



FUTURE OF EMPLOYABILITY IN THE **AGE OF AI**

INDIA'S PLAYBOOK FOR STUDENTS, INSTITUTIONS, AND INDUSTRY

A Deep-Dive Report by AI4India.org | January 2026



Alok Agrawal

Co-Founder
AI4India



Shashi Shekhar Vempari

Co-Founder
AI4India



Adarsh Lathika

Founder and CEO
Anatomy of Work

“

AI is poised to fundamentally reset the landscape of employability. While the change may not be immediate, the long-term impact is undeniable, requiring a radical transformation of our education system at its core. As traditional curricula and degrees lose their relevance, a new era of opportunity will emerge – one that demands a workforce defined by human ingenuity, creativity, and complex problem-solving. In this future, we must all become lifelong students, returning to learn and reskill as industry evolves. Drawing on its ancient heritage of academic excellence, India has the potential to lead the world in bridging the gap between education and employability in the age of AI.

”

“

With the rise of the algorithm-driven economy and automation, AI can enhance efficiency, productivity, and create new classes of jobs, even as some traditional roles become redundant. AI and job creation can go hand in hand, which is crucial for India given its young and growing workforce. Academia and Industry need to work together to create an enabling environment for students to be ready to join the AI-era workforce. This collaboration is crucial if India is to get ahead of the AI adoption curve and prepare our youth for the new class of jobs/opportunities that will arise in this era of co-intelligence with interplay between AI and innately human intuition, creativity, and instincts.

”

“

Across my conversations with students, educators, and industry leaders, one thing became clear: expectations from early-career professionals are rising faster than the structures meant to prepare them. Entry-level roles are becoming thinner, not because work is disappearing, but because routine work is being absorbed elsewhere. That shift quietly changes what employability means. Preparing people only to 'fit into jobs' may no longer be enough. We may need many more individuals who can create work – who can identify problems, assemble tools, and generate value rather than wait for formal roles to appear. This places new pressure on education, skilling, and leadership ecosystems to cultivate judgment, initiative, and responsibility, not just credentials. The question ahead is not whether AI will change work – it already has – but whether we are equipping people to adapt with agency rather than anxiety.

”

**Gopal Devanahalli**

President, Skilling
Wadhwani Foundation

“Just as the internet, in the 1990s, transformed how we live and work, artificial intelligence is poised to reshape our world three decades later. Its implications for youth, skills, and jobs are profound. Realising this opportunity will require a coordinated response from Indian policymakers, industry, educational institutions, and citizens alike.

”

**Krishnan Narayanan**

Co-Founder and President
itihaasa Research and Digital

“From my conversations with business leaders, two truths stand out. First, entry-level work is shifting from doing routine tasks to working with AI, including framing problems, checking outputs, and owning judgement. Second, enterprises must redesign roles: rewrite job descriptions, change hiring tests, and build fast, hands-on upskilling pathways so young people stay employable.

”

**Jai Asundi**

Executive Director
Center for Study of Science, Technology
and Policy (CSTEP)

“Artificial Intelligence (AI) has the potential to fundamentally transform society. How it impacts India is of great interest to many of us who have seen the potential and pitfalls of technology in a developing country such as ours. It is imperative that we pay particular attention to the details associated with the development, use, and governance of AI. An important perspective we must consider is the training and employability of students in this new age. We need to be able to match the interests of industry/employers and academia so that we are able to graduate model citizens of tomorrow. Leveraging AI for the same would be ideal, however, this will take strategic design and intent which this report covers. Great to see such efforts from AI4India.

”

Why This Study Now: The Convergence of Three Crises

Three intersecting crises make this research urgent and timely:

Crisis 1: The Institutional Response Gap

In 2025, higher education institutions remain largely unprepared for the AI-era learning despite having some policies in place. Secondary research shows:

- **60% of institutions¹** permit student AI usage, but **very few have structured AI-first pedagogies**.
- **Faculty readiness gaps** are significant. As of 2025, only **17% of faculty** rate their AI proficiency as "Advanced" or "Expert," and only **6%** feel satisfied with the AI literacy resources provided by their institutions¹.

This institutional hesitation, where institutions restrict or avoid AI integration without a clear pedagogical framework, creates a structural vacuum. Students fill this vacuum by learning informally, outside institutional structures, without mentorship or quality assurance.

Crisis 2: Job Market Anxiety Uncoupled from Reality

Popular discourse on AI and employment oscillates between apocalyptic warnings and techno-optimistic reassurance. Yet these macro-level projections obscure a critical reality: displacement risk is sector-specific, not universal. Some roles face genuine disruption; others face augmentation. Many face neither.

Industry surveys show that employers are not asking for "AI skills" per se, but rather adaptability, critical thinking, and the ability to work with AI tools effectively. Despite this evidence, widespread job market anxiety persists. Students struggle to translate macro employment forecasts into decisions about their own career paths. The misalignment between what employers are signalling and what students perceive is driving defensive career choices and missed opportunities.

Crisis 3: The Capability-Access Inversion

The third crisis is perhaps the most counterintuitive: **students have more access to AI than ever before, but less institutional guidance on how to develop genuine capability**.

53.5% of Indian students² use AI daily (free tools like ChatGPT). The result is a generation developing AI capability haphazardly, without feedback loops, without connection to domain expertise, without the institutional scaffolding that transforms tool access into genuine learning.

1 EY-FICCI AI Adoption Survey 2025

2 Lathika, A. (2025, December). The shadow curriculum: How students are rebuilding higher education with AI – Faster than institutions can respond.

The Inflection Point: Why January 2026 Matters for Employability in the AI Era

- 1. Policy Windows Are Opening:** Government and regulatory bodies are beginning to move from reactive (banning AI) to proactive (designing AI integration). The next 6-18 months will see policy framework creation that could either enable or constrain institutional transformation.
- 2. Employer Signalling Is Shifting:** Forward-thinking organisations are beginning to hire based on portfolio evidence and demonstrated capability rather than credentials. This trend will accelerate in 2026-2027. Institutions that prepare students for portfolio-based assessment now will have a competitive advantage; those that wait will miss the window.
- 3. Faculty Readiness Has a Time Horizon:** Faculty development is not instantaneous. A faculty member trained in AI-aware pedagogy in 2025 will refine their practice and teach 100+ students with improved methods over the next 3-5 years. A faculty member who delays until 2027 cannot catch up in time for the 2025-2028 graduating cohorts.
- 4. Tier 2/3 Institutional Support Requires Lead Time:** External support (government funding, trainer deployment, and infrastructure development) takes 6-12 months to mobilise. Without action now, these institutions will fall further behind structurally.
- 5. Student Expectations Are Rising:** Current students are increasingly aware that their institutions are not preparing them adequately for AI. This awareness will drive institutional pressure if not matched by an institutional response.

