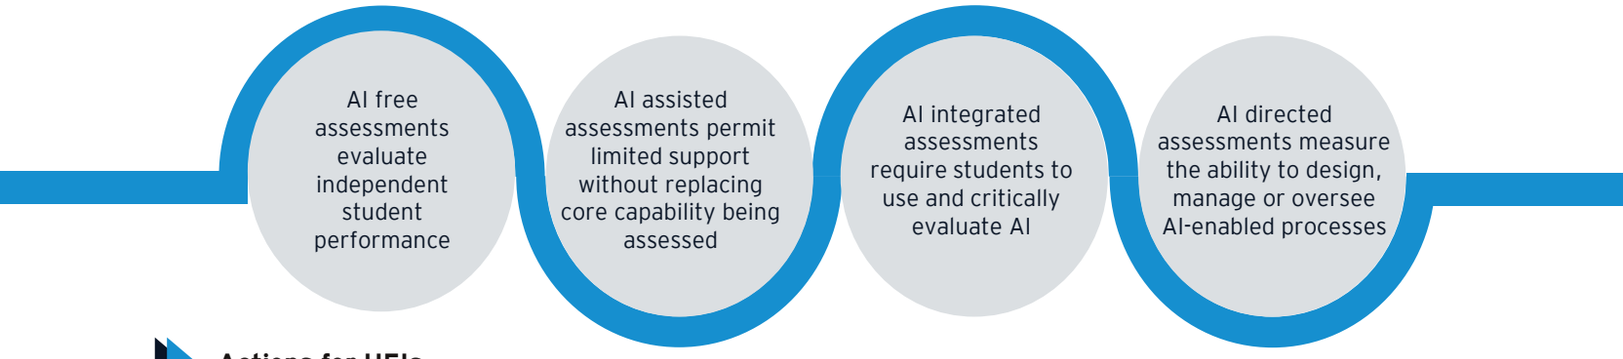
"Agility matters. Institutions that have the independence to revise curricula and engage closely with employers are much better positioned to keep programs aligned with industry needs." **- Mr Kalyan Kolachala, India Head, Symphony AI**

6.2.1 Clear rules on AI use in assessment

India requires a common national framework for AI use in assessment that provides clarity and consistency without imposing identical rules across all disciplines, courses and institutions. National regulators should establish baseline principles and minimum standards for academic integrity, transparency, disclosure and responsible AI use. Universities, departments and faculty should retain the flexibility to adapt these requirements to their disciplinary, pedagogical and assessment contexts.

Suggestions to HEIs

The national standard should require **every significant assessment to clearly specify permitted AI use**. A practical framework could distinguish between AI free, AI assisted, AI integrated and AI directed assessments. (Explained in detail in Chapter 5.)



Actions for HEIs

Institutions should publish clear assessment rules and provide faculty with practical templates that allow requirements to be adapted by discipline and learning objective.

AI detection tools should not serve as the sole basis for academic misconduct findings nor should arbitrary AI-content thresholds automatically be treated as evidence of plagiarism. Decisions should rely on corroborating evidence, transparent procedures and accessible appeal mechanisms.

Policy effectiveness should be assessed through student understanding and changes in assessment practice rather than through the existence of an institutional policy alone.

Source: EY-Parthenon analysis



As AI adoption expands, risk-based governance frameworks are needed to support trust, accountability and student protection

6.2.2 Model risk management standards

Chapter 5 establishes a risk-based framework for classifying AI applications according to the consequences of their use. The immediate policy priority is to translate that framework into consistent governance practices across institutions.

Risk classification should determine the level of oversight, transparency, documentation, validation and human review required. Applications with greater potential to affect student outcomes, rights, access to opportunities or wellbeing should be subject to stronger safeguards and clearer institutional accountability.

Institutional AI registers

Every medium and high-risk AI system should be recorded in an institutional AI register. At a minimum, the register should document:

- The purpose and intended users of the system
- The vendor or model provider
- The data inputs and processing arrangements
- The designated institutional risk owner
- Human-oversight and escalation mechanisms
- Known limitations and conditions of use
- The approval date and review schedule

Effective risk management is possible only when institutions have visibility into the AI systems operating across departments.

Impact assessment and approval

High-risk systems should undergo a formal impact assessment before deployment. This should examine:

- Whether the application is necessary and proportionate to the objective
- Accuracy and reliability within the intended educational context
- Privacy, security and data protection implications
- Potential bias and variation in performance across student groups
- Accessibility and implications for students with disabilities
- Effects on student rights, opportunities, reputation and wellbeing
- Whether less intrusive alternatives are available
- The adequacy of human review, appeal and redress mechanisms

Approvals should be time bound rather than indefinite. Reassessment should be required when there are material changes to the model, intended use, vendor, data inputs or processing arrangements.

Local validation

National standards should be complemented by local testing. Vendor benchmarks alone cannot demonstrate whether a system would perform reliably across Indian English, regional languages, disciplinary contexts or institutions with different student populations.

Education specific evaluation should examine risks such as hallucinated information, fabricated citations, false positives, inconsistent grading, biased recommendations and unequal performance across student groups. Institutions should establish that a system is fit for its intended purpose rather than relying solely on general claims of model accuracy.

Source: EY-Parthenon analysis

Governance must move beyond policy statements to embedded systems for evaluation, monitoring and continuous improvement

AI management-system approach

NIST's AI Risk Management Framework and its Generative AI Profile provide useful foundations because they embed trustworthiness throughout the design, deployment, use and evaluation of AI systems rather than treating approval as a one-time exercise.[76]

AI RMF use-case profiles represent tailored applications of the AI Risk Management Framework (AI RMF) functions, categories and subcategories to specific contexts, sectors or use cases. These profiles are developed based on an organization's objectives, risk tolerance, regulatory obligations and available resources. Examples may include AI RMF profiles for recruitment, housing, healthcare or education. By illustrating how AI-related risks can be identified, assessed and managed across different stages of the AI lifecycle, such profiles provide practical guidance for sector-specific implementation. They enable organizations to align AI risk management practices with their strategic goals, legal and regulatory requirements, industry best practices and broader risk management priorities.

ISO/IEC 42001 similarly adopts a management-system approach under which organizations maintain and continually improve their policies and processes for responsible AI use. India should adapt the underlying principles to the higher education context rather than require individual colleges to pursue costly international certification.

Shared evaluation capacity

The proposed AI Highway could host shared evaluation, audit and red teaming services. Institutions should be able to submit prospective tools for privacy, security, bias, accessibility and educational-performance review rather than having to build specialist teams independently.

A national, appropriately anonymized incident repository could also allow institutions to learn from failures identified elsewhere and reduce the repeated deployment of systems with known limitations.

76. Source: NIST's AI Risk Management Framework , AI Risk Management Framework (AI RMF) , ISO/IEC 42001, EY-Parthenon analysis

As AI adoption expands, higher education institutions require clear data governance frameworks to protect privacy, trust and institutional accountability

6.2.3 Data privacy frameworks for HEIs

The Digital Personal Data Protection Rules, 2025, establishes a national framework for personal data processing, security, user rights and breach management.[77] Higher education institutions now require sector specific guidance that translates these obligations into the contexts of teaching, research, assessment, student support and AI procurement.

Establish a common data taxonomy

The framework should begin with a common HEI data taxonomy covering:

- Public institutional information
- Administrative and operational data
- Student records and academic information
- Assessment submissions and learning data
- Research data and unpublished intellectual work
- Sensitive information, including health, counseling, disability, behavioral and disciplinary records

Each category should have clearly defined rules governing AI use, retention, vendor access, disclosure and model training.

Protect student and faculty work

Essays, projects, code, recordings, research outputs and other forms of intellectual work should not be used to train external models by default. Students and faculty should be informed when their work is retained, analyzed or used for system improvement. These practices should be supported by clear policies on intellectual property, research integrity, consent and institutional ownership.

Where AI use is optional, refusal should not result in academic disadvantage. Where data processing is necessary, institutions should explain the legal basis and purpose transparently.

Apply enhanced safeguards to sensitive data

Counseling, health, disability, disciplinary and behavioral data require stronger protections. AI-generated inferences based on such information should be transparent, contestable and subject to strict purpose limitations.

Institutions should also provide clear guidance on information that must not be uploaded to public AI models, including:

- Student records
- Examination materials
- Unpublished research
- Confidential institutional documents
- Proprietary or commercially sensitive information
- Sensitive personal data

Secure institutional alternatives must be provided if compliance is to be practical rather than merely aspirational.

77. Source: DPDP

Source: EY-Parthenon analysis

Effective oversight requires common procurement frameworks, vendor accountability and scalable compliance support

Strengthen procurement and vendor oversight

Cross-border storage, subcontracting, vendor access and model training arrangements should be addressed through common procurement guidance and shared technical and legal capability.

Institutions remain accountable for the AI tools they approve. They need to understand what data is collected, why it is processed, where it is stored, who can access it, whether it is shared with downstream parties and how users may exercise their rights.

National templates should be developed for:

- Data mapping
- Privacy notices and consent communications
- Vendor due diligence
- Data-processing agreements
- Breach response
- Retention and deletion
- Cross-border processing review

These resources would help smaller institutions comply more effectively and reduce their reliance on privacy terms defined primarily by vendors.

Adopt a proportionate compliance model

Smaller HEIs cannot be expected to maintain the same legal, privacy, cybersecurity and technical capacity as large universities.

National guidance should define the essential data-protection functions for which every institution remains accountable, while enabling shared implementation support through state university systems, regional hubs or other common service structures.

Source: EY-Parthenon analysis

Sustainable AI adoption requires funding models that support infrastructure, capability building, governance and equitable access

6.2.4 Funding rules for AI infrastructure

Central funding for AI readiness should extend beyond capital investments such as servers, labs and licenses. AI adoption depends on recurring expenditure, including cloud services, software subscriptions, cybersecurity, technical support and continuous training.

- Universal funding should enable all eligible institutions to meet minimum requirements, including approved AI environments, baseline AI literacy, data governance, assessment guidance and incident reporting.
- Competitive funding should support advanced research, indigenous AI applications, institutional platforms, centers of excellence and industry linked innovation.

Institutions serving disadvantaged, rural, remote or regional language populations, or managing large affiliated college networks, will require additional support to achieve comparable readiness levels.

Funding should be linked to implementation outcomes rather than procurement alone. Evidence may include redesigned assessments, demonstrated faculty capability, student AI literacy outcomes, equitable adoption, effective governance processes and independent review of high-risk systems. Institutions should also disclose the long-term costs of major AI deployments and maintain sustainability and exit plans to manage future financial commitments.

Source: EY-Parthenon analysis

A National AI Readiness Index can provide a structured framework for measuring institutional preparedness across higher education

6.3 A National HEI AI Readiness Index: Proposed framework

A 100-point framework could organize indicators across seven dimensions. The weights below should be tested through stakeholder research and piloting rather than treated as final.

Table 1:

Indicator	Description	Points
Governance and accountability	Indicators should examine whether the institution has a usable AI policy, a cross functional governance body, clear decision rights, an AI system register, risk classification procedures, incident reporting and functioning appeal mechanisms. Evidence should test whether these structures operate in practice rather than whether they exist in a document.	15 points
Curriculum, pedagogy and assessment	Indicators should examine the spread and depth of discipline-specific AI integration, the proportion of courses with clear AI use guidance, the extent of assessment redesign, the use of verification and disclosure and whether institutions preserve credible evidence of independent learning. Greater weight should be assigned to changed teaching practice than to the number of standalone AI courses.	20 points
Faculty capability	The Index should measure the proportion of faculty who have demonstrated baseline capability, received discipline-specific support and redesigned at least one course or assessment. Workshop attendance alone should receive limited credit. Advanced indicators may cover instructional design support, workload recognition and faculty participation across departments.	15 points
Student literacy, safety and voice	Institutions should demonstrate that students can identify hallucinations, verify sources, protect confidential information, disclose AI use and explain relevant institutional rules. Indicators should also cover student participation in governance, awareness of data practices, availability of redress and the experience of students who choose not to use optional AI tools.	15 points
Technology, data and model risk readiness	Indicators should cover connectivity, secure identity, approved tool access, data mapping, vendor review, access controls, cybersecurity, local model evaluation and human oversight of medium and high-risk applications. The Index should not reward raw computing expenditure without evidence that the infrastructure is secure, used and aligned with educational needs.	15 points
Equity, inclusion and accessibility	Indicators should measure differences in access and capability by institution location, student income, language, disability, discipline and device type. Institutions should receive credit for mobile and low bandwidth access, multilingual support, accessibility testing, subsidized tool availability and targeted assistance that reduces rather than merely documents disparities.	15 points
Research, industry alignment and demonstrated outcomes	Indicators should examine responsible AI research, applied projects, access to compute and datasets, industry-supported learning, student portfolios and evidence that AI adoption is improving research, employability or institutional effectiveness. The relatively low initial weight reflects the need to establish the governance and learning foundations before rewarding advanced outputs.	5 points

Source: EY-Parthenon analysis

Measuring AI readiness requires a balanced assessment of institutional reach, depth of adoption and governance maturity

Proposed design principles for an AI Readiness Index

The weighting methodology should balance horizontal adoption and vertical integration, reflecting the AI readiness framework outlined earlier in this report. Horizontal indicators should assess whether AI capabilities and support mechanisms reach students, faculty and departments across the institution. Vertical indicators should evaluate the depth of AI integration within teaching, assessment, research, governance and professional practice. This approach encourages institutions to build broad-based readiness rather than rely on isolated pockets of excellence.

The Index should be based on a combination of documentary evidence, stakeholder experience and independent verification. Institutions may submit policies, training records, assessment guidelines, system inventories and governance documentation. These submissions should be complemented by faculty and student surveys* to determine whether policies and processes are understood and consistently applied in practice. Independent audits and periodic reviews of courses and systems can help validate institutional claims and reduce the risk of overreliance on showcase examples.

Meaningful reporting requires transparent and contextualized metrics. For example, reporting the number of faculty trained is only valuable when accompanied by information on the total faculty population, training content and duration, departmental coverage and evidence of changes in teaching practice. Similarly, reporting AI license access should include utilization rates, participation across different student groups and the quality of support provided to users.

The Index should prioritize demonstrated capability over expenditure. Institutions should be assessed on how effectively AI is implemented and governed rather than on the size of technology investments. A resource-constrained institution that successfully deploys shared infrastructure and demonstrates broad and effective use may be more AI-ready than a well-funded university with extensive licenses but limited practical usage or governance.

Certain foundational safeguards should function as minimum eligibility requirements rather than scoring criteria. Institutions should not be classified as highly AI ready if they lack basic student guidance on AI use, have no mechanisms for handling sensitive data, permit high-risk automated decision-making without oversight or appeal processes or rely solely on AI detection systems as evidence of academic misconduct. Strong performance in one area should not compensate for critical shortcomings in student protection or governance.

Public reporting should provide a multi dimensional view of readiness, including overall scores, dimension level performance, methodology, peer group comparisons and data quality assessments. A single ranking would mask important differences between institutions, such as strengths in infrastructure but weaknesses in assessment practices or strong research capabilities alongside limited equity measures.

The Index should also recognize progress over time. Institutions that make significant improvements in governance, faculty capacity, infrastructure or student support should receive recognition even if they remain below more established universities in absolute rankings. This would position the Index as a tool for institutional improvement rather than solely as a mechanism for reputational comparison.

The link between the Index, accreditation and funding should be introduced through a phased approach. Initially, low scores should trigger advisory support, capacity building initiatives and targeted funding rather than punitive measures. As indicators mature and data quality improves, minimum readiness thresholds could gradually be incorporated into accreditation reviews, while demonstrated progress against identified gaps may inform funding decisions.

* While surveys provide useful insights, they should not be treated as definitive measures of institutional performance. Over time, survey findings can be supplemented with quizzes or other evidence sources, creating a more robust and continuously refined assessment process.

Source: EY-Parthenon analysis

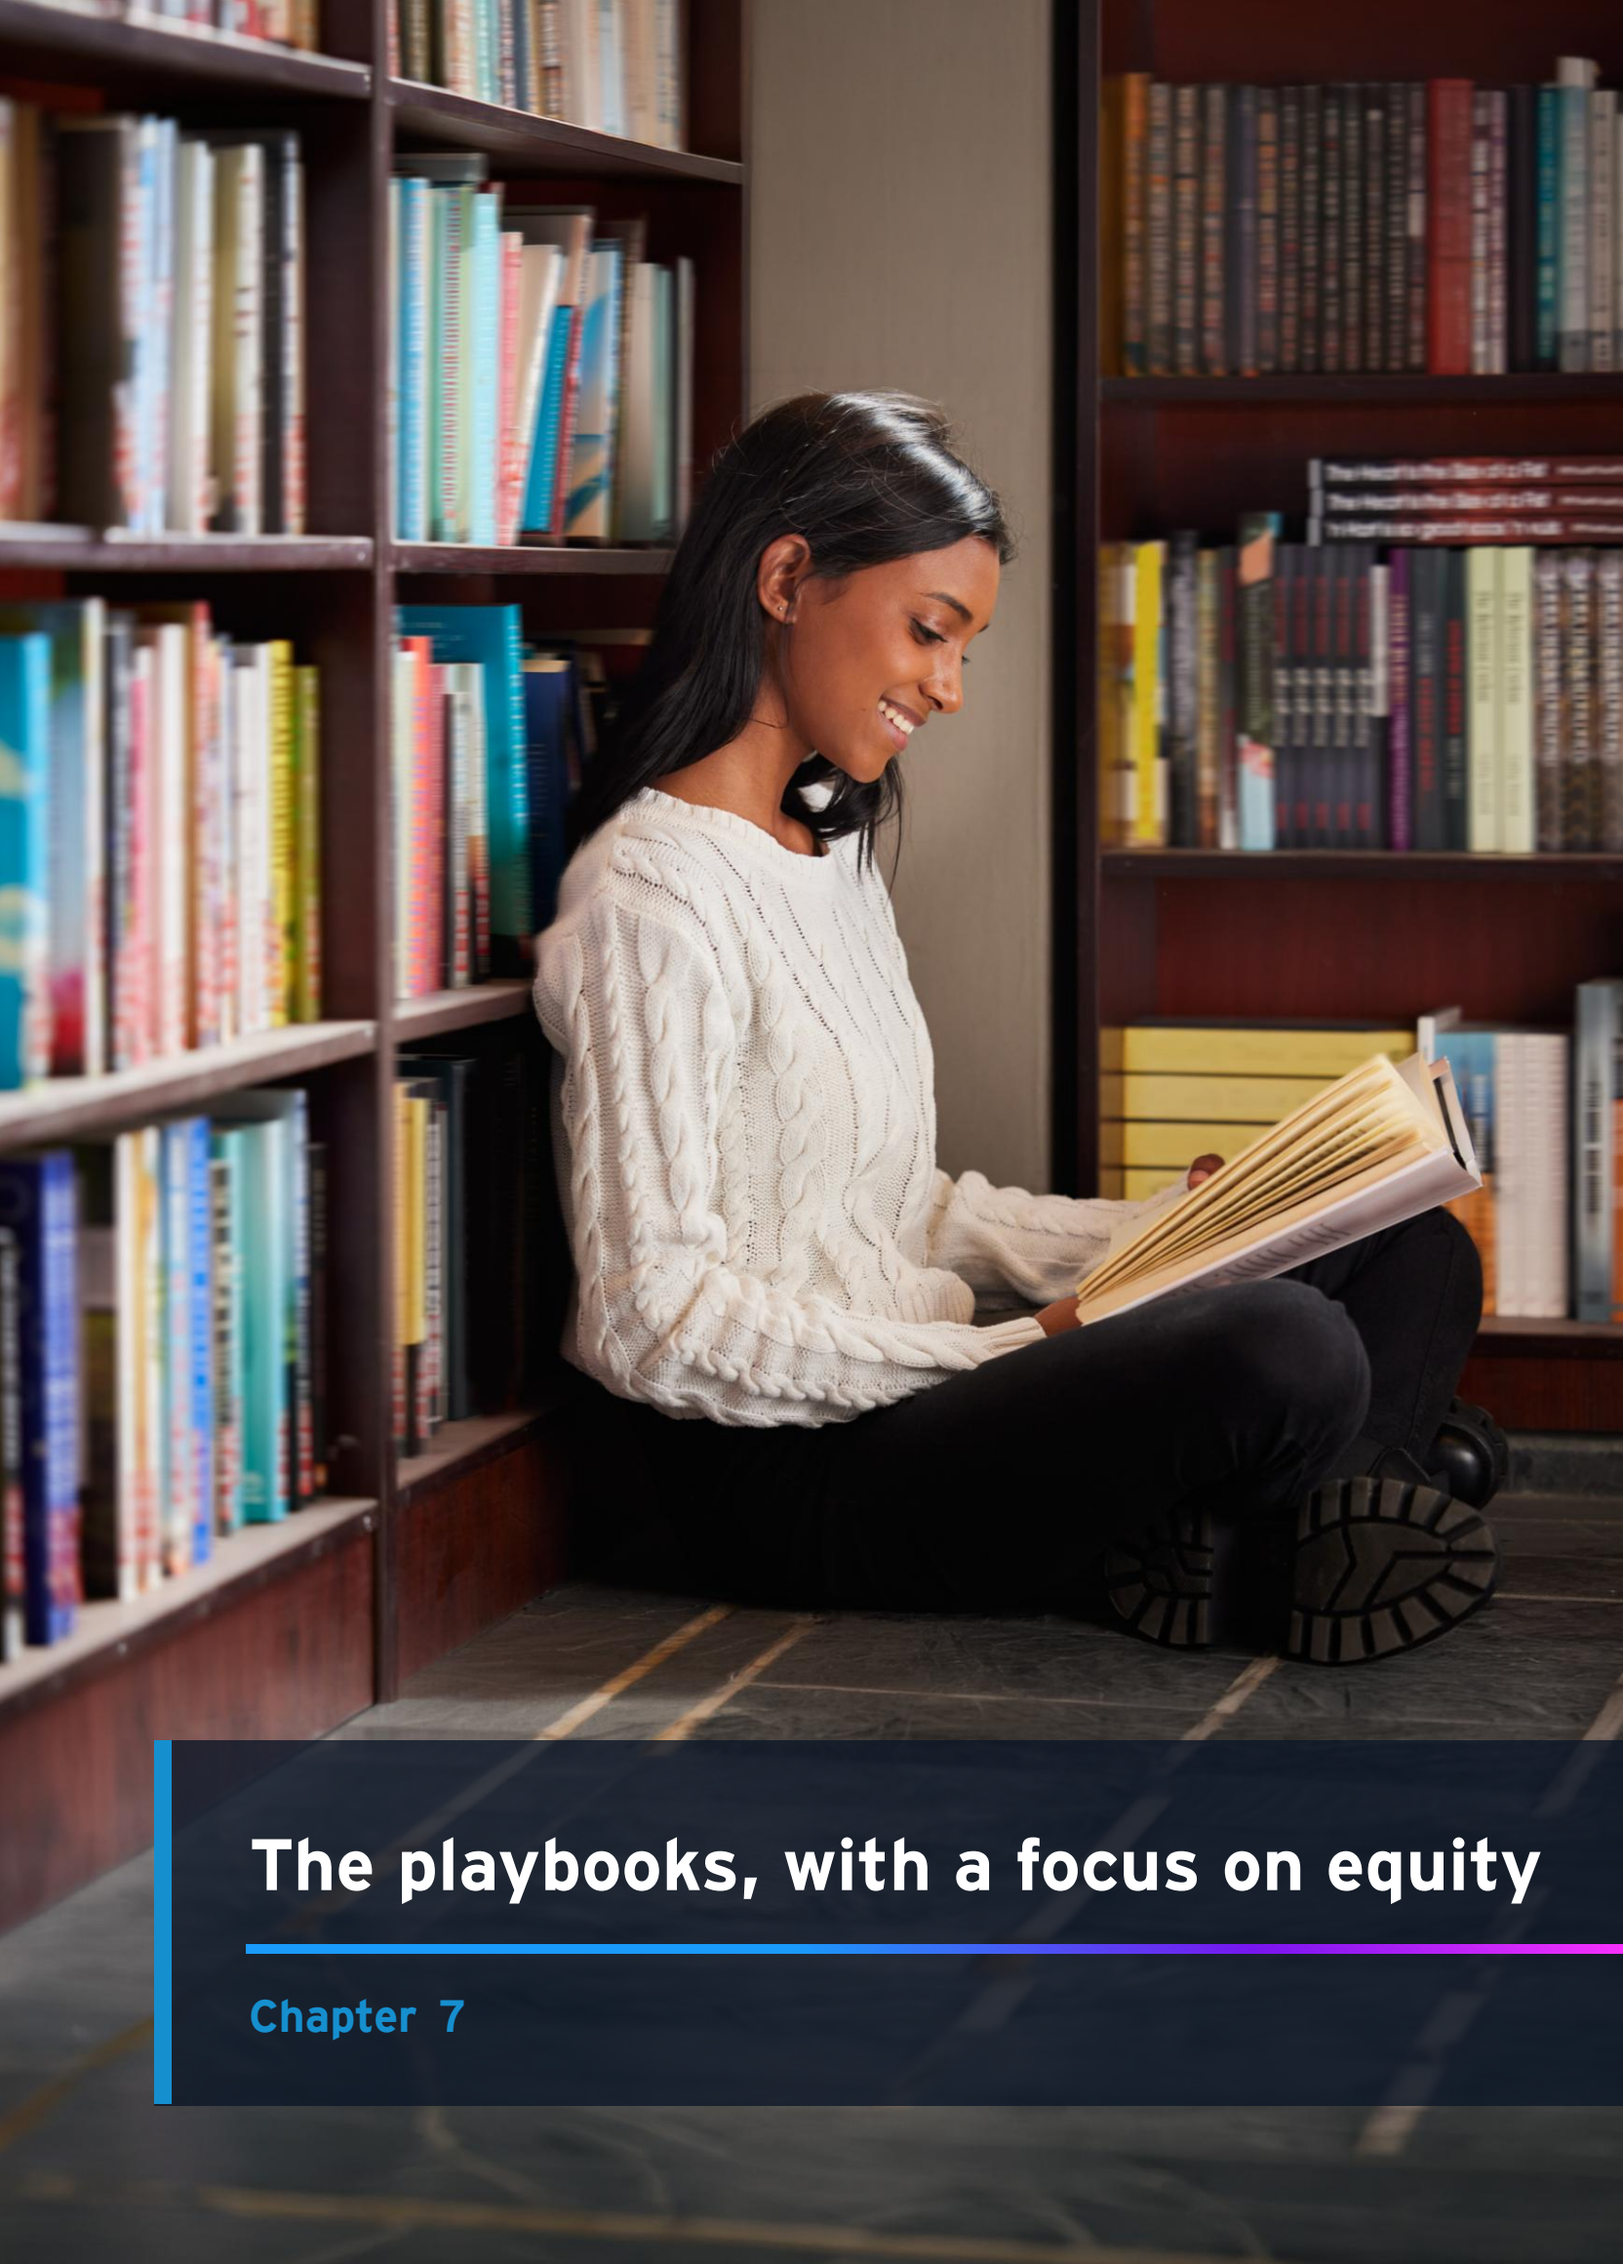
A dedicated AI Readiness Index should be piloted and refined before integration into accreditation and ranking systems

India's experience with the National Institutional Ranking Framework (NIRF) demonstrates the feasibility of collecting institution-level data and assessing universities across multiple weighted dimensions. However, AI readiness should not simply be incorporated as a minor component within existing rankings. A dedicated AI Readiness Index would allow indicators, evidence requirements and benchmarking methodologies to mature before consideration is given to integration with broader accreditation or ranking systems.

A phased rollout strategy should begin with a pilot involving a diverse mix of central, state, private, deemed, autonomous and affiliated institutions across regions. The pilot should assess the practicality of data collection, the verifiability of indicators, the reporting burden on smaller institutions and the risk of unintentionally favoring wealthier universities. Findings should be independently reviewed and used to refine the methodology before national implementation.

During the initial years, results should be reported through maturity bands and dimension-level assessments alongside rankings. Given the rapid evolution of AI technologies, policies and institutional practices, the methodology should be reviewed and updated annually. Indicators should remain technology-neutral, focusing on institutional capabilities and outcomes rather than prescribing the use of specific AI models, platforms or commercial vendors.

Source: EY-Parthenon analysis



The playbooks, with a focus on equity

Chapter 7

AI readiness is not about avoiding AI, but knowing when it enhances learning and when it replaces it

From AI principles to repeatable institutional practices

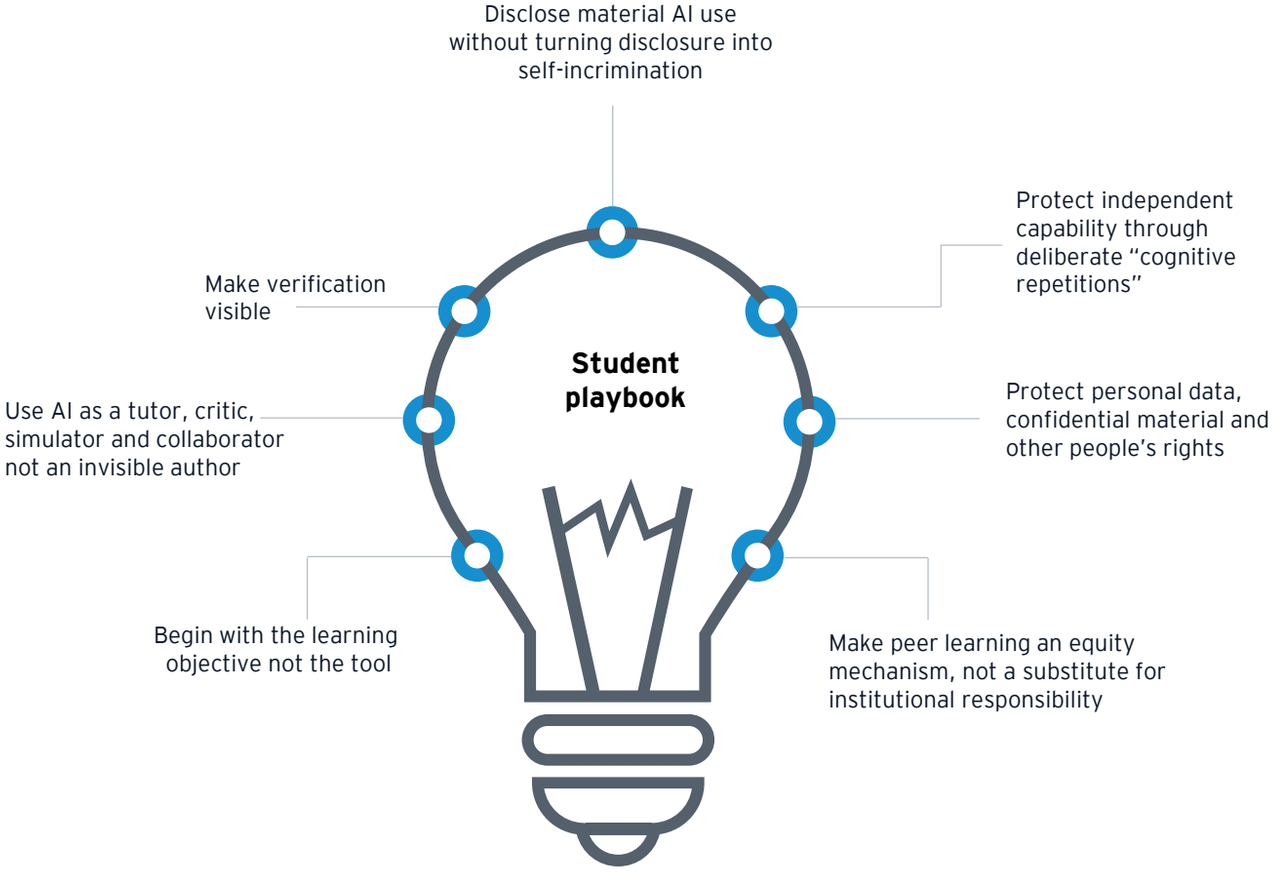
7.1 The student playbook

How to use AI ethically, responsibly and productively without surrendering intellectual agency

The student playbook replaces moralistic warnings and generic guidance on prompt engineering with a structured and repeatable learning framework. Rather than unrealistically expecting students to avoid AI, it helps them distinguish between situations where AI enhances their thinking and those where it substitutes for the capabilities they are expected to develop.

Our primary research indicates that AI use is often driven by behavioral habit rather than conscious ethical decision-making. Students frequently turn to AI because it is fast, readily accessible and continuously available, only later reflecting on whether meaningful learning has occurred. One student described a **shift from "thinking-first" to "AI-first,"** while another acknowledged repeatedly returning to AI despite attempts to reduce usage. Consequently, the playbook emphasizes the development of effective habits and task-design strategies rather than relying solely on intentions or awareness.