The evidence is clear: The choices made in the next 18 months will determine whether India's higher education system shapes an AI-era-ready generation or perpetuates and deepens existing inequalities.

NOTES

Preliminary findings from qualitative research commissioned by AI4India.org.

Primary Researcher: Adarsh Lathika

Publishers: Alok Agrawal and Shashi Shekhar Vempati

Reviewers: Jai Asundi, Rajesh Shenoy, Krishnan Narayanan, and Gopal Devanahalli

Research Period: November – December 2025

Methodology: This report is synthesised using a multi-method research design, centred on a systematic review of over 85 structured, primary conversations. Our contributors represent a cross-section of the Indian economic landscape, including leaders from industry, academia, government, EdTech, staffing agencies, and student communities. This “multi-stakeholder” lens ensures the findings reflect the friction between classroom learning and workplace requirements.

Disclaimer: The views expressed in this report are solely those of the authors and reviewers and do not reflect the positions of any affiliated employers.

Confidentiality Note: All company-specific data and quotes have been anonymised to maintain compliance with corporate disclosure policies and confidentiality agreements, ensure neutrality, and prevent any perception of commercial advancement or opinion.

AI Usage & Ethics Declaration: In the spirit of the “AI-augmented capability” discussed in this report, the authors employed Generative AI tools (including OpenAI’s GPT-5.2, Perplexity and Google Gemini) to assist in the thematic synthesis of interview transcripts, data structuring, and preliminary drafting. All AI-generated outputs were rigorously reviewed, fact-checked, and edited by human subject-matter experts to ensure accuracy, nuance, and contextual integrity. AI was used as a “copilot,” but the strategic insights and conclusions remain entirely human-led.



© 2026 DeepTech4Bharat Foundation. All rights reserved. No part of this publication—including text, layout design, and logos—may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the publisher.

Contents

Executive Summary	08
1. The AI Exposure Landscape - Mapping Student Agency	14
2. Industry Lens - What Employers Are Saying	28
3. Academia Lens - How Universities Are Responding	36
4. Policy Pathways & Suggested Recommendations	49
5. Conclusion: From Awareness to Action	64
Appendices	74



EXECUTIVE SUMMARY

A Moment of Inflection

India stands at an inflection point in its employability story. Artificial intelligence is no longer a future scenario; it is reshaping how work gets done, who gets hired, what skills command a premium, and where opportunity concentrates. Over the past few weeks³, we conducted 85+ in-depth conversations with industry leaders, faculty members, students, EdTech founders, and policymakers across Bengaluru, Delhi, Kolkata, and smaller cities. A single question unified the inquiry: **What does it take to be employable in India in the age of AI?**

The answer is neither reassuring nor pessimistic. It reveals a country at a fork in the road. One path leads to an AI-augmented talent pipeline that is inclusive, geographically distributed, and rooted in Indian contexts. The other path leads to deepening inequality, where elite institutions and Tier 1 metros capture most opportunity while Tier 2 and 3 regions remain structurally disadvantaged. **That choice is being made right now, in 2026, largely through institutional default rather than deliberate policy.**

This report synthesises four critical narratives: (1) the paradox of access without capability, (2) how work is changing, (3) the signals that will shape hiring and opportunity over the next five years, and (4) what must change immediately to prevent regional inequality from hardening.

Industry: Compressed Roles and New Hiring Signals

Workflows are Compressing, Not Disappearing

Across sectors, AI is compressing workflows, reducing layers and changing what remains, rather than eliminating functions entirely. In planning, proposal writing, software engineering, and data analysis, work once distributed across large junior teams is now handled by smaller teams working with AI copilots and automation. There are fewer entry-level slots doing routine execution, but the jobs that remain require deeper judgment, contextual understanding, and system-level thinking.

From Credentials to Capability and Mindset

Employers consistently report that **technical depth and degrees matter less than adaptability, curiosity, and cross-domain thinking**. A moderate level of technical skill coupled with strong problem-solving, communication, and learning agility is often more valuable than advanced training without these traits. Portfolios such as GitHub repositories, case studies, design work, documented experiments are increasingly being used alongside or instead of resumes to evaluate readiness.

At the same time, business uncertainty is slowing AI adoption. Many companies are still refining what "AI-ready talent" actually means for different roles, leading to vague job descriptions, legacy interview formats, and few AI projects scaling up beyond the pilot stage. The employers moving fastest are those that have answered three questions clearly: what AI will handle, where human judgment remains, and which capabilities matter most in an AI-augmented team.

Students: Access Without Capability

Tools Are Ubiquitous; Understanding Is Rare

India is among the world's highest users of AI in higher education: 53.5% of students report daily AI tool usage, and another 23.5% use AI weekly. Free-tier Gemini, Perplexity, open-source models, and older versions of Claude are available on most smartphones, thus levelling the access gap between a student in a Tier 3 town and a student in a top metro college.²

However, usage depth is shallow. A majority of students use AI for summarising, quick explanations, and last-minute assignment support, while only 10–15% operate at a level where they iteratively refine outputs, critique model reasoning, and deploy AI in structured projects. This gap is primarily due to the absence of mentorship, curricular scaffolding, and exposure to high-quality problems with adequate real-world context.

Hardware Inequality Defines the Scope of Learning

Device access has become a limiting factor on the quality of AI learning. Among engineering students in Tier 1 institutions, around **95% have a laptop or desktop**, whereas among non-engineering students in **Tier 3 institutions, only about 25%** do. A student with only a smartphone cannot meaningfully code, train models, or build deployable projects, effectively excluding themselves from much of what employers recognise as employment-ready AI skill or capability.⁴

Institutional access to paid AI tools is also skewed. Even in the most privileged cohorts, only a minority have institution-provided access to premium AI platforms; in many Tiers 2/3 colleges, especially in non-engineering programs, students rely entirely on free tiers with limited compute and features. Device and compute inequality thus translates directly into "scope inequality": two equally motivated students diverge simply because one has access to premium tools on an appropriate compute device while the other lacks access to both.

Anxiety–Reality Mismatch

Nearly half of Indian students believe heavy reliance on AI is reducing their preparedness for the workplace. At the same time, their anxiety often does not map to actual job risk: students in relatively stable domains such as healthcare, education, and law frequently express as much fear as those in high-disruption areas like basic coding or content generation. Without grounded, domain-specific guidance, counselling and mentorship, students are abandoning promising paths to chase hype-driven courses, instead of building the durable capabilities that will matter most.