This finding is closely linked to the employer implications discussed in Section 3. While AI can help produce polished assignments, presentations, projects and resumes that may generate short term recognition, such outputs are unlikely to withstand oral examinations, technical interviews or workplace challenges if students cannot articulate the assumptions, sources, calculations and decisions underpinning their work.



Source: EY-Parthenon analysis

While AI can accelerate task completion, effective learning still depends on student reasoning and intellectual engagement

7.1.1 Begin with the learning objective, not the tool

Before using AI, students should first clarify the purpose of the task by asking:

- What knowledge, understanding or capability should I develop through this activity?
- Which aspects of the task require my own reasoning and judgment?
- Which repetitive, administrative or supporting activities can be appropriately accelerated through AI?
- How will I determine whether AI has enhanced my learning rather than simply helped me complete the task?

A **"think-first" principle** should be applied to foundational and unfamiliar tasks. Before consulting AI, students should make a brief independent attempt to engage with the problem, for example, by outlining the issue, proposing an initial hypothesis, completing the first step, identifying uncertainties or drawing upon prior knowledge.

This is not intended as a rigid prohibition on "AI-first" approaches. Rather, it establishes a baseline from which students can evaluate AI's contribution and prevents AI systems from framing the problem before meaningful learner engagement has occurred.

The application of this principle should vary according to the learner's stage of development. A first-year programming student, for instance, may need to independently write and debug basic code before relying on AI generated solutions, whereas a final year student may appropriately use AI for routine implementation while retaining responsibility for system architecture, testing and security considerations. Similarly, a novice writer may draft content independently before using AI for revision, while a more advanced student may employ AI to challenge, refine or critique an already developed argument.

Importantly, manual effort should not be equated with learning. Students need not perform repetitive tasks solely because they can be completed without automation. The key consideration is whether a particular activity develops a durable capability that remains valuable when AI is unavailable, produces incorrect outputs or cannot be appropriately applied.

7.1.2 Use AI as a tutor, critic, simulator and collaborator-not an invisible author

Students should be taught productive AI roles through practical and repeatable workflows:

Tutor	Critic	Simulator	Research assistant	Production assistant	Comparator
Explain concepts at different levels, provide worked examples, generate practice questions and respond to an attempted solution	Identify weak assumptions, missing evidence, counterarguments and alternative interpretations	Rehearse interviews, oral examinations, client conversations, negotiations, clinical reasoning or code reviews	Surface terminology, candidate sources and competing hypotheses, while requiring the student to find and read the underlying material	Support formatting, translation, transcription, routine coding, visualization or editing after the core thinking is complete	Generate several approaches that students need to evaluate rather than accepting the first plausible answer

Source: EY-Parthenon analysis

While AI can accelerate learning and productivity, responsible use depends on verification and transparent disclosure

7.1.3 Make verification visible

Verification should be embedded as a core component of AI use rather than treated as a final reminder to "check the answer." Students require a structured approach that distinguishes between factual claims, interpretations, calculations, citations and recommendations.

Effective verification practices include:

- Consulting original sources rather than relying on AI-generated summaries
- Prioritizing primary evidence and official datasets
- Independently reproducing calculations, validating time-sensitive information against current references
- Testing codes across normal and edge-case scenarios
- Assessing whether evidence is relevant to the appropriate geography, population and time period

Students should also learn to recognize responses that appear authoritative but are supported by weak or insufficient evidence. The level of verification should be proportionate to the stakes involved. Errors in brainstorming exercises carry different consequences from inaccuracies in medical advice, financial models, legal analyzes, research outputs or employment applications.

We found that some universities' AI policies require students to submit their prompt history as a disclosure mechanism for AI use. Beyond this, faculty can make verification a teachable skill by requiring targeted AI claim audits rather than comprehensive prompt logs. For example, students could be asked to identify the most consequential AI-generated claims in their work and demonstrate how each was validated, refined or rejected.

Insights from our primary research suggest that responsible verification practices do exist but remain inconsistent. While some students reported verifying all AI-generated outputs, faculty also encountered instances of technically polished submissions that failed to address the underlying question. Together, these findings indicate that verification remains an individual practice rather than an established institutional norm.

7.1.4 Disclose material AI use without turning disclosure into self-incrimination

Students require a proportionate and standardized disclosure framework that promotes intellectual transparency without automatically creating suspicion. A concise **AI Contribution Note** should document the tool used, the purpose of its use, the nature of its contribution and the steps taken by the student to review, modify and verify the output.

For example: *"I used an approved generative AI tool to explore alternative structures for the literature review and test counterarguments. I selected the final structure, independently reviewed all cited sources, rejected unsupported claims and rewrote the analysis in my own words."*

Disclosure requirements should be tiered according to the level of AI involvement. Routine uses, such as spell checking or accessibility support, may require no disclosure or a standard declaration. More substantive uses involving outlining, analysis, coding or content generation should require an AI Contribution Note. High-stakes projects involving significant AI collaboration may warrant additional evidence, such as selected prompts, version histories, error logs or oral defense.

Source: EY-Parthenon analysis

The goal is not to avoid AI, but to enable students to think, explain and perform beyond it

Importantly, disclosure should not legitimize prohibited uses nor should students be penalized for transparently declaring permitted use. If disclosure is treated as evidence of misconduct, institutions risk incentivizing concealment rather than fostering academic integrity. Process evidence should also form part of student portfolios.

In an AI-enabled labor market, employers are expected to increasingly value not only the final output but also the ability to demonstrate how a problem was framed, how AI was used, how errors were identified and addressed and how a defensible conclusion was reached.

7.1.5 Protect independent capability through deliberate "cognitive repetitions"

Students should periodically practice core competencies without AI assistance. The objective is not to preserve pre-AI approaches to education, but to equip students to exercise independent judgment, recognize errors and perform effectively when AI support is unavailable or inappropriate.

Such practices may include:

- Summarizing a reading from memory before consulting AI
- Independently drafting an argument or code structure
- Explaining completed work orally without notes
- Solving representative problems under time constraints
- Identifying the strongest counterargument without first relying on the model

These activities help reinforce foundational capabilities while enabling students to use AI as a complement to, rather than a substitute for, critical thinking and domain knowledge.

A simple self-test is: **Could I explain, defend or reproduce the essential reasoning if the tool disappeared tomorrow?** A negative answer does not mean AI use was wrong; it means the learning process is incomplete.

Dependency signals include:

- Inability to begin without opening an AI tool
- Discomfort with one's own writing or reasoning
- Accepting answers because the model sounds more confident
- Inability to detect obvious errors
- Anxiety in AI-restricted assessments
- Avoiding readings, lectures or conversations because AI feels faster

Primary focus group discussions reveal a nuanced impact of AI on student wellbeing. While AI can reduce workload, improve efficiency and alleviate academic stress, it can also create anxiety about the depth of learning achieved and concerns that peers using AI gain an unfair advantage through increased speed and productivity. This, in turn, can create pressure to use AI even when students may prefer not to.

When asked to imagine what campus life would be like if ChatGPT and similar tools disappeared for a month, several students responded with a single word: "chaos." These responses suggest that responsible AI use is important not only for maintaining academic integrity but also for preserving students' confidence in their own abilities, judgment and capacity for independent learning.

As students increasingly use AI, they should protect sensitive information and have equitable access to AI learning opportunities

7.1.6 Protect personal data, confidential material and other people's rights

Students should be guided by a clear and consistent principle, personal data, confidential institutional information, identifiable research data, unpublished intellectual property, internship documents and sensitive information relating to others should never be uploaded to unapproved external AI tools.

To support responsible decision-making, institutions should distinguish between different categories of information:

- **Public information:** Generally suitable for processing using AI tools
- **Institutional or course-related material:** Subject to faculty guidance, licensing terms and copyright requirements
- **Personal or confidential information:** Permissible only within approved and appropriately governed environments
- **Highly sensitive information:** Prohibited from use in generative AI systems unless explicit institutional authorization has been granted

Students should also understand that free consumer AI tools and institutionally licensed systems may differ significantly in their contractual protections, data retention practices and administrative safeguards. At the same time, institutional licensing alone should not be assumed to guarantee that all forms of data use are appropriate or risk-free.

Guidance for research projects, internships and industry collaborations should explicitly address obligations relating to client confidentiality, intellectual property and proprietary information. An organization's general acceptance of AI use does not constitute permission to upload internal documents, confidential materials or commercially sensitive data to publicly accessible AI models.

7.1.7 Make peer learning an equity mechanism, not a substitute for institutional responsibility

Student societies, senior students, online creators and peer networks increasingly function as informal channels for AI learning and capability development. While these communities provide valuable support, foundational AI literacy should be underpinned by structured, institution-led initiatives to maintain consistency, quality and inclusivity. Such measures are essential to enabling the fair distribution of opportunities to benefit from AI across the student community.

One approach is the establishment of **peer AI clinics**, where trained student mentors support their peers in understanding approved AI tools, developing verification practices, leveraging multilingual and accessibility features, building discipline specific workflows, recognizing unsafe data practices and preparing for both AI permitted and AI restricted assessments. To maintain quality and accountability, peer mentors should be trained using institutional resources, operate under faculty or learning-design supervision and receive formal recognition or compensation where their contributions are substantial.

Without appropriate structure and oversight, peer-led programs may inadvertently reinforce the inequalities they are intended to reduce, with highly confident STEM students disproportionately shaping norms of responsible AI use for disciplines and learners with different needs and contexts. Accordingly, an equitable AI ecosystem should include an institutionally supported baseline tool or equivalent capability, mobile friendly and low bandwidth access options, computer-lab facilities for students without adequate personal devices, multilingual interfaces and learning resources, accessibility testing and assistive features, non-AI alternatives where AI use remains optional and safeguards against assignments that effectively require students to purchase premium AI subscriptions. Care essential to helping opportunities to benefit from AI be distributed fairly across the student community.

Key student playbook KPIs

7.1.8 Key student playbook KPIs

Coverage

Share of first-year students completing an assessed baseline AI literacy module, not merely an orientation

Policy clarity

Share of students able to correctly classify uses such as brainstorming, translation, first-draft generation

Verification capability

Performance on exercises requiring detection of fabricated sources, unsupported claims

Independent capability

Difference between comparable AI-assisted and secure independent performance

Process transparency

Share of relevant submissions with complete, proportionate AI contribution notes

Meaningful access

Share of students with an approved tool, a suitable device, reliable connectivity and basic support

Access parity

Gap in usage and capability between students with paid subscriptions and those relying on institutional or free access

Dependency and self-efficacy

Student-reported ability to begin, complete and explain core tasks without AI

Safety

Number and nature of incidents involving personal data, confidential material, fabricated citations or unsafe advice

Student voice

Share of AI policy reviews, tool evaluations and assessment-redesign processes with meaningful student participation

Source: EY-Parthenon analysis

As AI reshapes teaching and learning, institutions need to equip faculty to redesign courses

7.2 The faculty playbook

How to scale AI literacy and redesign courses and assessments without lowering academic standards

Faculty are not a homogeneous group of adopters or resistors. Their responses to AI are shaped by disciplinary context, workload, confidence, incentives and legitimate concerns about preserving foundational skills. Faculty do not need to become AI specialists, but they require sufficient understanding to set appropriate boundaries, design credible assessments, use AI responsibly and recognize when specialist support is needed.

Our primary research indicates that faculty are already using AI for lecture preparation, content development, literature discovery and administrative tasks, while remaining cautious about its role in high-stakes assessment. Many also report encountering students who submit work they are unable to explain. This practical experience should serve as the foundation for faculty AI capability-building efforts.

Faculty playbook

1. Start with the learning outcome and create an activity-level AI map
2. Build a program-level assessment architecture, not a collection of defensive assignments
3. Redesign assignments so that the process carries academic value
4. Establish clear rules for faculty use of AI
5. Replace generic workshops with discipline-specific course-redesign support
6. Align workload, incentives and recognition with the scale of change
7. Build equity and accessibility into course design

Source: EY-Parthenon analysis

As AI becomes embedded in learning and assessment, institutions should align AI use with learning outcomes and focus on verifiable evidence of student competence

7.2.1 Start with the learning outcome and create an activity-level AI map

Every course should include a concise AI learning and assessment map that links each major learning outcome to the appropriate role of AI in teaching, learning and assessment. Using the assessment framework outlined in Section 6, activities should be classified as AI-free, AI-assisted, AI-integrated and AI-directed.

For each learning activity, the map should clearly identify the capability being developed, the aspects that current AI systems could complete without demonstrating that capability, the forms of AI use that would strengthen learning, the evidence required to demonstrate genuine understanding and the alternatives available to students who require accommodation or lack access to AI tools.

AI expectations should be articulated at the assignment level, rather than relying solely on institution-wide policies. This aligns with emerging international practice, including guidance from Stanford University, which advocates course-specific AI policies that balance academic integrity, student success and teaching effectiveness. To minimize ambiguity, policies should include discipline-specific examples. Broad statements such as "AI may be used for editing" should be accompanied by clear guidance on whether editing extends to restructuring arguments, rewriting substantive content or making analytical changes.

7.2.2 Build a program-level assessment architecture, not a collection of defensive assignments

Rather than attempting to eliminate AI use in assessments, institutions should enable students to demonstrate achievement of learning outcomes through reliable and verifiable evidence. This shift in focus allows assessments to remain authentic while preserving confidence in student competence.

Each major learning outcome should be supported by two complementary forms of evidence:

- Secure evidence of learning: Supervised application, oral defense, live demonstrations, practical exercises, in-class reasoning or other forms of independently verified performance.
- Open evidence for learning: Authentic projects and assignments in which students may use appropriate tools, including AI, provided they document, justify and critically evaluate their use.

Implementation approaches need to be practical and scalable, taking into account large class sizes, affiliated college structures and limited faculty capacity. Feasible mechanisms may include:

- Short oral defenses for a sample of students, with random escalation where concerns arise
- Group presentations followed by individually directed questions
- Secure concept checks at key progression stages
- Laboratory demonstrations using unseen inputs or scenarios
- Version histories and staged assignment submissions
- Structured peer questioning and review
- Rotating viva panels supported by teaching assistants
- Program-level secure assessments, reducing the need to replicate secure evaluation tasks across every course

This approach shifts the focus from preventing AI use to enabling institutions to reliably verify learning outcomes while helping students develop the skills required to use AI responsibly and effectively in professional contexts.

Source: EY-Parthenon analysis



The challenge is no longer detecting AI use, but designing assessments where learning remains visible

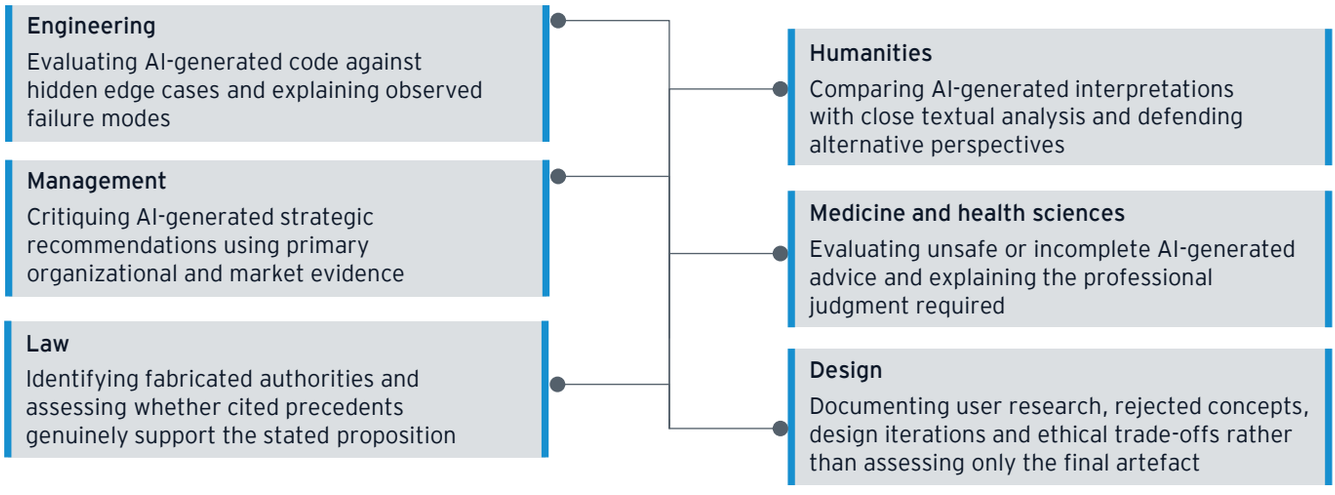
7.2.3 Redesign assignments so that the process carries academic value

Assignments should be structured so that they are inherently difficult to outsource, as student judgment, contextual understanding, iterative thinking and the ability to justify decisions remain central to successful completion.

Effective assessment approaches include:

- Diagnosing and correcting errors in AI-generated outputs
- Comparing multiple plausible responses and justifying the preferred solution
- Interpreting local, current or institution-specific data rather than simply summarizing information
- Integrating field observations, interviews, experiments or original data collection
- Assessing problem definition, evidence gathering, drafting and revision as distinct stages of the learning process
- Requiring documentation of revisions made in response to feedback
- Applying concepts to unfamiliar contexts
- Critically reflecting on where AI added value, where it failed and where it may have introduced bias
- Complementing submitted work with oral examinations or practical demonstrations

To keep assessment reform relevant across disciplines and prevent it from becoming overly STEM-focused, discipline-specific approaches are essential. Examples include:



Our primary research indicates that several faculty members are already experimenting with practices such as open AI assignments followed by oral examinations, surprise tests, debugging exercises, scenario-based questions, presentations and "fix rather than create" activities. These emerging practices represent promising examples of innovation within the Indian higher education context. However, further evaluation is needed to determine whether such approaches can be scaled institution-wide and embedded systematically, rather than remaining isolated initiatives led by individual innovators.

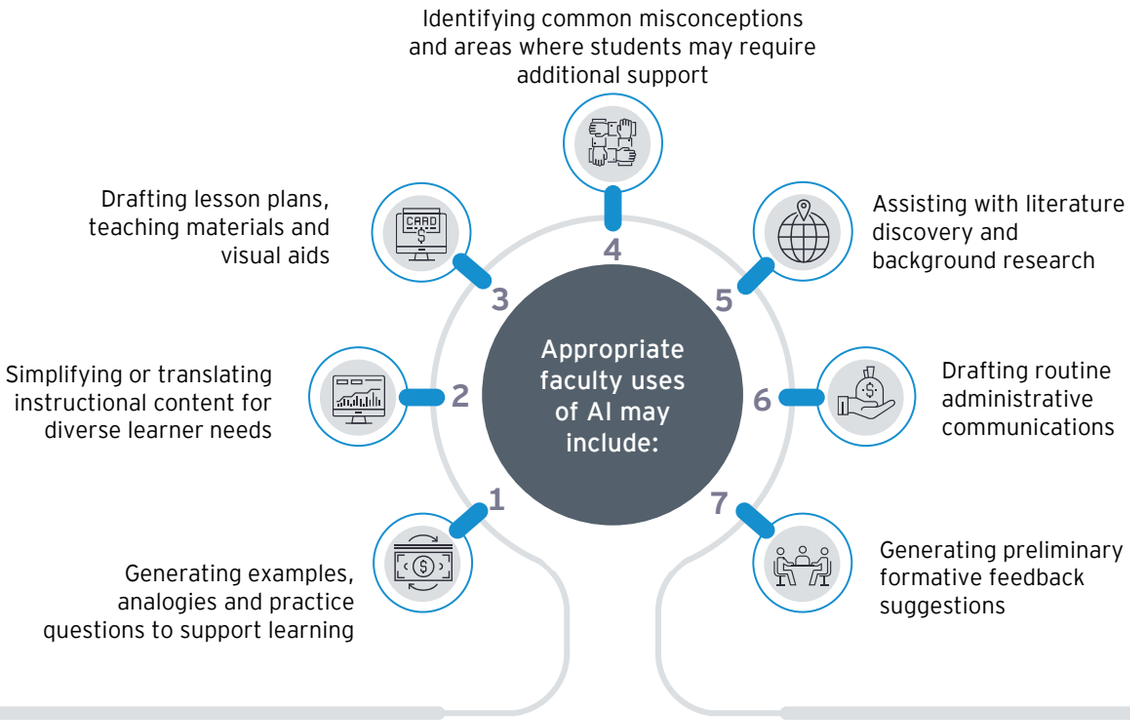
Primary interviews suggest that some faculty members are placing greater emphasis on direct interaction to make student thinking visible. Examples include asking students to explain their projects on a whiteboard, walk through how they approached the problem, justify key decisions and discuss how AI was used during the process. Interestingly, faculty noted that the rise of AI has made human interaction more important, not less, as direct engagement is often required to verify student understanding and reasoning. However, several faculty members noted that such approaches can be difficult to implement at scale in large classrooms without teaching assistants or additional support. As a result, the feasibility of these practices often depends on class size and institutional resources.

Source: EY-Parthenon analysis

The challenge is not whether faculty use AI, but whether its use is transparent, verifiable and accountable

7.2.4 Establish clear rules for faculty use of AI

Faculty use of AI should be governed by the same principles of transparency, verification and accountability expected of students. Institutions cannot credibly require students to disclose AI use while permitting its invisible use in the creation of course materials, feedback or assessment processes.



Faculty should disclose material AI use where it influences published teaching content, research outputs, assessment practices or student-facing decisions. Such disclosure need not be burdensome, however, faculty members should be able to explain, justify and take responsibility for outcomes informed by AI systems.

However, higher risk applications require stronger governance controls and safeguards. These include:

- Uploading identifiable student work to unapproved AI systems
- Generating high-stakes grades or assessment outcomes without meaningful human oversight
- Treating AI detection outputs as conclusive evidence of misconduct
- Providing consequential feedback without verifying it against the student's actual work
- Using AI to support decisions relating to student progression, disciplinary action or wellbeing interventions

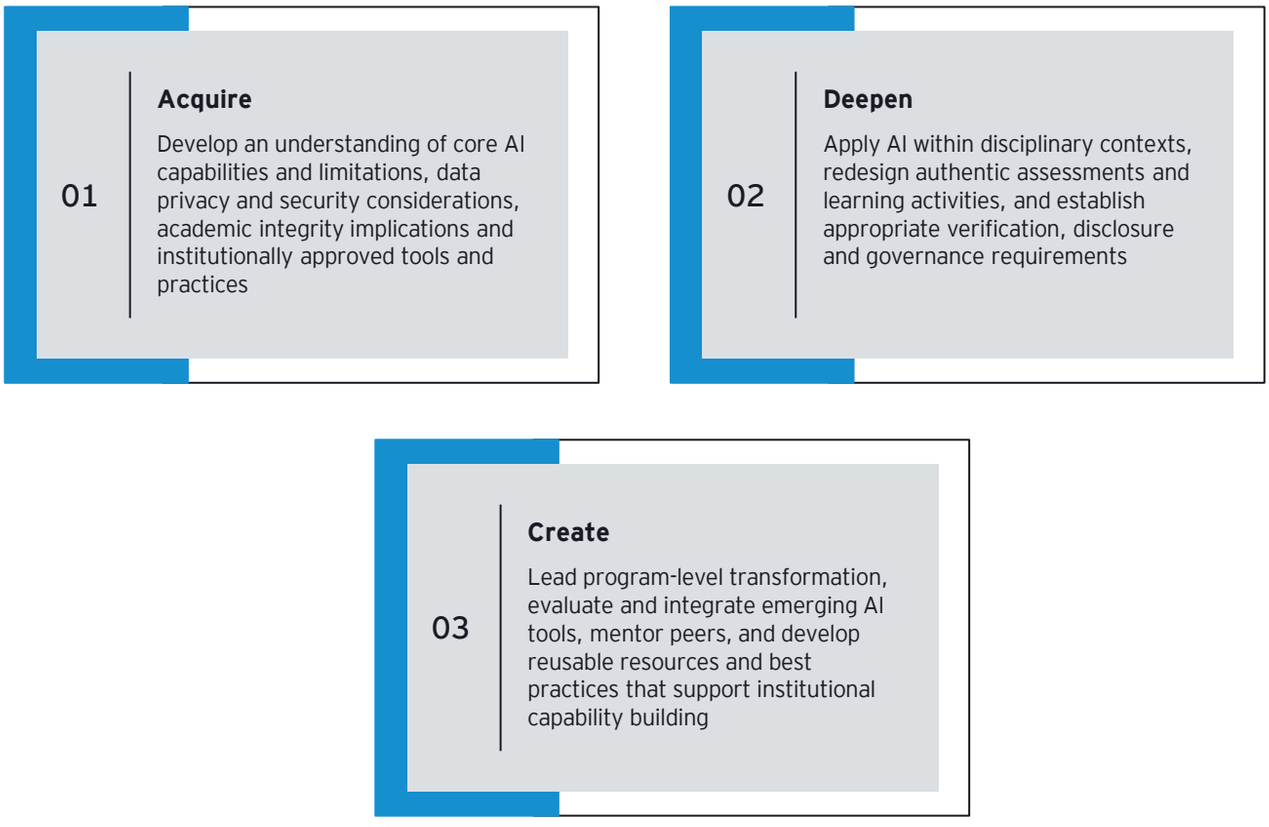
Our primary research indicates that faculty are generally willing to use AI to improve efficiency and reduce administrative burden, while remaining appropriately cautious in contexts involving interpretation, handwriting, subjectivity or high-stakes decision-making, where errors may be difficult to detect or rectify. A risk-based policy framework should recognize this caution as a governance strength rather than a barrier to adoption. By distinguishing between low-risk and high-risk applications, institutions can enable responsible innovation while maintaining academic standards, professional accountability and trust.

Source: EY-Parthenon analysis

As AI becomes embedded in teaching, institutions need to move beyond generic training to practical faculty capability development

7.2.5 Replace generic workshops with support to redesign courses for specific disciplines

Faculty knowledge development should progress through three levels:



Primary interviews consistently emphasized the need for hands-on, use-case-specific training, regular capability updates and stronger support for non-technical disciplines. As a result, workshop attendance alone should not be considered a valid indicator of AI readiness; readiness should be assessed through demonstrated changes in teaching practice, assessment design and responsible AI use.

Some primary interviews suggest that AI is enabling new forms of pedagogy rather than simply improving efficiency. Faculty reported using AI to create videos, presentations, visualizations and gamified learning activities that help students engage more deeply with course content and better understand complex concepts.

The learning support for faculty should prioritize:

- Short, practice-oriented sessions anchored in real teaching and learning challenges
- Discipline-specific examples that demonstrate relevant applications and limitations
- Course redesign workshops where faculty revise existing assessments, learning activities and feedback approaches
- Peer review mechanisms for evaluating course policies, assessment frameworks and grading rubrics
- Access to instructional designers, educational technologists and librarians to support implementation
- Dedicated office hours and experimentation spaces for testing and evaluating AI tools
- Communities of practice that sustain peer learning and knowledge sharing beyond individual workshops
- Curated repositories of approved, adaptable examples, templates and institutional resources

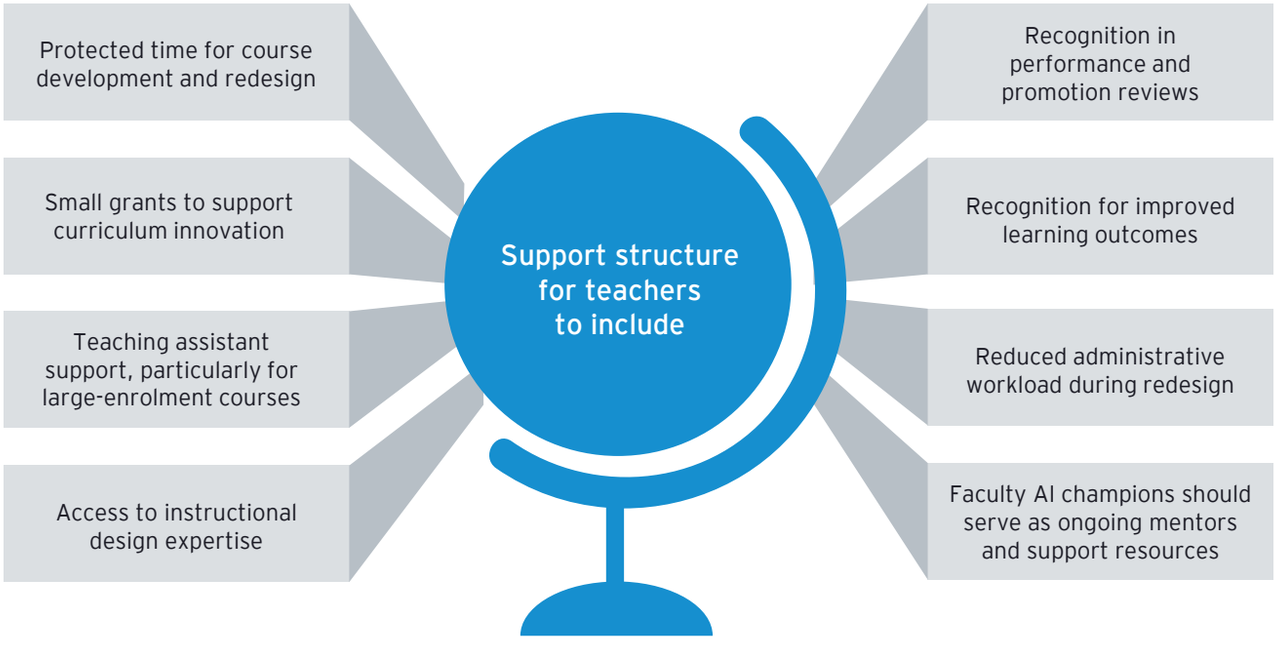
Source: EY-Parthenon analysis

As AI reshapes teaching and assessment, institutions need to recognize course redesign as a strategic academic contribution

Some leading institutions are already beginning to adopt this approach. Shiv Nadar University's 2026 Faculty Professional Development Programme combined hands-on AI training with a cross-disciplinary "Festival of Pedagogies", where faculty from English, economics, mathematics, chemistry, arts, computer science and engineering demonstrated how they were applying AI in teaching, assessment and student engagement. The model illustrates how faculty development can move beyond generic tool training towards peer-led sharing of classroom practices that can be adapted across disciplines.

7.2.6 Align workload, incentives and recognition with the scale of change

Course redesign extends far beyond learning how to use an AI tool. It requires faculty to revisit learning outcomes, assess evolving AI capabilities, redesign assignments, develop appropriate assessment rubrics, communicate expectations and usage policies, incorporate oral or practical validation mechanisms and continuously update course designs as technologies evolve. Institutions should recognize this work as a significant academic contribution rather than treating it as invisible labor.

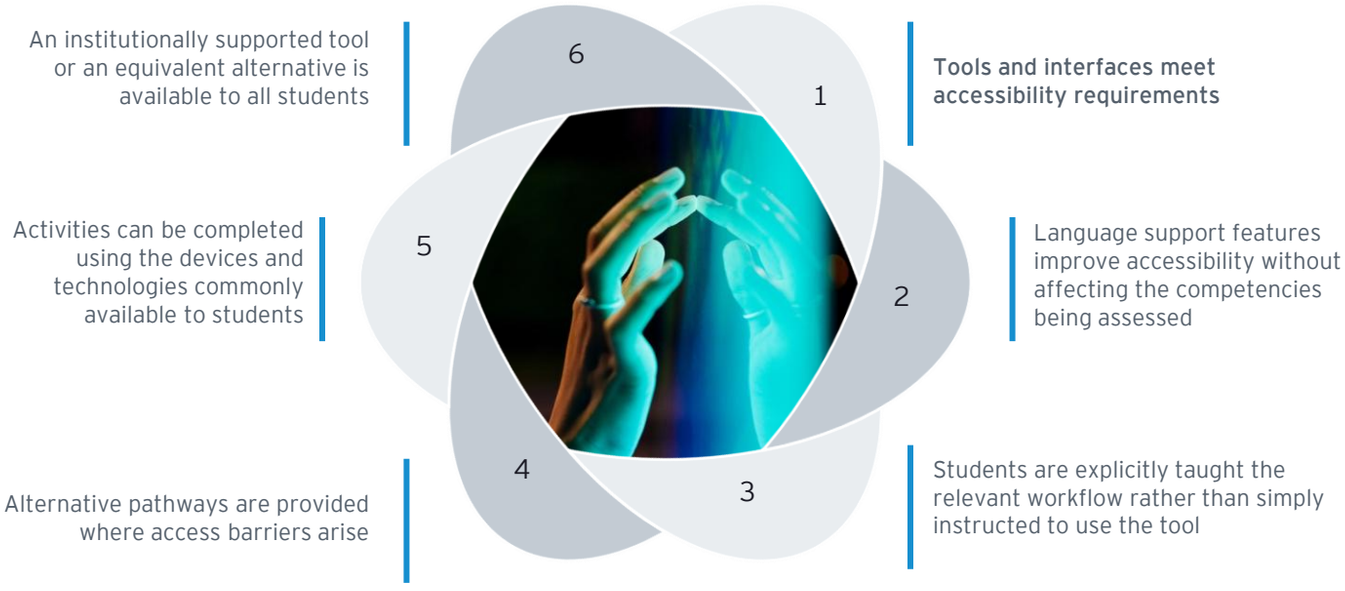


Institutional leadership should also foster a culture of psychological safety in which faculty feel comfortable acknowledging uncertainty. Superficial adoption often emerges when AI proficiency is publicly celebrated without creating space for open discussion of unresolved questions, limitations or areas where the implications of AI tools are not yet fully understood.