Academic Institutions: Fear, Fragmentation, and the Assessment Crisis

Denial, Policing, and Paralysis

Faced with rapid student adoption of AI, many institutions have responded with denial ("AI is not our problem"), policing (bans, detection thresholds, punitive policies), or paralysis (committees and circulars without meaningful change). This combination drives AI usage underground: students

⁴ Primary Research Synthesis (Nov–Dec 2025). These figures represent a qualitative synthesis of reported access rates from stakeholder interviews with EdTech providers and institutional leadership. The 70-percentage-point delta is a thematic deduction based on platform telemetry and faculty observations, highlighting the "Compute Ceiling" that restricts regional non-engineering cohorts to mobile-only, low-depth AI interactions.

continue to use AI tools unsupervised without guidance, thus losing the chance to receive meaningful feedback on their use of AI.

The result is a trust gap with students treating institutional rules as obstacles as they look to work around them. There is little by way of guidance in these institutional rules on how to think and act responsibly with AI. Faculty, meanwhile, report genuine uncertainty: they are not trained in AI-aware pedagogy, have little clarity on how to redesign assignments, and often feel personally threatened by tools that automate parts of their expertise.

Geographic Inequality as the Primary Fault Line

The dominant divide in India's AI readiness is not simply rich versus poor, but **Tier 1 metros versus Tier 2/3 towns**. Tier 1 institutions have better access to industry mentors, updated curricula, GPU-enabled labs, and local employers who can feed real projects into classrooms; Tier 2/3 institutions often lack all four. Students in Tier 2/3 areas, therefore, rely more heavily on social media, bootcamps, and informal networks for AI learning. This makes them more vulnerable to hype and makes it harder for them to derive substantial value from AI tools.

Over time, this gap translates into a compounding advantage for Tier 1 students. A Tier 1 student with high-quality mentorship and project experience enters the job market several steps ahead, and this gap widens with each career transition. Left unaddressed, this geography-driven inequality will harden into a structural divide that cannot be closed by individual effort alone.

Assessment Has Decoupled From Reality

Underneath these patterns lies a deeper issue: **assessment systems are measuring the wrong parameters**. Traditional closed-book, memory-focused exams rewarded recall in a pre-AI world; they now say little about whether a student can frame problems, work with AI tools, evaluate outputs, or communicate reasoning. Banning AI during exams or adding AI-detection tools on top of unchanged assessments does not address this misalignment between assessment systems and the imperatives of AI.

Only a few institutions have begun shifting to process-based evaluation, requiring students to submit prompts, intermediate outputs, and reflections alongside final answers. But these are still exceptions. Without broad assessment reforms, academic degrees will increasingly lose their significance with employers, regardless of institutional prestige.

Infrastructure: Four Factors That Shape Opportunity

The report identifies four infrastructure factors that now strongly influence who can develop AI-era capabilities:

1. **GPU access determines what students can learn:** Without access to GPUs or shared cloud compute, students cannot meaningfully experiment with training, fine-tuning, or deploying current generation AI systems; they remain confined to chat interfaces and small-scale tasks.
2. **Device access determines the scope of learning:** Students without laptops are effectively excluded from coding, systems design, and project-based learning; this "device divide" is especially pronounced for non-engineering students in Tier 2/3 colleges.
3. **Localised datasets determine relevance:** Learning AI on Indian health, agriculture, logistics, and

financial datasets builds domain awareness and employable skills; relying only on generic global datasets produces capabilities that are disconnected from local needs and realities.

4. **Mentorship infrastructure determines the pace of learning:** Students with access to practitioners who show how AI fits into real workflows progress in weeks to a few months; those without such mentorship could waste valuable time guessing which tools and skills matter, thus inordinately delaying their pace of learning.

These factors are all addressable through coordinated investment and policy decisions and do not require deep technological breakthroughs.

What Students Can Do Now

Students need not remain passive during this transition; they can act even before systems fully adapt:

- **Use AI as a thinking partner, not a shortcut:** Treat AI tools as collaborators that help explore ideas, generate alternatives, and test understanding, while keeping human judgment at the centre.
- **Build portfolios that show real work:** Convert class assignments, internships, and self-initiated experiments into visible artefacts such as code, models, case write-ups, and design mockups that demonstrate how AI was applied to meaningful problems.
- **Go deep into any one domain–AI combination:** Rather than chasing every new tool, pick a domain they care deeply about (finance, health, law, agriculture, education, design) and learn how AI is actually being used within that domain.
- **Seek mentors and peer communities:** Join or create groups on campus or online where one can share prompts, critique outputs, and get feedback from seniors, alumni, and professionals.
- **Replace headline-driven fear with informed action:** Study credible sector-specific trends and align preparation with the capabilities employers repeatedly highlight: critical thinking, adaptability, communication, and effective collaboration with AI tools.

Imperatives for Institutions, Industry, EdTech, and Policymakers

The report proposes a set of imperatives that align actions so that all stakeholders move in step rather than in isolation or at cross-purposes:

Policymakers should:

- Declare AI literacy a national baseline and integrate it across disciplines.
- Fund shared compute and device-support schemes to close hardware gaps.
- Incentivise assessment reform, not AI bans.
- Support Indian language AI ecosystems and Indian datasets. Set a goal for sovereign foundational models being developed by IndiaAI Mission to enable code generation using AI in any of the multiple Indian languages.
- Professionalise AI pedagogy through recognised faculty certification.

Universities and colleges should:

- Shift from policing AI to teaching with AI via redesigned assignments and transparent usage norms.
- Guarantee a minimum level of device and compute access for all students.
- Invest in faculty confidence and communities of practice around AI-augmented teaching.
- Bring real Indian problems into classrooms through long-term regional partnerships.
- Align curricula and assessments with the skills and mindsets employers actually seek.

Industry and CHROs should:

- Rewrite job descriptions to describe AI-augmented responsibilities and required judgment.
- Adopt portfolio- and task-based hiring processes.
- Launch AI apprenticeships that give early-career talent structured exposure to real workflows. Encourage employees to mentor students at local colleges.
- Co-design flexible micro-curricula with universities, using pre-approved "flex slots" inside core courses.
- Share anonymised use-case libraries and decision frameworks with educators and students.

EdTech and skilling platforms should:

- Reorient instruction away from standalone tool tutorials and toward capability-building that emphasises reasoning, evaluation, and multi-tool orchestration.
- Integrate real Indian projects and sectoral challenges into learning pathways.
- Create role-specific AI journeys for different professions rather than generic "AI for everyone" courses.
- Design mobile-first, multilingual experiences for low-bandwidth, Tier 2/3 contexts.

The Choices Ahead

India's future in AI employability remains open to deliberate choice. It is, fundamentally, about the choices that are currently being made about who gets access to meaningful learning, who has access to mentors and infrastructure, and which problem-sets are prioritised. Today, those choices are deepening a divide by concentrating opportunities in Tier 1 metros while creating structural disadvantages elsewhere.

However, better choices can be made to bridge this divide by leveraging tools, frameworks, and early success stories that already exist. By aligning students, institutions, industry, EdTech, and policymakers around a shared goal, the opportunity exists to turn **India's latent talent into an AI-augmented dividend that is broad-based, regionally inclusive, and globally competitive**. The challenge lies in coordinating action across stakeholders and in ensuring speed of execution to keep pace with the fast-changing AI technology landscape.



1

THE AI EXPOSURE LANDSCAPE - MAPPING STUDENT AGENCY

1.1 Reading the Ground Reality

If you stand in a classroom at IIT Delhi and ask students about AI, you hear stories of weekend hackathons, deployed projects, and GitHub portfolios with dozens of repositories. Walk into a regional engineering college in a Tier 3 town and ask the same question, and you hear uncertainty: "We've heard of ChatGPT, but we are not sure how to use it for our subjects." The gulf between these experiences is not a difference in intelligence, motivation, or even internet access. It is due to systemic gaps in **infrastructure, mentorship, institutional support, and exposure to contextualised problems**.

This section maps India's AI exposure landscape across four dimensions: what students are actually doing with AI, how their capability varies by geography and institution type, where the anxiety-reality gaps lie, and what infrastructure constraints are quietly shaping who gets to participate. The picture that emerges is both promising and troubling. It is promising because AI usage is widespread. However, it is troubling because capability depth is concentrated within an elite minority while the majority of students engage with AI superficially.

1.2 What the Data Shows: Usage Is High, Depth Is Low

Among Indian students⁴, **53.5% use AI tools daily**, another 23.5% use them weekly, and only 4.7% report rare or no usage. This places India among the countries with the highest AI adoption rates globally for higher education.

Yet beneath this high usage rate lies a critical nuance: **what students are doing with AI is overwhelmingly superficial**. When asked about their dependency on AI, 52.9% of Indian students report that they "use AI often but can manage without it," while 16.5% say they "struggle to work without AI"⁹. This dependency pattern suggests habitual usage rather than deep capability. While students rely on AI for routine tasks (summarisation, assignment shortcuts, and quick explanations), they have not developed the evaluative, iterative, or problem-framing skills that would make them genuinely AI-augmented learners.

Institutional support remains minimal. **78.8% of Indian students rely entirely on personal or free-tier AI tools**; only 8.8% receive full institutional access to premium tools, and another 6.5% have limited trial access⁹. This means most students are learning how to use AI through trial-and-error, peer sharing, and social media influencers, and not through structured pedagogy or mentored experimentation.

Perhaps most revealing is the workplace preparedness concern. **45.9% of Indian students believe that heavy reliance on AI is reducing their preparedness for the workplace**⁹. This anxiety is not irrational. It reflects an intuition that their current usage patterns (accepting AI outputs without evaluation, using AI as a shortcut rather than a thinking partner) are not building durable capability. Students sense they are becoming dependent without becoming competent.

1.3 The Maturity Divide: Five Levels of Capability

To move beyond the binary categories of "AI user" versus "non-user," this study proposes a five-level maturity model based on depth of engagement, evaluative capacity, and transfer to real-world contexts. The model is grounded in patterns observed across dozens of student conversations, faculty interviews, and industry hiring signals.

Student AI Usage Maturity - Distribution across institutional tiers

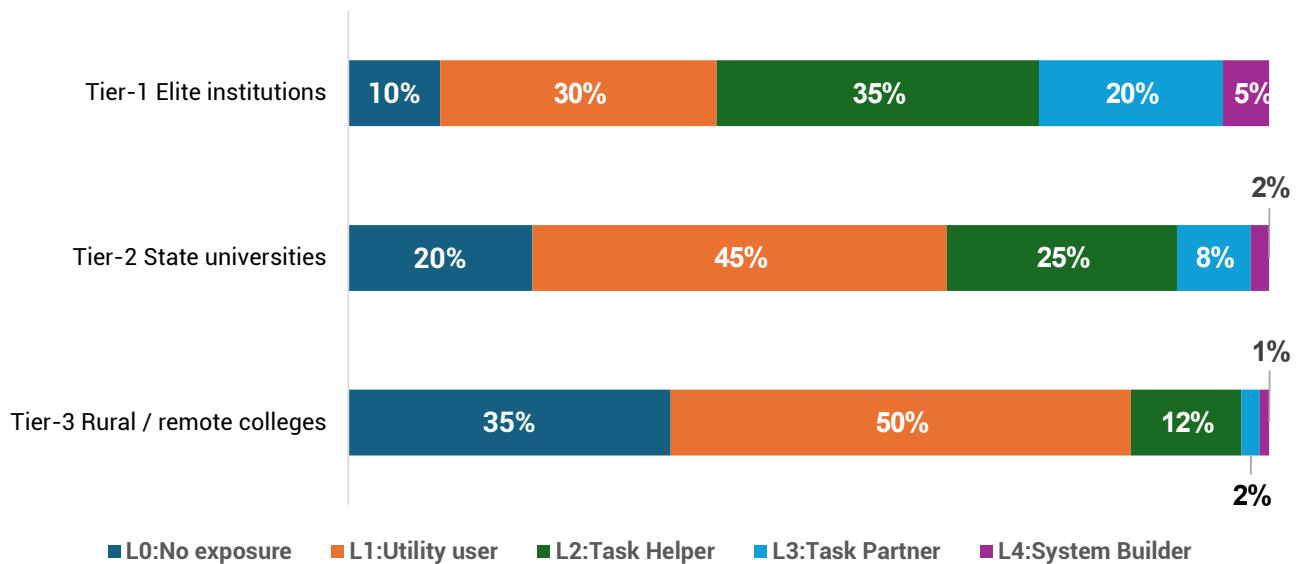


Figure 1: Student AI Usage Maturity distribution across Tier 1, 2, and 3 institutions in India, showing stark differences in capability levels⁵.

The distribution across these five maturity levels reveals India's employability challenge in stark terms. Approximately **20% of students remain at Level 0** (no meaningful AI exposure), mostly concentrated in Tier 3 institutions, non-engineering programs, and regions with device constraints. Another **40-50% are at Level 1** (utility users who treat AI as search engine upgrade). **25-30% reach Level 2** (task helpers who use AI for assignments but lack depth). **Only 8-10% achieve Level 3** (task partners who iterate, evaluate, and problem-solve with AI), and a mere **2-5% reach Level 4** (system builders who design workflows and deploy projects).

This distribution correlates tightly with **geography, institutional tier, discipline, and device access**. A computer science student at an IIT with a laptop, mentorship from industry-connected faculty, and peer communities focused on building projects is far more likely to reach Level 3 or 4. A commerce student at a Tier 3 college with only a smartphone, who is exposed to faculty that discourage AI usage, and has no access to real-world problem sets, is far more likely to remain at Level 1.