Source: EY-Parthenon analysis

While AI creates new opportunities for learning, institutions should address the risks of unequal access and inappropriate enforcement

Faculty should not assume that all students have access to laptops, premium AI subscriptions, reliable internet connectivity, equivalent English language proficiency or prior experience with AI tools. Before incorporating AI into learning activities, faculty should verify that:



Assessment policies should clearly distinguish between accessibility and language support functions and the use of AI for intellectual substitution. Tools such as speech to text, translation or text simplification that assist students in expressing independently developed ideas should not automatically be treated as equivalent to AI-generated content.

Faculty should also exercise caution when using AI detection systems, particularly in multilingual contexts where writing styles, translation support or standard academic conventions may increase the risk of false positives. Suspected misuse should prompt further academic review and evidence-based inquiry rather than automatic disciplinary action.

Source: EY-Parthenon analysis

From policy to practice: Metrics to track AI integration in teaching

7.2.7 Key faculty playbook KPIs

01 **Course policy coverage:** Proportion of active courses with clearly defined AI use policies aligned to learning outcomes

02 **Assessment alignment:** Proportion of programs that have redesigned assessment strategies to evaluate learning outcomes in AI-enabled environments

03 **Practice transformation:** Proportion of faculty who redesign at least one course component or assessment, evidenced through documented artefacts

04 **Faculty capability:** Performance on practical scenarios assessing responsible AI use, hallucination detection, data protection, disclosure practices and assessment design

05 **Student clarity:** Proportion of students reporting clear, consistent and understandable guidance on AI use, compared with faculty perceptions of the guidance provided

06 **Learning assurance**: Evidence that students can independently explain, justify or demonstrate competencies reflected in AI-assisted work

07 **Disciplinary reach:** Distribution of trained faculty, redesigned courses and approved AI use cases across disciplines, rather than concentration within computing and engineering fields

08 **Faculty workload impact:** Faculty time devoted to course redesign, assessment validation, academic integrity reviews and AI-related support, alongside institutional resources provided

09 **Equity of learning experience:** Differences in access, policy clarity, academic outcomes and support across language, device, disability and socioeconomic groups

10 **Academic integrity case quality:** Volume of AI-related misconduct cases, quality of supporting evidence, appeal outcomes and the proportion of cases primarily dependent on AI-detection scores

Source: EY-Parthenon analysis

Successful AI implementation depends on visibility, accountability and a risk proportionate governance model

7.3 THE ADMINISTRATOR PLAYBOOK

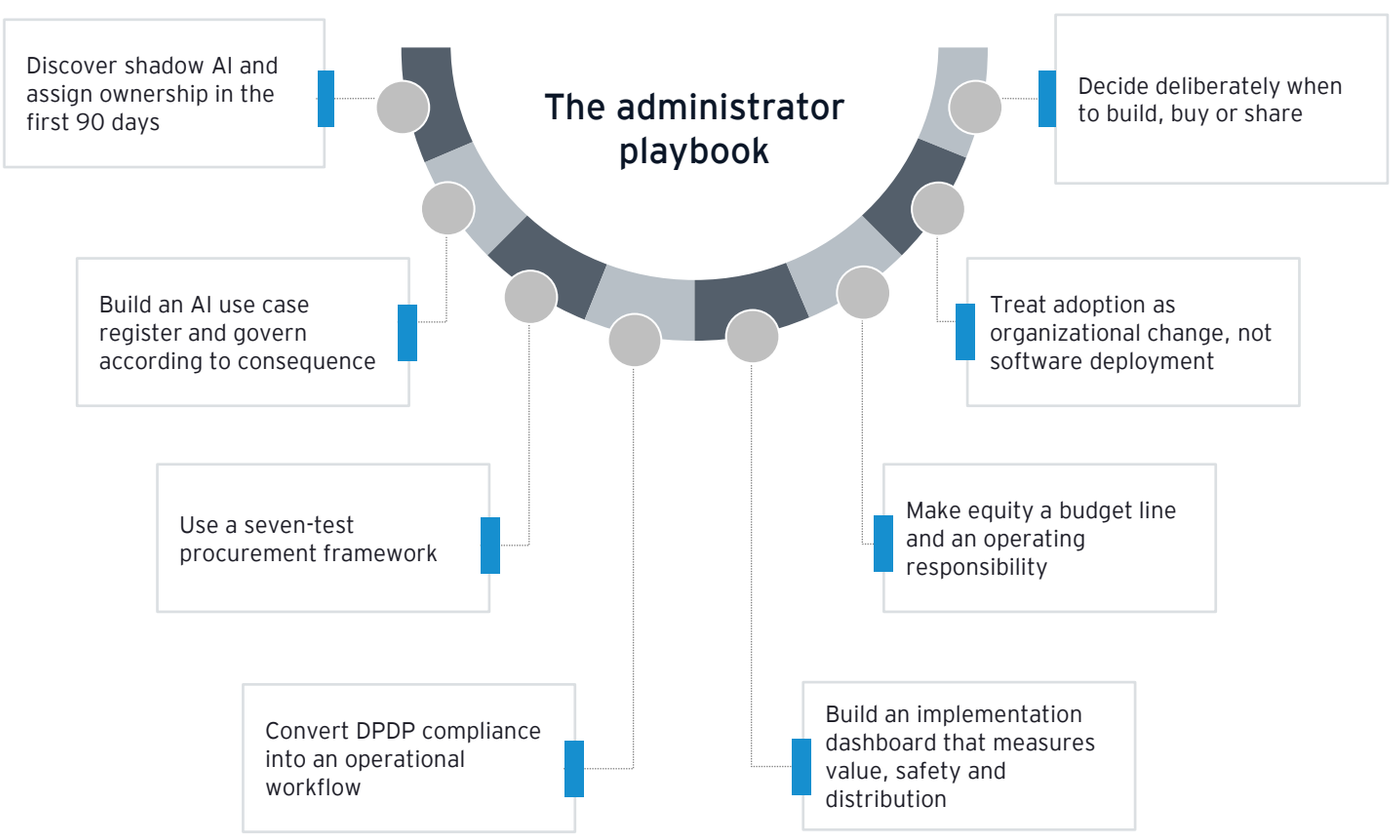
How to turn scattered experimentation into governed, affordable and equitable institutional implementation

Primary interviews indicate that AI is already being used across teaching, research, admissions, IT support, financial aid, placements, document creation and routine communication. However, the extent of institutional oversight varies considerably. While some leading universities have established approval processes, governance committees and department level budgets to guide AI adoption, others rely on ad hoc and decentralized initiatives. As a result, AI implementation often remains fragmented, with limited visibility into how tools are being selected, used and monitored across the institution.

The starting point is not a multi-year strategy document; rather, it is visibility, clear decision rights and a small number of governed implementations. Institutions cannot govern what they have not inventoried, evaluate what they have not defined, or scale what they have not measured.

The playbook outlines a three-tier capability model:

- **Baseline model:** A foundational readiness baseline applicable to all institutions
- **Medium and high-risk systems:** Additional governance and risk controls for medium and high-risk systems
- **Advanced capabilities:** That can be delivered through shared regional infrastructure, such as the AI Highway, rather than being independently developed by every institution.



Source: EY-Parthenon analysis

Effective AI governance begins with visibility, ownership and clear decision rights

7.3.1 The first 90 days: Discover shadow AI and assign ownership

Step 1: Institutions should begin with a rapid institutional inventory covering:

- Tools used by students, faculty and staff
- Departmental procurements, pilots and experimentation
- Teaching, research and administrative use cases
- Data entering each system
- Contractual arrangements and account types
- Designated owners or responsible functions
- Known incidents, complaints and operational limitations

This inventory should also include informal and individual usage. Even in the absence of enterprise AI contracts, institutions may face significant exposure through personal accounts, faculty uploads of student work, or the use of external AI platforms for institutional materials.

Step 2: A cross functional AI governance body should be established with clearly defined decision rights Membership should include academic leadership, IT and cybersecurity teams, legal and data protection specialists, faculty and student representatives, procurement, library and learning support services, accessibility and inclusion experts and placement or industry relations teams where relevant.

The objective is not to review every low-risk classroom experiment, but to establish a clear governance framework that defines:

- Activities that remain within faculty discretion
- Activities requiring departmental approval
- Use cases requiring institutional review for technical, privacy, or risk considerations
- Authority to suspend systems
- Responsibility for incident management
- Accountability for communicating changes to students and staff

Step 3: A formal responsibility matrix should identify accountable owners for AI strategy, tool approvals, data governance, assessment guidance, procurement, incident response, user support and periodic evaluation.

Primary interviews highlight a common set of challenges: unclear governance structures, case-by-case approval processes, strong institutional enthusiasm for AI adoption, but limited central guidance and support. Participants also emphasized the need for a common platform to address implementation questions and guide tool selection decisions. IT leaders further described the ongoing tension between enabling innovation and maintaining security, compliance and risk controls. **Collectively, these findings support the case for a focused 90-day governance sprint to establish foundational oversight and decision-making mechanisms.**

Source: EY-Parthenon analysis

Effective AI governance requires visibility into use cases and oversight proportionate to risk

7.3.2 Build an AI use case register and govern according to consequence

Institutions should record every institutional AI use case in a central register containing:

- The problem addressed
- Intended users and affected individuals
- Data inputs Vendor or model
- Risk classification
- Accountable owner
- Human review arrangements
- Known limitations
- Approval and review dates
- Performance and equity indicators
- Exit or suspension conditions

Use cases should be classified according to their potential consequences rather than the technology provider or model used. A risk-based approach enables governance effort to be directed where institutional impact is greatest.

The **NIST AI Risk Management Framework** provides a useful foundation for operationalizing this approach through its four core functions: **govern, map, measure and manage**. rather than adopting a generic technical framework in its entirety, institutions should adapt these principles to the higher education context, ensuring that governance processes reflect the academic, administrative and student facing risks associated with AI deployment.[78]

78. Source: NIST AI Risk Management Framework , EY-Parthenon analysis

AI procurement should be driven by institutional needs, evaluated through risk and designed for long term sustainability

7.3.3 Use a seven-test procurement framework

AI procurement should begin with a clearly defined institutional challenge rather than a vendor demonstration. Procurement decisions should be guided by evidence of educational or operational value, alignment with institutional priorities and the ability to manage risks effectively. Every proposed solution should be assessed against these seven criteria:

1 Educational and operational value

What specific problem will the tool address? What is the current baseline and what measurable improvement is expected in learning, administration, research, or student outcomes?

2 Data protection and privacy

What data is collected, retained, processed or transferred? Is institutional data used for model training? Can data retention be disabled and what third parties or sub-processors are involved?

3 Model quality and risk management

How has the system been evaluated for accuracy, hallucinations, bias, accessibility and performance across relevant languages, disciplines and user groups? What limitations or constraints are explicitly acknowledged by the provider?

4 Technical architecture and security

Does the solution support institutional identity management, role-based access controls, audit logging, secure integrations and appropriate segregation of institutional data?

5 Accessibility and equity

Is the solution accessible on commonly used devices and lower bandwidth networks? Does it support relevant Indian languages and assistive technologies? Can users access comparable functionality without requiring premium personal accounts?

6 Commercial sustainability

What is the total cost of ownership, including licensing, integration, training, support, monitoring and anticipated price changes? Where usage-based pricing applies, are future costs predictable and manageable?

7 Exit strategy and portability

Can the institution retrieve its data, prompts, configurations and outputs in a usable format? How easily can the institution migrate if the provider changes its pricing, service model, or underlying technology?

Procurement processes should rely on time-bound pilot implementations with clearly defined baselines, success metrics, user feedback mechanisms and termination criteria. Positive user perception alone is insufficient grounds for institution-wide deployment. Decisions should be supported by demonstrable improvements in educational, operational, or student outcomes.

Primary interviews indicate that licensed institutional arrangements are often perceived as offering stronger safeguards for confidentiality and data protection than free consumer applications. At the same time, many faculty members, staff and students acquire AI tools independently because institutional access remains fragmented or unavailable. These findings highlight the need for procurement strategies that provide a secure institutional baseline while reducing the risks associated with uncoordinated and uncontrolled tool adoption.

DPDP compliance should be embedded into everyday institutional processes rather than treated as a standalone regulatory exercise

7.3.4 Convert DPDP compliance into an operational workflow

Compliance with data protection requirements should be approached as a structured sequence of governance and operational activities. Institutions should:

- Map where personal data enters AI systems
- Define the purpose and lawful basis for processing
- Provide clear and accessible privacy notices
- Minimize data collection and processing to what is necessary
- Establish retention, deletion and archival policies
- Conduct due diligence on vendors and subcontractors
- Implement appropriate access controls
- Establish incident management and breach response procedures
- Enable applicable rights for individuals
- Maintain clear accountability and documentation mechanisms

For all medium- and high-risk AI applications, institutions should develop data flow diagrams that clearly document the movement of information between students or staff, institutional systems, vendors, cloud providers and downstream applications. Such mapping supports both regulatory compliance and effective risk management.



Institutions should deliberately choose what to build, buy or share based on strategic value, capability and risk

7.3.5 Decide deliberately when to build, buy or share

Institutions should evaluate AI investments through a **structured build buy share framework**, selecting the most appropriate model based on strategic importance, organizational capability, cost and risk.

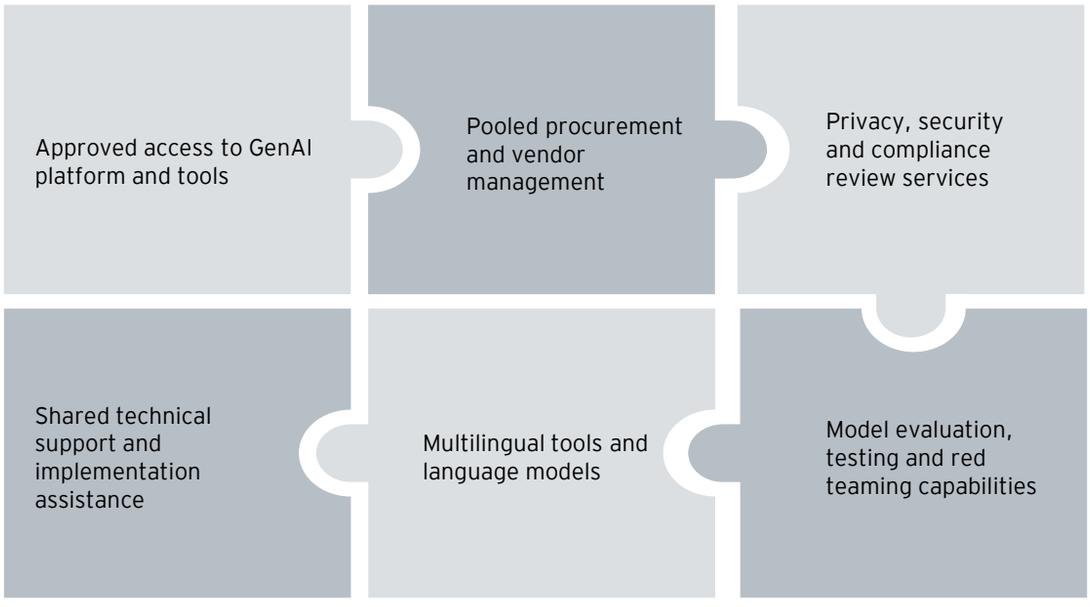
- **Buy:** Adopt mature third-party solutions where standard functionality, enterprise-grade security, vendor support and systems integration can be delivered more efficiently than through internal development
- **Build:** Develop solutions internally where technical capability exists, the use case is strategically important, sensitive data requires greater control, or local customization can substantially enhance effectiveness and user outcomes
- **Share:** Pool capabilities across institutions where common needs exist but individual institutions lack the resources to sustain specialized talent, computing infrastructure, evaluation functions, or procurement processes independently

Build decisions should consider the full lifecycle of ownership, including maintenance, security updates, model improvements, documentation, governance and workforce requirements, rather than focusing solely on initial development costs. Primary interviews indicate that several technically capable institutions are already building workflow automation solutions and administrative AI tools. However, many continue to face challenges in recruiting, developing and retaining specialized AI expertise within institutional IT functions.