The employability implications are direct. Industry hiring leaders consistently describe requiring graduates at Level 3 or above - people who can iterate on AI outputs, evaluate relevance to business context, and apply AI to domain-specific problems. Yet **only 10-15% of Indian students are reaching that threshold**. The remaining 85-90% have "AI exposure" on their resumes but lack the depth that employers value.

1.4 The Infrastructure Constraint: Devices Determine Scope

One of the most consequential yet underappreciated barriers to AI-era employability is **device access**. A student with only a smartphone cannot meaningfully participate in many AI learning activities: coding, running local models, deploying projects, participating in systems design exercises, or building portfolio-ready artifacts.

⁵ Qualitative deductions synthesised from 85+ primary stakeholder conversations, including 40+ CHROs, academia, edtech, staffing firms, and 45 students.

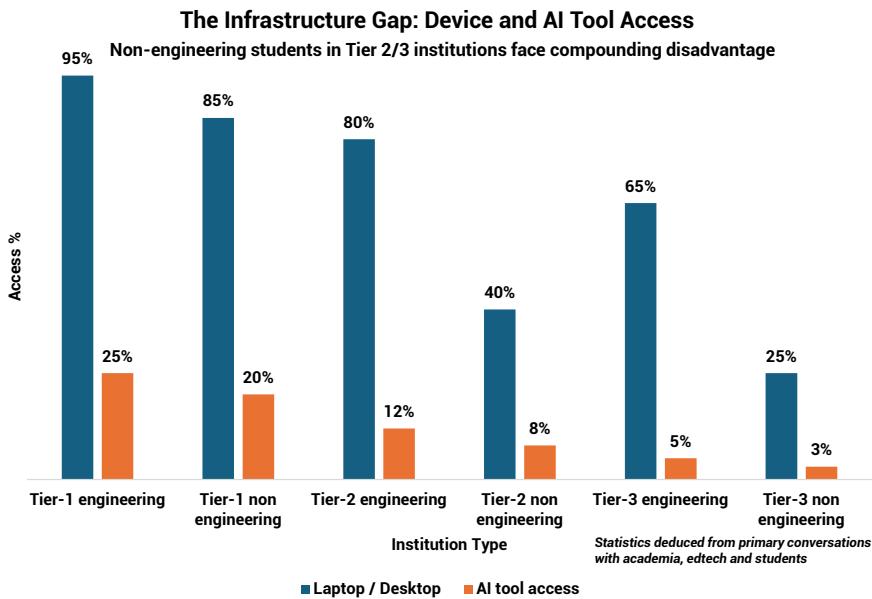


Figure 2: Infrastructure access disparities showing dramatic gaps in laptop ownership and institutional AI tool access, particularly affecting non-engineering students in Tier 2/3 colleges.

The data reveals a stark gradient. Among engineering students at Tier 1 institutions, **95% have laptop or desktop access**. Among non-engineering students at Tier 3 institutions, **only 25% have laptop access**. This 70-percentage-point gap directly impacts AI learning. A Tier 3 non-engineering student relying on a smartphone can use ChatGPT only for text queries. These students cannot code, debug, train models, experiment with parameters, or build deployable projects. Effectively, they are locked out of 80% of what constitutes "AI-augmented capability" in the job market.

Lack of institutional access to premium AI tools compounds the problem. Even among Tier 1 engineering students (the most privileged cohort), only **25% have institutional access to paid AI tools**. For Tier 3 non-engineering students, institutional access drops to 3%. This means students in under-resourced institutions must rely entirely on free-tier tools with limited compute, rate limits, and restricted features. A student trying to run an agentic workflow or train a custom model on free-tier tools will face limits to learning and experimentation when compared to a peer with institutional access to Claude Pro or enterprise Gemini.

Device inequality is the new digital divide. Just as internet access once determined who could participate in online learning, laptop access now determines who can participate in AI-augmented learning. Yet institutional responses have been minimal. Few Tiers 2 or 3 colleges have device subsidy programs, nor do they have adequate government funding programs to bridge this gap. Few institutions have GPU-enabled labs. Even those that permit "bring your own device" to students acknowledge that a large portion of their student population is unable to bring personal devices.

1.5 The Anxiety-Reality Mismatch: Fear Does Not Map to Risk

Misinformation Drives Career Choices

While 45.9% of students express fear that AI will reduce their employability, this anxiety is distributed unevenly and often irrationally across fields. Students in relatively stable domains like healthcare diagnostics, law, and education report anxiety levels comparable to those in high-disruption fields like content writing, graphic design, and entry-level coding. This suggests that **misinformation and**

generalised fear, rather than field-specific risk assessment, are driving student concerns.

As a result:

- Students are avoiding certain career paths not because AI genuinely threatens those roles, but because they lack clarity on how AI will reshape them.
- Career counselling and placement offices are not equipped to provide nuanced, field-specific guidance on AI's actual impact.
- The students who need the most help, especially those in Tier 2/3 institutions with limited industry exposure, are the least likely to receive accurate information.

What is needed: Career services and faculty must shift from vague reassurances ("AI won't replace you") to specific, evidence-based discussions of how AI is changing roles in each domain and what new skills become important.

Student anxiety about AI rendering them unemployable is pervasive across all disciplines. Yet anxiety levels do not correspond to actual displacement risk. This mismatch creates poor career choices, defensive skill-building, and erosion of confidence at a time when adaptability is most needed.

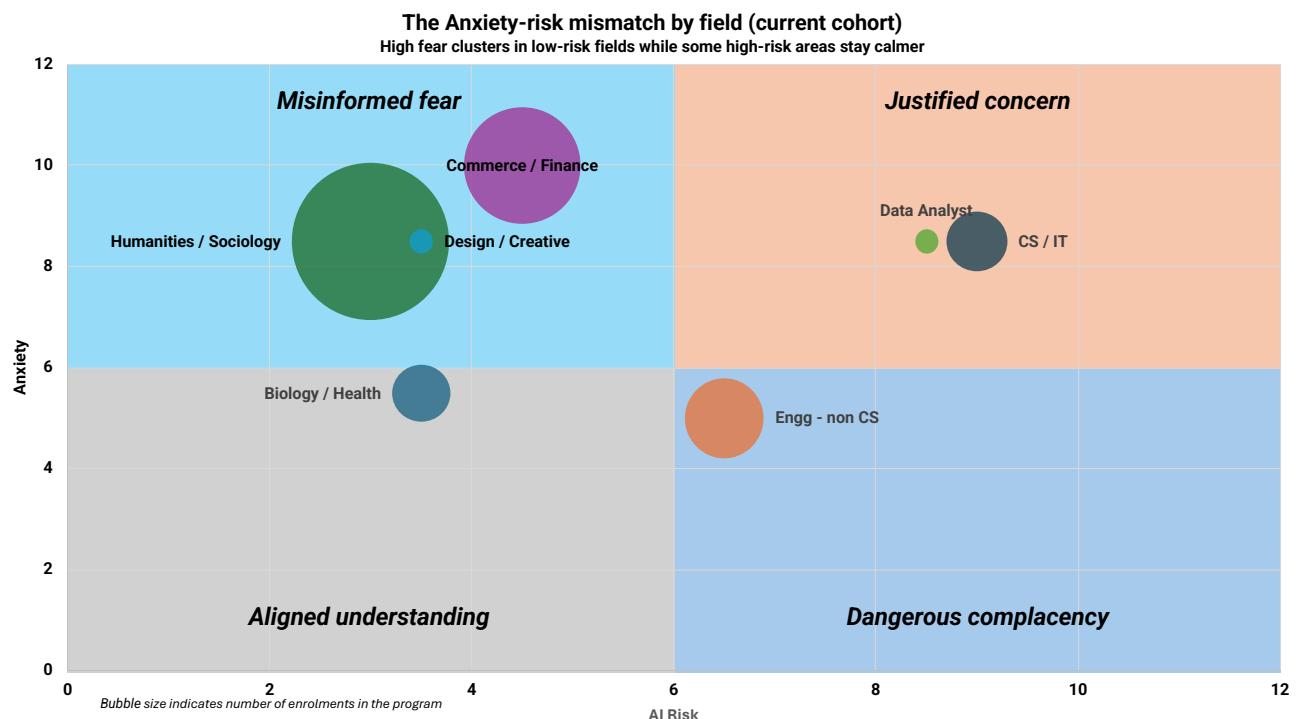


Figure 3: Illustrative - Job anxiety vs actual AI risk matrix revealing significant misalignment, with Commerce and Humanities students experiencing disproportionate fear relative to actual displacement risk.

The matrix reveals four zones:

Zone 1: Misinformed Fear (High Anxiety, Low Actual Risk)

This is where most students cluster. Commerce, finance, and management students express very high anxiety, often saying "AI has taken over finance jobs" or "I'm reconsidering my career because AI will automate everything I was planning to do." Yet the actual risk is moderate and domain-specific. Routine financial reporting and basic analysis are indeed being automated. But advisory,

risk assessment, client relationship management, and strategic financial planning are expanding. Students are abandoning promising career paths based on misinformation.

Humanities and social sciences students express similar fears that "AI will write all the articles", "AI will analyse all the surveys" or "There won't be jobs for us". Yet fields requiring cultural context, ethical judgment, synthesis across sources, and human-centred interpretation face low displacement risk. AI augments these capabilities; it does not replace them. As an example, a historian using AI to process archives faster produces better scholarship. AI does not replace the historian's interpretive judgment.

Design and creative students worry about "AI-generated designs replacing us." Yet employers hiring designers report they need distinct points of view, curatorial taste, and the ability to translate ambiguous client needs into coherent visual languages. These are all low-automation-risk skills. AI-generated designs are commoditising basic execution (like logos and mockups), raising the premium on conceptual thinking and originality.

Zone 2: Justified Concern (High Anxiety, High Risk)

Computer science and IT students pursuing routine coding roles face real threats in the job market. Entry-level roles focused on boilerplate code, unit testing, and scaffolding are shrinking. Data analysts focused on reporting and dashboard creation face similar challenges. Anxiety in these zones is justified, but the response should be upskilling toward systems design, architecture, and interpretive work, rather than abandoning the field altogether.

Zone 3: Aligned Understanding (Low Anxiety, Low Risk)

A small minority of students, often those with mentorship or industry exposure, have realistic assessments of AI's impact. Healthcare students pursuing patient-facing roles recognise that clinical judgment and care remain human domains. They are not anxious about displacement; they are, however, focused on learning how AI can assist diagnosis and documentation.

Zone 4: Dangerous Complacency (Low Anxiety, High Risk)

A few students pursuing routine roles (data entry, basic reporting, simple coding) express low anxiety because they lack information about how their roles are changing. This is less common but consequential, as they will be displaced without preparation.

The mismatch in Zones 1 and 4 implies an urgent need for **disciplined, evidence-based career guidance**. Students need to shift their thinking from anxiety over "will AI take my job?" to an understanding of "which tasks in my field are automating, which are augmenting, and what capabilities should I build?"

Regional Use Cases: The Missed Opportunity

One of the sharpest insights from employer and EdTech conversations is that **generic AI training is becoming commoditised; the premium is in localised, domain-specific application of AI**. A healthcare startup hiring for patient triage workflows wants graduates who understand how AI can improve outcomes in rural Indian clinics with limited infrastructure. To this startup, a student who completes a generic "Healthcare AI" course using international datasets is not of relevance. A logistics company wants graduates who can apply AI to Indian supply chains with unpredictable demand, monsoon disruptions, and fragmented last-mile delivery. Once again, mere training on optimised Western supply chain models are not of relevance to this logistics firm.

Yet curricula remain stubbornly generic, lacking in the Indian local context. Students learn machine learning on ImageNet. They study finance using American case studies. They analyse healthcare data from NIH datasets. When they graduate and apply for Indian roles, employers find they lack the contextual grounding in Indian domestic realities to apply their knowledge.

Regional use cases are a competitive advantage that India is not leveraging. Consider what localised AI training could look like:

- **Agriculture:** Students in Punjab learning AI for wheat disease detection using local crop varieties, soil data from Indian farms, and monsoon-adjusted yield models. Upon graduation, they are immediately hireable by AgriTech startups, government extension services, and farmer cooperatives.
- **Healthcare:** Students in Kerala learning AI for diagnosis workflows tailored to Indian disease burdens (diabetes, tuberculosis, maternal health), using data from Indian hospitals, optimised for low-resource settings. They graduate understanding not just "AI in healthcare" but "AI in Indian healthcare contexts."
- **Financial inclusion:** Students in Maharashtra learning AI for credit assessment of unbanked populations, using alternative data (mobile payments, utility bills, social networks), calibrated for Indian risk profiles. They become immediately valuable to fintech startups, microfinance institutions, and NBFCs.
- **Urban services:** Students in Bengaluru learning AI for traffic prediction, water distribution optimisation, and power grid management using Bengaluru's actual infrastructure, data, and constraints rather than models tuned for Singapore or San Francisco.

The infrastructure for this exists. EdTech platforms can partner with local industries to curate problem libraries. Colleges can collaborate with regional companies to scope capstone projects. The government can fund dataset creation (anonymised, ethically sourced) for regional sectors. Yet this remains largely undone. The result: graduates trained on generic problems competing in global talent markets where they have no advantage, while Indian employers complain about a talent shortage for India-specific problems.