Ashoka University recently used AI to audit, extract and validate key document information. This process guarantees data veracity, ensuring financial aid is disbursed to deserving candidates fairly and equitably. The initiative reduces manual effort, improves processing accuracy and features built-in data privacy and risk considerations.

> *"The challenge is deciding what to build ourselves and what to buy. Not every institution has the specialized AI skills needed in-house, so partnerships with external vendors often become necessary."* ***- University Leader, Primary Interview***

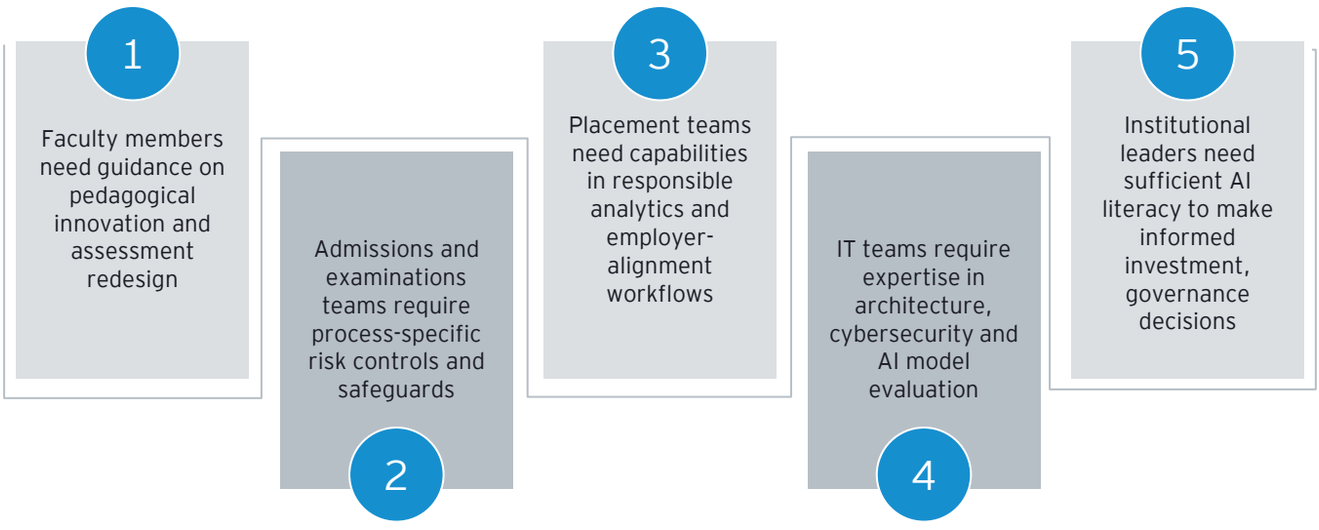
The AI Highway, proposed in Section 5, provides a mechanism for making the share model viable at scale. Shared services could include:



AI transformation requires capability building, leadership engagement and continuous communication

7.3.6 Treat adoption as organizational change, not software deployment

Different stakeholder groups require tailored capacity-building and support mechanisms.



To enable consistent adoption and governance, institutions should establish a comprehensive support ecosystem comprising:

- An approved AI tool catalogue
- Concise use case guidance
- Dedicated support channels such as office hours or help desks
- Role specific training programs
- Departmental AI champions
- Incident reporting mechanisms
- Regular communications to reflect evolving technologies, policies and institutional priorities

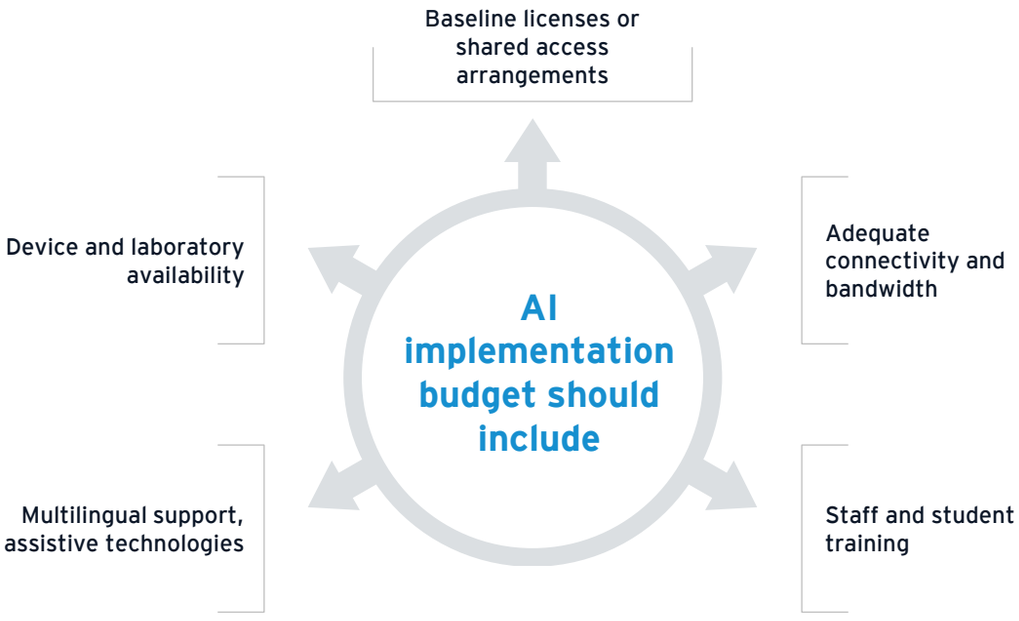
Leadership development programs should incorporate reverse mentoring and structured engagement with student and staff AI usage. As highlighted in a primary leadership interview, senior leaders may benefit from insights provided by younger or more technologically immersed colleagues to better understand evolving behaviors, expectations and patterns of AI adoption. Formalizing these exchanges transforms them from informal practices into an institutional governance mechanism that supports informed decision-making.

Communication should be treated as an ongoing implementation function rather than a one-time policy exercise. Publishing a policy alone is insufficient. Institutions must continuously communicate what has changed, the rationale behind those changes, the actions expected of different stakeholders, the support available to them, and the channels through which decisions, concerns or challenges can be raised and addressed.

AI equity requires dedicated resources, ongoing monitoring and institutional accountability

7.3.7 Make equity a budget line and an operating responsibility

Every AI implementation budget should explicitly account for the cost of equitable participation. This includes



Equity cannot be achieved by providing identical access credentials to all users; institutions should monitor patterns of activation, usage, support requests, disengagement and realized benefits to identify and address disparities in adoption and outcomes.

Institutions should also take deliberate steps to prevent persistent inequalities between departments. While early experimentation may be concentrated in well-resourced units, successful workflows, training approaches and resources should be systematically transferred across the institution to establish broader capability development and adoption.

Student-led innovation and peer support can play an important role in accelerating AI adoption. However, they should complement rather than substitute formal institutional support. Primary evidence of student-led workshops for staff, occurring alongside limited structured support mechanisms, highlights both the potential of collaborative learning and the governance gaps that institutions must address.

Shared access models and regional support structures should prioritize affiliated colleges, smaller state institutions and geographically underserved campuses. The funding mechanisms proposed in Section 6 should support these efforts, while the administrator playbook should provide practical guidance on how institutions can access and implement available resources effectively.

Institutions should measure AI success through outcomes, risk management and equitable benefit realization

7.3.8 Build an implementation dashboard that measures value, safety and distribution

The dashboard should contain a small set of decision useful indicators rather than long activity lists.

Table 2:

Indicators	Actions
Coverage and adoption	▪ Activation and sustained monthly usage of approved AI tools ▪ Patterns of AI adoption across user roles, academic and administrative departments and student segments ▪ The proportion of institutionally relevant AI activity conducted within approved, secure and governed environments
Operational value	▪ Improvements in turnaround time, backlog, error rates and staff effort ▪ Verified time savings and their redeployment to higher value activities ▪ Use case specific satisfaction and outcome measures, rather than general perceptions of AI
Learning and academic practice	▪ Courses with clearly defined AI use policies and expectations ▪ Assessments redesigned to reflect AI-enabled learning environments ▪ Mechanisms to verify student authorship and capability ▪ Independent evidence demonstrating genuine learning and competency attainment
Risk and governance	▪ Proportion of medium and high-risk AI systems formally registered, assessed and reviewed ▪ Data privacy, security, or governance incidents ▪ Outstanding complaints ▪ Consequential decisions subject to appropriate human oversight and review ▪ Time required to suspend, remediate, or correct identified system failures or risks
Commercial sustainability	▪ Total cost per active user and per successful use case ▪ license utilization rates and unused licenses ▪ Unexpected compute, integration, or operational costs ▪ Vendor concentration and dependency risks ▪ Switching, migration and exit costs
Equity	▪ Gaps in AI access, adoption and capability across user groups ▪ Participation through mobile only and low bandwidth channels ▪ Multilingual usage patterns and performance outcomes ▪ Accessibility barriers and compliance issues ▪ Distribution of benefits across departments, campuses and user segments

AI adoption should progress through phased implementation, measurable outcomes and continuous improvement

7.3.9 A phased institutional implementation sequence

The dashboard should contain a small set of decision useful indicators rather than long activity lists.

Table 3:

Timeline	Actions
Days 0 - 90 **Establish visibility and minimum control**	▪ Conduct an institution-wide inventory of AI tools, systems and use cases ▪ Establish clear governance structures, accountability mechanisms and decision rights ▪ Issue interim guidance on data management, privacy and acceptable AI use ▪ Identify and prioritize the oversight of medium- and high-risk AI systems ▪ Provide a clearly defined channel for user support, enquiries and incident reporting ▪ Select a limited number of high-priority pilot initiatives to demonstrate value and inform scale up
Months 3 - 6 **Build the minimum operating system**	▪ Establish and maintain an institutional AI register ▪ Provide access to approved AI tools and platforms ▪ Complete initial data mapping and classification exercises ▪ Implement standardized procurement, risk assessment and impact assessment frameworks ▪ Deliver targeted training for priority user groups ▪ Initiate course redesign and assessment transformation efforts ▪ Establish baseline indicators for equity, access and AI capability development
Months 6 – 12 **Demonstrate value and correct weaknesses**	▪ Evaluate pilot initiatives against predefined learning, operational, risk and equity outcomes ▪ Scale only those solutions that meet established performance thresholds ▪ Retire underperforming, redundant, or low-value tools ▪ Publish an institutional AI readiness and progress report ▪ Identify capability gaps requiring regional, sectoral or national support ▪ Integrate lessons learned into budgeting, accreditation planning and the Section 6 AI Readiness Index

Measuring institutional readiness to govern, scale and manage AI responsibly

7.3.10 Key administrator playbook KPIs

01 **Governance:** Share of material use cases with a named owner, risk classification and current review date

02 **Visibility:** Estimated share of institutional AI use captured in the tool and use case register

03 **Data readiness:** Share of medium- and high-risk systems with completed data maps, vendor review and retention rules

04 **Human accountability:** Share of consequential AI influenced decisions receiving competent human review and offering an accessible appeal route

05 **Procurement quality:** Share of contracts meeting standard clauses on data use, security, auditability, accessibility, product changes and exit

06 **Operational value:** Verified change in turnaround time, error rates and staff effort for each scaled administrative use case

07 **Support effectiveness:** Resolution time, repeat queries, training completion and demonstrated capability after support

08 **Cost sustainability:** Total expenditure per active user and validated use case, including licenses, compute, integration, security, support and monitoring

09 **Equity:** Gaps in access, use and benefit across campuses, departments, student groups and staff categories

10 **Risk:** Data incidents, inaccurate outputs reaching users, suspended systems, grievances and time to resolution

Employers must move beyond hiring talent to shaping the capabilities needed in an AI-enabled workforce

7.4 The employer playbook

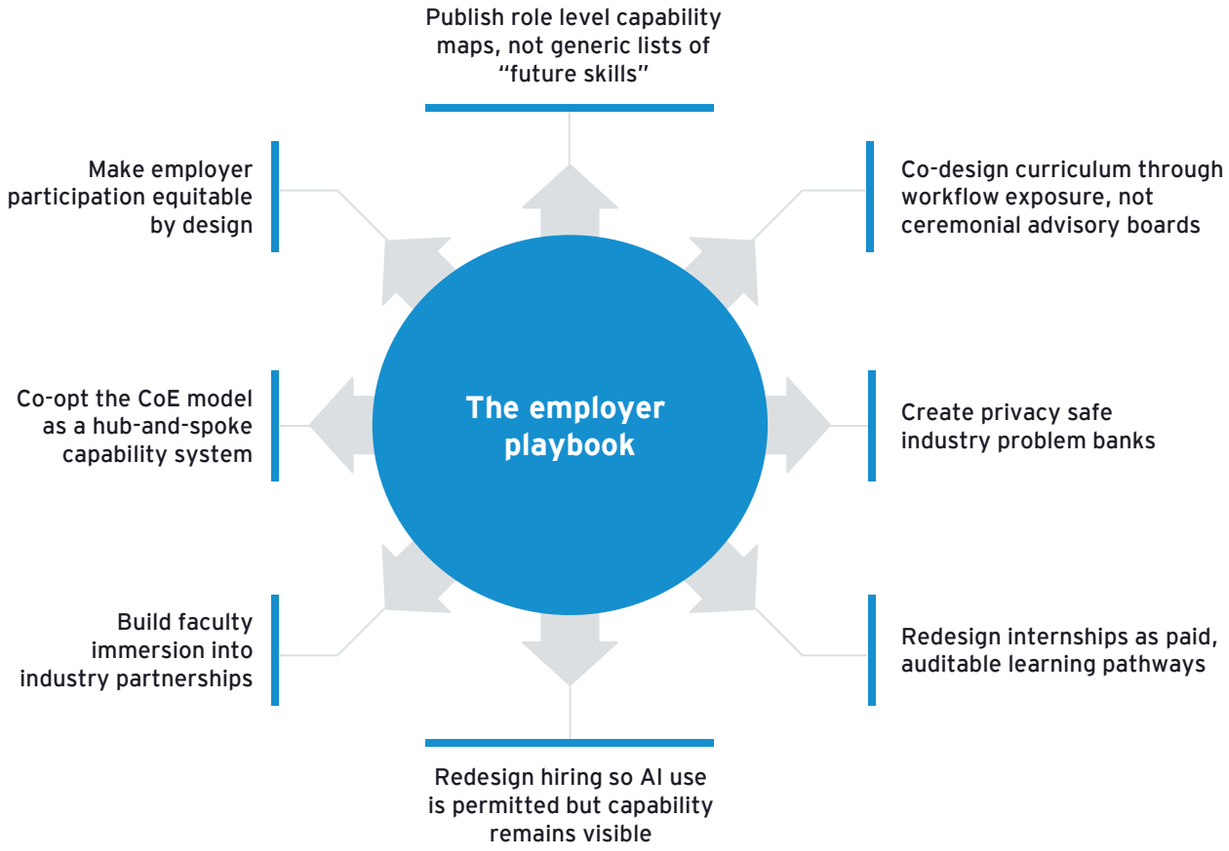
How can industry shift its role from being a critic or recruiter to becoming a co-developer of AI capabilities?

The Employer Playbook directly addresses the challenges identified in Section 3. Employers cannot reasonably expect higher education institutions to produce job-ready, AI-capable graduates while limiting their own involvement to placement presentations, guest lectures, advisory-board participation and end-stage recruitment.

At the same time, employers should not seek to shape curricula around short-term technology or tool requirements. Their role should instead be to make changes in workplace practices visible, contribute authentic business problems and feedback, expose faculty to evolving industry workflows, create meaningful work-based learning opportunities and help validate student capabilities against real-world expectations.

Employer engagement should be structured across four complementary functions:

- Defining the capabilities and competencies required for specific roles and occupations
- Providing authentic workplace experiences through which students can develop and demonstrate these capabilities
- Exposing faculty to emerging technologies, industry practices and changing workflows
- Redesigning recruitment and assessment processes to recognize and reward demonstrable capability rather than credential attainment alone



As AI reshapes jobs, employers must clearly define role-level capabilities and continuously communicate evolving workforce expectations to institutions

7.4.1 Publish role-level capability maps, not generic lists of "future skills"

Employers should clearly articulate how AI is reshaping specific entry-level roles. This requires identifying tasks that are increasingly automated or AI-assisted; capabilities that remain essential for foundational skill development; emerging requirements for supervision and verification; decisions that continue to require domain expertise and professional judgment; relevant confidentiality, data and regulatory considerations; and the forms of evidence that demonstrate an individual's ability to perform effectively in the role.

Capability frameworks should distinguish among three levels of proficiency: baseline AI literacy expected of all graduates, role-specific AI fluency required for AI-augmented professional work, and specialized technical capabilities needed for roles in AI, data, model development and product management.

The urgency of this effort is underscored by the World Economic Forum's *Future of Jobs Report 2025*, which projects that 39% of workers' core skills will change by 2030. While AI and big data rank among the fastest-growing skill areas, analytical thinking, creativity, resilience, and collaboration remain critical. This reinforces the need for AI capabilities to augment, rather than replace, human judgment and expertise.

Employers should regularly review and update these capability frameworks and communicate evolving requirements to partner institutions. At present, industry feedback is often transmitted informally through recruiters or placement teams and may influence only subsequent recruitment cycles. A more structured and continuous mechanism would allow institutions to align curricula, learning experiences and assessment approaches with changing workforce requirements without requiring frequent program-wide redesign.

Industry partnerships should be built around authentic workflows and measurable learning outcomes, not symbolic engagement

7.4.2 Co-design curriculum through workflow exposure, not ceremonial advisory boards

Industry engagement should focus on authentic AI-enabled workplace workflows and real-world application including



- **Employers should be engaged as strategic partners in curriculum development,** bringing current workplace practices, role-specific case studies, industry projects, co-delivered learning experiences and structured feedback on student work into academic programs. Their involvement can help institutions identify emerging skill requirements, understand common AI-related challenges encountered in professional settings, validate program outcomes and confirm that learning remains relevant to evolving workforce needs.
- **Industry engagement should translate into tangible improvements in teaching and learning.** Effective partnerships should contribute to curriculum updates, authentic project-based assessments, refined competency frameworks, industry-informed evaluation criteria, faculty immersion opportunities and stronger pathways to internships, live projects and graduate employment. The value of industry collaboration should therefore be assessed through its impact on learning outcomes rather than the number of engagement activities conducted.
- At the same time**, academic institutions must retain responsibility for learning outcomes, curriculum design, pedagogy and assessment.** The aim is not to tailor programs to the immediate needs of individual employers or technology providers, but to equip students with enduring capabilities, critical thinking skills and professional judgment that remain relevant across changing workplace contexts.
- **Institutions should also adopt a more outcome-oriented approach to evaluating partnerships.** Activities such as memoranda of understanding, guest lectures, advisory boards and employer interactions should be recognized as enabling mechanisms rather than measures of success in themselves. Partnerships should be reviewed periodically against clear indicators, including curriculum enhancements, student participation in industry projects, faculty capability development, employer feedback, internship conversion rates and graduate outcomes. Focusing on these measures helps distinguish collaborations that meaningfully strengthen educational quality from those that have limited impact on student learning and employability.

Industry problem banks can bridge the gap between classroom learning and workplace practice

7.4.3 Create privacy-safe industry problem banks

In hackathons, case competitions, industry challenges, capstone projects and other forms of applied assessment, institutions face a recurring tension: providing students with authentic, industry-relevant problems while protecting confidential data, intellectual property and commercially sensitive information. To address this, institutions and employers should collaborate to create privacy safe industry problem banks that enable students to work on realistic scenarios without exposing proprietary information.

Employers should contribute authentic business challenges that students can engage with over an extended period. Each problem bank entry should include a clearly defined operational or strategic challenge, relevant contextual information, access to an anonymized, synthetic, or appropriately governed dataset, guidance on permitted AI tools and data use restrictions, expected deliverables, assessment criteria, access to industry mentorship, and an explicit link to professional practice. Project-based learning should assess not only technical outputs but also the underlying capabilities required for effective AI-enabled work. throughout the project lifecycle, alongside the quality of the final deliverable.



Problem banks should be designed to support institutions with varying levels of capability and resources. The same challenge can be adapted across different levels of complexity, enabling advanced model development in some institutions and application, evaluation, or implementation-focused work in others.

Assessment should make the process of work visible, rather than focusing solely on final outcomes. Employer evaluation should consider how students defined the problem, allocated responsibilities, incorporated AI tools, identified and addressed errors, and responded to feedback throughout the project lifecycle, alongside the quality of the final deliverable.

Internships should be structured as equitable learning experiences that make capability development visible

7.4.4 Redesign internships as paid, auditable learning pathways

High-quality internships should be designed around clearly defined learning and capability outcomes. Each internship should specify the competencies to be developed, the nature of the work assignments, expected deliverables, AI and data-governance requirements, supervision and feedback mechanisms, evaluation criteria and the evidence students may appropriately retain in a professional portfolio.

Internships should not become environments in which AI-generated work is submitted without adequate oversight or learning. Effective supervision should examine the extent of AI assistance used, the processes followed to verify outputs, the errors identified and addressed, the decisions that required human judgment and the capabilities students developed independently of AI tools.

Paid internships should be treated as a matter of equity as well as employability. Unpaid or poorly structured opportunities disproportionately benefit students who can absorb the associated financial and opportunity costs.

Employers and higher education institutions should therefore evaluate internship quality, not merely placement volume. Completion certificates alone do not demonstrate meaningful work-based learning. Substantive assignments, structured supervision, regular feedback and demonstrable capability development are more credible indicators of internship effectiveness.

7.4.5 Redesign hiring so AI use is permitted, but capability remains visible

Hiring processes should be designed on the assumption that candidates may use AI in unsupervised settings. Prohibiting AI without credible mechanisms for supervision or verification is likely to encourage concealment while generating limited evidence of genuine capability.

A more effective approach would combine authentic work samples with structured validation. This could include AI-permitted take-home assignments, disclosure of how AI was used, live questioning on assumptions and evidence, evaluation of deliberately flawed AI outputs, supervised or oral assessments, portfolio reviews, and reflection on issues such as confidentiality, bias, and appropriate AI use.

This shift is consistent with broader labor-market trends. The NACE Job Outlook 2026 survey found that 70% of participating employers used skills-based hiring practices, while reliance on GPA screening had declined. Although these findings are drawn from the United States, they reinforce a wider movement towards assessing demonstrated capability alongside formal credentials.

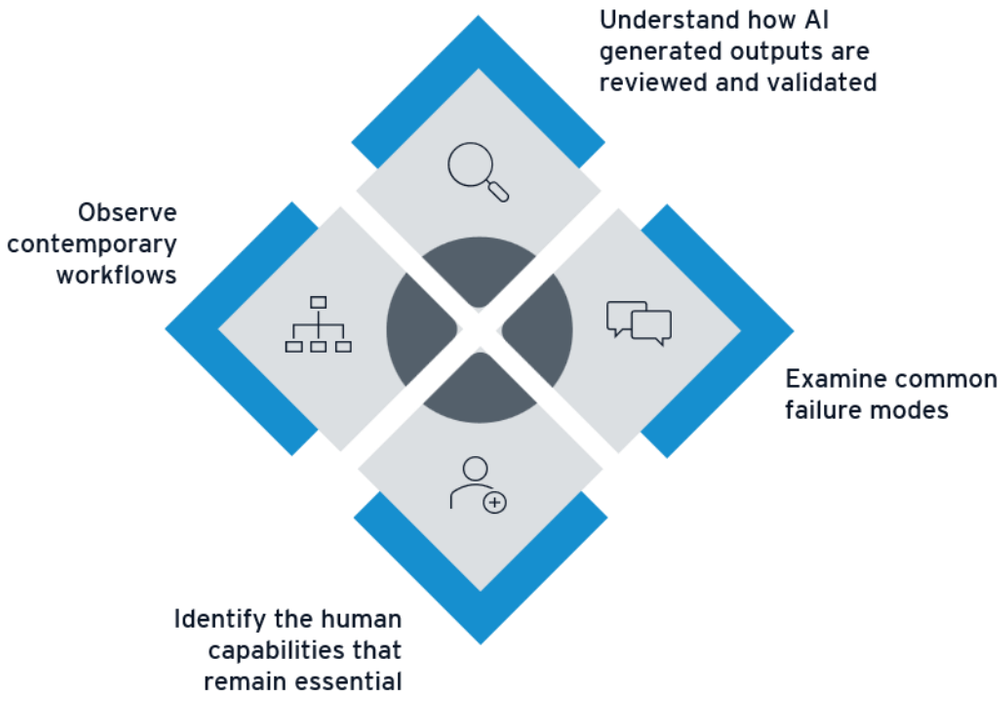
Primary interviews with students and placement stakeholders suggest that resumes, project descriptions and interview preparation are already extensively AI-assisted. Some employers are also embedding AI into recruitment and screening processes. These developments increase the importance of work-sample validation, follow-up questioning and capability verification, rather than excluding candidates solely because their applications were AI-assisted.

Employers should communicate assessment expectations transparently. When candidates understand that their process, judgment, verification and ability to explain their decisions will be evaluated, they are more likely to use AI to strengthen capability rather than conceal gaps.

Faculty need direct exposure to AI-enabled workflows to keep teaching, assessment and curriculum relevant

7.4.6 Build faculty immersion into industry partnerships

Faculty require direct exposure to how AI is changing professional practice. Employers should provide structured immersion opportunities that allow faculty to observe contemporary workflows, understand how AI-generated outputs are reviewed and validated, examine common failure modes, identify the human capabilities that remain essential, and translate these insights into improvements in curricula, pedagogy and assessment.



Each immersion should produce a tangible institutional output, such as a revised case study, workflow, assessment, module, or faculty-development program. This would establish that the insights gained are systematically embedded in teaching and learning practices.

Industry experts can also add significant value by co-teaching the applied components of academic programs. Institutions should, however, avoid building partnerships only around individual volunteers. Effective collaboration requires designated employer sponsors, committed mentoring capacity and mechanisms that maintain continuity beyond a single academic term.

Faculty immersion should extend across disciplines rather than remaining confined to technical fields. AI is changing professional practice in finance, marketing, law, design, healthcare, operations, research, human resources and administration. Restricting industry engagement to computing and data science would reinforce the disciplinary disparities identified in Section 2 rather than support institution-wide AI readiness.

Shared CoEs can accelerate talent development, faculty capability and applied innovation

7.4.7 Co-opt the CoE model as a hub-and-spoke capability system

The Employer Playbook translates the Centre of Excellence model described in Section 3 into a practical framework for sustained collaboration between industry and higher education institutions.

Rather than operating as isolated innovation hubs, AI-focused CoEs should function as long-term capability-building platforms that connect teaching, research, workforce development and industry engagement. At a minimum, they should support four core functions:

- Curriculum and assessment co-design
- Applied research and problem-solving
- Faculty capability development
- Identification and development of future talent

This approach aligns with the Government of India's broader AI agenda, including the establishment of AI Centres of Excellence in healthcare, agriculture and sustainable cities. Within higher education, however, CoEs should serve a wider purpose. They should not only advance innovation within host institutions but also strengthen AI capability across the broader education ecosystem.

A hub-and-spoke model can help scale impact while enabling participation from institutions with different levels of maturity and resources. Leading universities could serve as hubs for research, evaluation and faculty development. Regional universities could drive domain-specific applications and knowledge sharing, while Tier-2 and Tier-3 institutions could act as centers for implementation, testing and talent development. Industry partners could contribute real-world problems, mentorship, datasets, funding and recruitment opportunities. Such a model would create multiple pathways for participation rather than concentrating opportunities within a small group of institutions.

For CoEs to become sustainable institutional assets, participation must be formally recognized and rewarded. Student involvement can be integrated into academic credit structures, while faculty contributions should be reflected in workload allocation, professional recognition, performance evaluation and career progression.

CoEs should not operate as niche extracurricular initiatives accessible only to highly motivated participants. They should be embedded within mainstream teaching, learning and workforce development activities.

Industry engagement should also be structured around clear and measurable commitments. Employers can contribute authentic business problems, provide mentorship and faculty-immersion opportunities, offer paid internships, participate in the evaluation of student work and create employment pathways for successful participants.

Importantly, employers should be encouraged to engage with a diverse range of institutions, including those outside their traditional recruitment networks. This would broaden both access to opportunity and the pool of talent available to industry.

Finally, CoEs should promote greater consistency and transparency in how talent is identified and assessed. Shared assessment rubrics, externally reviewed portfolios, common evaluation frameworks and industry-validated competency standards can help employers recognize high-potential students across a wider range of institutions. Over time, this could expand opportunities, strengthen equity across the talent ecosystem and move employer engagement beyond recruitment from a limited group of institutions.

Expanding opportunity requires industry partnerships that reach beyond traditional talent pipelines

7.4.8 Make employer participation equitable by design

Industry partnerships should not remain concentrated within nationally prominent institutions simply because they are easier to administer. Equitable participation requires deliberate mechanisms, including regional project cohorts, remote mentoring models, shared problem banks, common capability frameworks, paid internships, open portfolio-based challenges, travel and connectivity support and explicit participation targets for Tier-2 and Tier-3 institutions.

Broadening engagement is not solely an equity imperative. It also provides employers with access to overlooked talent pools, reduces dependence on a limited set of institutions, strengthens geographic diversity in hiring and creates opportunities to influence workforce preparation before recruitment begins.

To enable these opportunities to deliver meaningful outcomes, remote engagement should complement, rather than replace, high-quality experiential learning. Programs should include clearly defined work assignments, structured supervision, regular feedback, and credible pathways to internships, employment and career progression.

Employers should also assess whether AI-enabled recruitment processes create unintended disadvantages for candidates from particular institutions, linguistic backgrounds, or those with less polished digital profiles. Where AI is used in screening or shortlisting, human oversight and periodic bias audits remain essential to enable fairness, transparency, and accountability.

Equity should ultimately be evaluated not by the number of institutions participating but by the extent to which students from diverse backgrounds gain access to high-quality opportunities and the outcomes those opportunities produce in terms of learning, employability and career progression.

Key employer playbook KPIs

7.4.9 Key employer playbook KPIs

Active partnership ratio

Partnerships producing assessed student work, faculty development or hiring outcomes as a share of AI-related MoUs

Capability map coverage

Share of priority entry-level roles with defined AI-assisted tasks, human judgment requirements, risks and evidence of competence

Problem bank contribution

Number of current, usable and privacy safe industry problems, including reuse across institutions and disciplines

Mentor intensity

Mentor hours per student team and completion of agreed feedback milestones

Faculty immersion

Number completing structured exposure and share translating it into a revised course, case or assessment

Faculty and student feedback

Whether employer contributions are current, educationally useful and connected to real work

Internship quality

Share of AI related internships that are paid, supervised, linked to defined deliverables and formally evaluated

Opportunity equity

Participation and completion by institution tier, geography, gender and socioeconomic background

Skills-based evaluation

Share of entry-level hiring processes using observable work samples, live verification, portfolios or structured capability rubrics

Project to opportunity conversion

Share of project participants progressing to internships, apprenticeships, interviews or other substantive pathways

Internship to offer conversion

Conversion to employment or extended project work, interpreted alongside quality and equity

Applied outcome

Number of projects progressing to further testing, implementation, research collaboration or product development



The future we need to build

Chapter 8

A shared agenda for responsible AI transformation

In the report, Sections 2 and 3 examine the gap between widespread AI adoption and genuine institutional readiness, Section 4 benchmarks India's higher education ecosystem against global approaches to AI implementation, Sections 5 and 6 outline the institutional and national enablers required to bridge these gaps and Section 7 translates these foundations into a clear agenda for coordinated stakeholder action. Ultimately, the success of these efforts would be measured not by the proliferation of licenses, policies, pilots or partnerships, but by their ability to advance deeper learning, strengthen credible capabilities and expand equitable access to opportunity.

8.1 Vision for 2030: Human + AI campuses, not AI-driven campuses

The campus of 2030 is not defined by the extent to which AI makes decisions or generates content, but by its ability to remove low-value friction and create capacity for higher-value human engagement. In such an environment, students are empowered to think more critically, faculty are able to teach more responsively and administrators can devote greater attention to complex human tasks. Intellectual authorship, pedagogical judgment, empathy and accountability remain fundamentally human responsibilities, while AI serves as a tutor, translator, critic, research assistant and administrative support tool rather than an invisible substitute for learning, reflection or care.