1.6 Access to Datasets and Compute: The Silent Gatekeepers

Beyond devices, two other infrastructure constraints are quietly limiting who can develop deep AI capabilities and who cannot. These constraints are **"access to quality datasets"** and **"access to compute resources"** (GPUs, cloud credits).

Dataset access matters because real AI learning takes place when students work with messy, real-world data, not clean Kaggle competitions. Yet most institutions lack partnerships with industries that can provide anonymised datasets. Students in elite institutions benefit from faculty research collaborations (accessing hospital data, financial transaction data, logistics datasets through faculty connections). Students in Tier 2/3 institutions work (if at all) only with public datasets, which are either international (not relevant to Indian contexts) or sanitised (not representative of real-world messiness).

This creates a capability gap. A student who has cleaned, validated, and built models on real hospital data from Indian clinics knows how to handle missing fields, inconsistent formats, regional language

variations, and data quality issues. A student who has only worked with Kaggle's pre-cleaned datasets does not. Employers hiring for applied AI roles can immediately distinguish between the two.

Compute access determines what experimentation is possible. A student with GPU lab access or cloud credits can run dozens of model variations, test different architectures, train custom agents, and experiment with parameters. A student limited to free-tier tools with CPU-only compute cannot. While such a student can conceptually understand AI, he or she cannot build fluency through hands-on iteration.

Tier 1 institutions increasingly provide GPU labs and cloud partnerships (AWS Educate, Google Cloud credits). Tier 2 and 3 institutions rarely do. The result: Tier 1 students build portfolios with deployed projects trained on real compute. Tier 3 students build theoretical understanding but cannot demonstrate practical capability. In portfolio-based hiring, this gap can be decisive.

Karnataka is piloting a promising model: **state-level Technology Business Incubators (TBIs) with shared GPU clusters**. For example, a Tier-3 college in the Hubballi–Dharwad–Belagavi cluster could, in principle, access high-end GPUs via a state-funded TBI or AI CoE, instead of building its own GPU lab. This is exactly the kind of shared model India needs at scale. This shared infrastructure model could nationally rather than expecting each college to build its own GPU labs. Yet few states have adopted this model.

The Iceberg of Latent Innovation⁶

While the Indian higher education system is often critiqued for its rigid curricula and slow pace of change, a massive, unobserved transformation is occurring beneath the surface. To understand the 18-month strategic window of opportunity, it is important to look beyond the "Visible Tip" of institutional data and engage with the "Submerged Reality" of student behaviour.

I. The Visible 20%: Superficial Productivity

The "Visible Tip" represents the AI usage that faculty and administrators currently see and, in many cases, attempt to police. This layer is characterised by **Superficial Productivity**:

- Usage is concentrated in basic tasks such as summarising long PDFs, cleaning up grammar for emails, or generating boilerplate code for introductory lab assignments.
- Because many institutions still operate under a "policing" mindset, students use AI primarily to bypass traditional hurdles. This results in a "cat-and-mouse" game of AI-detection software, which provides no pedagogical value and only increases institutional friction.
- At this level, AI is treated as an "external addon" rather than a core thinking partner. It is used to save time, but not necessarily to deepen domain expertise.

⁶ The "Iceberg of Latent Innovation" is a thematic deduction synthesised from 85+ primary stakeholder conversations. It illustrates the disconnect between institutional visibility and the "Shadow Curriculum" identified in Lathika, A. (2025), *The Shadow Curriculum: How Students Are Rebuilding Higher Education with AI*.

THE AI LEARNING ICEBERG

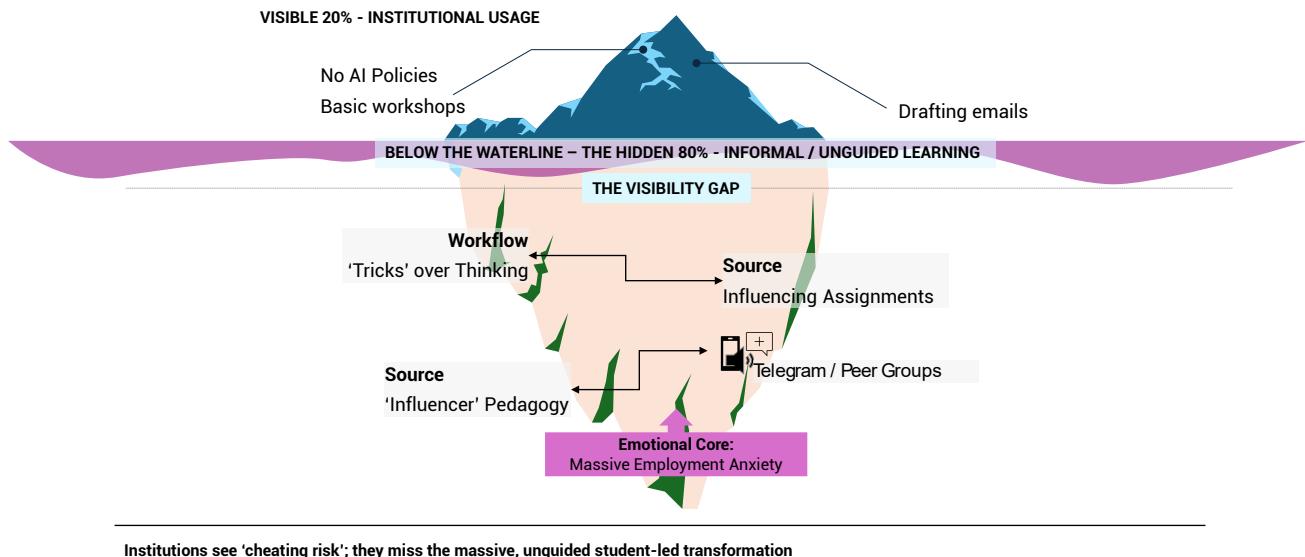


Figure 4: The Learning Iceberg: While institutional frameworks focus on the visible 20% of AI usage, a massive reservoir of student innovation remains unguided below the surface. Bringing this 'hidden mass' above ground is the key to India's AI Dividend.

II. The Hidden 80%: The Shadow Curriculum

Beneath the surface lies the "Shadow Curriculum"⁴ - a massive, peer-driven educational layer where students are rebuilding their own learning frameworks faster than their colleges can respond. This submerged reality is driven by both **Fear and Resourcefulness**.

- **The 45.9% Anxiety Gap:** Nearly half of the student body views AI through a lens of "Employment Anxiety". This fear of obsolescence acts as a high-octane fuel, driving students to spend 4–6 hours daily on YouTube, Discord, and GitHub, teaching themselves the tools they believe their college will never provide.
- **Unguided Peer Networks:** In the absence of faculty mentorship, students have built informal "Intelligence Circles." Here, a Tier 2 student in a regional town might be learning agentic workflows from a peer in a different city, bypassing the formal classroom entirely.
- **Tactical Hustlers:** This layer is where the "Hustle-Focused Multitaskers" live. They are using free-tier tools and shared mobile data to build portfolios that bypass traditional credentials. They aren't looking for a degree; they are looking for **Market Relevance**.