As AI makes information, explanation and content production cheaper and more abundant, the distinctly human dimensions of higher education become more valuable, not less. Mentorship, dialogue, collaboration, intellectual disagreement, belonging, and the experience of having one's reasoning tested by others cannot be reduced to content delivery. The Human + AI campus should therefore use technology not to substitute for these interactions, but to create greater space for them.

The true measure of an AI-ready campus is not universal tool adoption, but whether AI enhances learning without diminishing students' ability to think critically, communicate clearly and exercise independent judgment. Insights from student focus group discussions and stakeholder interviews underscore the importance of defining responsible AI use in practice, identifying the capabilities that must continue to be developed independently, and assessing whether AI functions as a genuine equalizer for learners facing educational, linguistic or resource constraints, rather than reinforcing existing patterns of dependency and exclusion.

8.2 India's opportunity to lead global AI adoption in education

India's opportunity lies in demonstrating that responsible AI adoption can be implemented effectively across languages, disciplines, geographies, institutional tiers and resource environments. Together, the AI Highway, national readiness standards, shared Centres of Excellence, and industry integrated learning pathways present a distinctly Indian model of AI transformation, one that combines excellence at the frontier with the broad distribution of capability across the higher education system.

India's large, adaptive and digitally engaged student population provides a strong foundation for AI-enabled transformation. However, widespread informal use of AI should not be mistaken for institutional readiness. Global leadership will depend on the ability to convert students willingness to experiment into structured capability through faculty development, credible assessment practices, shared digital infrastructure and effective governance. Innovation remains concentrated within institutions already equipped with resources and capacity cannot be considered national transformation. India's leadership is likely to be realized only when the benefits of AI extend to state universities, affiliated colleges, regional language learners and mobile-first students, ensuring that opportunity is distributed as broadly as capability itself.

The question is no longer whether AI will shape higher education, but whether institutions can harness it in ways that strengthen learning, human capability and educational equity

8.3 Closing reflections – The courage to tell the truth

Students are already engaging with AI, irrespective of whether institutional policies have evolved at the same pace. Evidence from this study suggests that AI adoption is often driven not only by a desire for efficiency but also by attempts to navigate unclear teaching practices, repetitive assessment tasks, heavy academic workloads and limited feedback mechanisms. At the same time, faculty cannot be expected to redesign teaching and assessment through isolated professional development interventions. Equal access to AI tools does not automatically translate into equal capability and employers cannot reasonably expect AI-ready graduates while remaining largely detached from the learning process until the point of recruitment.

AI is simultaneously expanding student agency, masking learning gaps in some contexts, and bringing greater visibility to institutional challenges across pedagogy, assessment, faculty preparedness and governance. Insights from students, faculty, administrators and employers may not always lead to comfortable consensus, but they provide a critical means of distinguishing genuine educational progress from performative adoption. The question is no longer whether AI will become part of Indian higher education—it already has. The more important question is whether institutions are willing to undertake the structural and pedagogical changes necessary to respond to this reality in a deliberate, equitable and accountable manner.

Looking ahead

Artificial intelligence in higher education is evolving too rapidly for any static report to fully capture the pace, breadth and direction of change. Institutional policies, technology capabilities, pedagogical practices, employer expectations and student behaviors are all shifting in parallel, often faster than conventional research and publication cycles can reflect. This report should therefore be viewed as a point-in-time assessment and a foundation for continued inquiry, rather than a definitive account of a space that will continue to evolve.

Going forward, the ambition for FICCI and EY-Parthenon is to build on this foundation through a broader and more continuous program of engagement with higher education institutions, students, faculty, employers, technology providers and policymakers. A larger and more diverse sample would allow subsequent work to test emerging findings more systematically, deepen quantitative evidence, and capture differences across institution types, disciplines and geographies. To support this, we propose quarterly addenda to the report that track material developments, revisit key questions, incorporate new evidence and stakeholder perspectives, and progressively build a richer longitudinal view of how Indian higher education is responding to AI.

Team

EY-Parthenon



Dr. Avantika Tomar
Partner



Aannshul Bandiya
Senior Consultant



Sanjna Vivek
Associate Consultant



Prarthana Madan
Assistant Manager,
India Insights CoE

FICCI Higher Education Team



Dr. Rajesh Pankaj
Senior Director & Head
Education & Skills



Partheshwar Kumar Dubey
Senior Assistant Director



Vidit Grover
Assistant Director



Khushi Kumari
Research Associate

We also acknowledge the contributions of **Tarun Tangri (Risk Management), Shweta Sharma and Sushmita Yadav (Editorial), and Jasmeet Joshi and Saloni Jain (Design)** in the making of this report.

About FICCI Higher Education

FICCI has been playing a proactive role in the Higher Education sector supported by the Higher Education Committee, comprising key representatives from leading Higher Education Institutions/Universities, Industry and the Government. Over the past two decades, FICCI has emerged as a prominent catalyst in shaping India's higher education landscape and has been playing a pivotal role in bringing about the requisite policy changes in the sector. With effective interface between higher educational institutions, industry, and government agencies, FICCI Higher Education Committee has not only created a dynamic forum for critical discourse but has also propelled the sector forward through progressive dialogue, vibrant knowledge exchange, and robust policy advocacy for strategic reforms.

Some of the key initiatives of the Committee are:

1. **Thought Leadership:** The Committee publishes several industry reports, policy briefs, and papers that intricately capture the evolution and progress of the higher education sector. Some of these include Reports on 'Vision 2030 for Higher Education,' 'Future of Jobs and its Implications in Indian Higher Education,' 'Higher Education in India: Vision 2047,' and 'Leapfrogging to Education 4.0.' Notably, the 'Higher Education Vision 2030' report significantly influenced the design of the National Education Policy (NEP) 2020.
2. **Policy Advocacy:** The HE committee collaborates with key stakeholders including government and industry, in actively identifying opportunities and addressing gaps. Through intensive consultations with a diverse array of stakeholders, this collaborative effort unfolds across various conferences, discussions, events, and forums. Brimming with shared wisdom and expertise, these engagements have proven pivotal in propelling the education sector reforms in India.
3. **International Delegation of Higher Education Leaders:** The Delegation of senior Higher Education Leaders to focussed countries has led to exposure to global best practices, facilitated effective teaching and research partnerships, and helped develop a comprehensive global education policy landscape.
4. **Higher Education Forum:** FICCI's Higher Education Forum serves as a premier platform for advancing dialogue, collaboration, and excellence in India's higher education ecosystem through its flagship initiatives—the FICCI Higher Education Summit (HES), Higher Education Excellence (HEE) Awards, and the FICCI Industry–Academia Conference (FIAC).
5. **Capacity Building –** FICCI Leadership Development Program (LDP): A residential leadership programme designed to strengthen institutional leadership and future readiness through expert-led sessions, case discussions, and peer learning. The programme covers governance, innovation, internationalisation, technology, research, and strategic leadership, with sessions delivered by leading experts from academia, policy, and industry.

Contact

Dr. Rajesh Pankaj

Senior Director & Head-Education & Skills

rajesh.pankaj@ficci.com

About our practice

EY-Parthenon, our strategy arm, consists of a team of 50+ partners and 1000+ consultants in India and ranks 4th in the Vault Consulting 50 list of top consulting firms. EY-Parthenon is the leading strategic advisor to the education sector globally, with our education team completing 300+ projects annually.

The EY-Parthenon education consulting strategists help clients negotiate the changing currents in the sector so that they not only adapt but also adopt strategies in terms of globalization-driven skill sets and new collaborations.

With broad experience and deep sector knowledge, the education strategy consulting professionals at EY-Parthenon are helping leaders overcome challenges with bespoke, all-encompassing growth strategy plans, due diligence services and implementation support.

We have dedicated consultants in the following five segments of the sector:

Governments and foundations	Pre-K and K-12 school chains	Higher education institutions and TVETs	Indian and global ed-tech companies	Global investors
Our clients include central and state ministries of education, supporting organizations and foundations. We have supported in developing short-term and long-term growth strategy plans to reform systems.	Our teams provide services such as market needs assessment, strategic planning, performance analytics, operational improvement, financial advisory and organizational redesign.	Our teams help HEIs identify opportunities for differentiation through various modes, using our insights from global best practices. We also help TVETs formulate end-to-end strategies and help executing the same.	We provide competitive landscaping, market analyses, go-to-market strategies, support on organic and inorganic growth like fundraising, acquisitions, partnerships, joint ventures or divestments.	We provide due diligence services to investors. From the pre-contract stage through the eventual integration or separation, we help guide decision-making and provide execution assistance.

Contact

Dr. Avantika Tomar

Partner and L&D Head

Avantika.Tomar1@parthenon.ey.com

Acknowledgements

We sincerely thank the Institutional leaders, faculty members, senior administrative staff, industry representatives, and students from the following institutions for their valuable time, insights, and contributions. Their perspectives has been instrumental in shaping the findings and recommendations presented in this Knowledge Report.

- Indian Institute of Technology (IIT) Roorkee
- Jio Institute
- Vellore Institute of Technology
- BITS Pilani
- Manipal Academy of Higher Education (MAHE)
- Dayalbagh Educational Institute
- Symbiosis International (Deemed University)
- Thapar Institute of Engineering and Technology, Patiala
- Banaras Hindu University (BHU)
- Jigyasa University
- Biju Patnaik University of Technology
- Atmiya University
- Charotar University of Science and Technology (CHARUSAT)
- Nagesh Karajagi Orchid College of Engineering and Technology
- Vivekanand Education Society Business School (VES Business School)
- Vidyavardhaka College of Engineering (VVCE)
- KCG College of Technology
- Sri Ramakrishna Institute of Technology
- Adamas University
- Ashoka University
- Indian Institute of Technology (IIT) Mandi
- Deakin University, GIFT City – Gujarat
- Tamil Nadu Music and Fine Arts University

Notes



Our offices

Ahmedabad
22nd Floor, B Wing, Privilon
Ambli BRT Road, Behind Iskcon
Temple, Off SG Highway,
Ahmedabad - 380 059
Tel:+ 91 79 6608 3800

Gandhinagar
8th Floor, Building No. 14A
Block 14, Zone 1
Brigade International Financial
Centre
GIFT City SEZ
Gandhinagar – 382 355
Tel:+ 91 79 6608 3800

Bengaluru
12th & 13th Floor
"UB City", Canberra Block
No.24 Vittal Mallya Road
Bengaluru - 560 001
Tel:+ 91 80 6727 5000

Ground & 1st Floor
11, 'A' wing
Divyasree Chambers
Langford Town
Bengaluru - 560 025
Tel:+ 91 80 6727 5000

3rd & 4th Floor
MARKSQUARE
#61, St. Mark's Road
Shantala Nagar
Bengaluru - 560 001
Tel:+ 91 80 6727 5000

1st & 8th Floor, Tower A
Prestige Shantiniketan
Mahadevapura Post
Whitefield, Bengaluru - 560 048
Tel:+ 91 80 6727 5000

Ecospace
1st Floor, Campus 1C
Ecospace Business Park
Outer Ring Road,
Bellandur - Sarjapura Area,
Varthur Hobli,
Bengaluru Urban – 560103

Bhubaneswar
8th Floor, O-Hub, Tower A
Chandaka SEZ, Bhubaneswar –
751024
Tel: + 91 674 274 4490

Chandigarh
Elante offices,
Unit No. B-613 & 614
6th Floor, Plot No- 178-178A
Industrial & Business Park,
Phase-I
Chandigarh - 160 002
Tel:+ 91 172 6717800

Chennai
6th & 7th Floor, A Block,
Tidel Park, No.4,
Rajiv Gandhi Salai
Taramani, Chennai - 600 113
Tel:+ 91 44 6654 8100

Delhi NCR
Aikyam, Ground Floor
67, Institutional Area
Sector 44, Gurugram - 122 003
Tel: + 91 124 443 4000

3rd & 6th Floor, Worldmark-1
IGI Airport Hospitality District
Aerocity, New Delhi - 110 037
Tel: + 91 11 4731 8000

4th & 5th Floor, Plot No 2B
Tower 2, Sector 126
Gautam Budh Nagar
Noida - 201 304
Tel:+ 91 120 671 7000

Hyderabad
THE SKYVIEW 10
18th Floor, "SOUTH LOBBY"
Survey No 83/1, Raidurgam
Hyderabad - 500 032
Tel:+ 91 40 6736 2000

THE SKYVIEW 20
2nd Floor, 201 & 202
Right Wing, Survey No 83/1
Raidurgam, Hyderabad –
500 032
Tel:+ 91 40 6736 2000

Jaipur
9th floor, Jewel of India
Horizon Tower, JLN Marg
Opp Jaipur Stock Exchange
Jaipur - 302018
Tel:+ 91 141 4618415

Kochi
7th Floor, ABAD Nucleus
NH-49, Maradu PO
Kochi - 682 304
Tel:+ 91 484 433 4000

Kolkata
22 Camac Street
3rd Floor, Block 'C'
Kolkata - 700 016
Tel:+ 91 33 6615 3400

6th floor, Sector V,
Building Omega, Bengal
Intelligent Park, Salt Lake
Electronics Complex, Bidhan
Nagar, Kolkata - 700 091
Tel: + 91 33 6615 3400

Mumbai
14th Floor, The Ruby
29 Senapati Bapat Marg
Dadar (W), Mumbai - 400 028
Tel:+ 91 22 6192 0000

5th Floor, Block B-2
Nirlon Knowledge Park
Off. Western Express
Highway, Goregaon (E)
Mumbai - 400 063
Tel:+ 91 22 6192 0000

3rd Floor, Unit No.301
Building No.1, Mindspace-
Gigaplex
IT Park, MIDC, Plot No. IT-5
Airoli Knowledge Park
Airoli West, Navi Mumbai -
400 708
Tel: + 91 22 6192 0003

18th Floor, Altimus
Pandurang Budhkar Marg
Worli, Mumbai - 400 018
Tel: + 91 22 6192 0503

Pune
C-401, 4th Floor
Panchshil Tech Park,
Yerwada
(Near Don Bosco School)
Pune - 411 006
Tel: + 91 20 4912 6000

10th Floor, Smartworks
M-Agile, Pan Card Club Road
Baner, Pune - 411 045
Tel: + 91 20 4912 6800

Ernst & Young LLP

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multidisciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EYG member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

About EY-Parthenon

Our unique combination of transformative strategy, transactions and corporate finance delivers real-world value – solutions that work in practice, not just on paper.

Benefiting from EY's full spectrum of services, we've reimagined strategic consulting to work in a world of increasing complexity. With deep functional and sector expertise, paired with innovative AI-powered technology and an investor mindset, we partner with CEOs, boards, private equity and governments every step of the way – enabling you to shape your future with confidence.

EY-Parthenon is a brand under which a number of EY member firms across the globe provide strategy consulting services. For more information, please visit www.ey.com/parthenon.

Ernst & Young LLP is one of the Indian client serving member firms of EYGM Limited. For more information about our organization, please visit www.ey.com/en_in.

Ernst & Young LLP is a Limited Liability Partnership, registered under the Limited Liability Partnership Act, 2008 in India, having its registered office at Ground Floor, Plot No. 67, Institutional Area, Sector - 44, Gurugram - 122003, Haryana, India.

© 2026 Ernst & Young LLP. Published in India.
All Rights Reserved.

EYIN2610-003

This publication contains information in summary form and is therefore intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. Neither EYGM Limited nor any other member of the global Ernst & Young organization can accept any responsibility for loss occasioned to any person acting or refraining from action as a result of any material in this publication. On any specific matter, reference should be made to the appropriate advisor.

SJ

About FICCI

Established in 1927, FICCI is the largest and oldest apex business organisation in India. Its history is closely interwoven with India's struggle for independence, its industrialization, and its emergence as one of the most rapidly growing global economies.

A non-government, not-for-profit organisation, FICCI is the voice of India's business and industry. From influencing policy to encouraging debate, engaging with policy makers and civil society, FICCI articulates the views and concerns of industry. It serves its members from the Indian private and public corporate sectors and multinational companies, drawing its strength from diverse regional chambers of commerce and industry across states, reaching out to over 2,50,000 companies. FICCI provides a platform for networking and consensus building within and across sectors and is the first port of call for Indian industry, policy makers and the international business community.

Address:
FICCI, Federation House, Tansen Marg, New Delhi, 110001

www.ficci.in

ey.com/en_in

EY_India EY EY India EY Careers India ey_india