III. The Leadership Risk: Missing the "Judgment Layer"

Student use of AI **without professional scaffolding** means that 80% of innovation happens "underground". It lacks the **Verification Layer** that only faculty can provide. Without mentorship, the "Shadow Curriculum" produces graduates who are technically proficient but lack the **Professional Judgment** to know when an AI output is contextually wrong, ethically compromised, or technically flawed.

IV. Strategic Mandate: Mainstreaming the Submerged 80%

The goal for institutional leadership over the next 18 months is not to dismantle the Shadow

Curriculum, but to mainstream it.

- **Validation over Policing:** Acknowledge the "underground" learning. When a student builds a project via the Shadow Curriculum, provide an institutional pathway to "Audit and Credit" that works.
- **Reducing the Anxiety:** Bringing AI into the formal classroom allows institutions to convert the 45.9% experiencing employment anxiety into strategically empowered graduates.
- **Closing the Substantial Gap:** The Iceberg is largest in Tier 2 and Tier 3 institutions, where the silence from administration is the loudest. Bridging this gap requires moving the "Shadow Curriculum" into the light of the formal syllabus.

1.7 Persona Map: Who India's Students Are

Beyond statistics and aggregate patterns, it is useful to map the distinct student archetypes that emerged across dozens of interviews, focus groups, and survey responses. These personas are not caricatures, but they are composites of real students whose experiences reveal how the intersection of **access to resources** and **AI capability development** creates divergent employability trajectories.

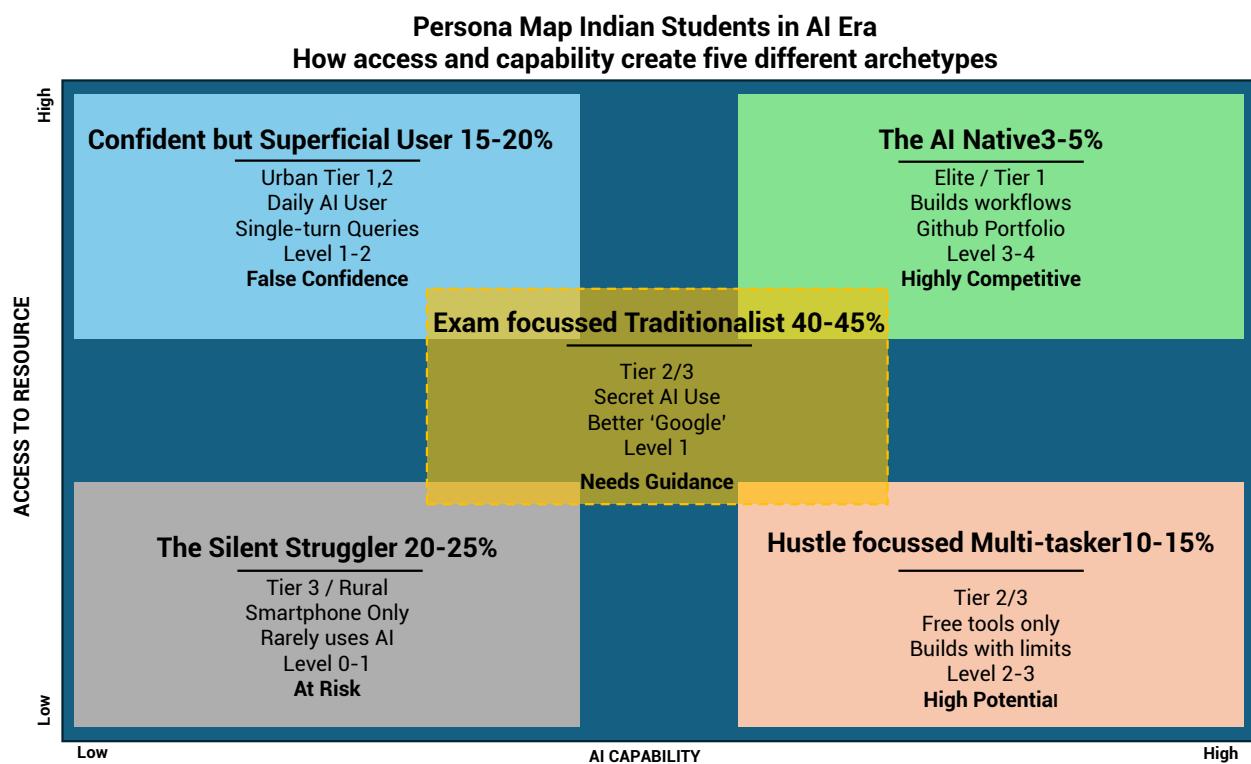


Figure 5: Student persona matrix mapping five archetypes based on access to resources and AI capability levels, revealing that 60-70% of students cluster in low-capability zones despite varying access levels.⁴

The matrix reveals five distinct archetypes, each with different advantages, vulnerabilities, and intervention needs:

The AI-Native

Top-Right Quadrant: High Access, High Capability | Population: 3-5% of students

Priya is a third-year computer science student at a Tier 1 engineering college. She uses Claude for complex reasoning tasks, Perplexity for research, and runs local models for experimentation on her laptop. She has built an AI-powered research assistant that processes academic papers, generates summaries, and suggests connections across literature. Her GitHub has 23 repositories. She mentors juniors at weekend hackathons and has clarity on the roles she wants: AI solution architect or applied AI researcher at a product company.

What gives her an advantage? A compounding set of privileges: a Tier 1 institution with GPU labs, faculty conducting cutting-edge research who share industry problems, peer communities obsessed with building, mentors from alumni networks at top companies, a laptop with adequate RAM, and an institutional culture that encourages experimentation rather than policing AI usage.

Employability outlook: Highly competitive. She will receive multiple offers, likely with premium compensation. Her challenge is not so much about getting hired, but about choosing among the multiple options.

Intervention needs: None for improving employability, but Indian language model access and local datasets exposure will give her the required depth. Institutions should, however, study what enables her relative advantage and work to democratise those conditions.

The Confident but Superficial User

Top-Left Quadrant: High Access, Low Capability | Population: 15-20% of students

Rahul is a second-year business management student at a well-resourced college in Bengaluru. He uses ChatGPT extensively for drafting emails, creating presentation slides, summarising case studies, generating ideas for group projects, and preparing for interviews. He owns a MacBook, pays for ChatGPT Plus, and uses multiple AI tools daily. He considers himself "AI-proficient" and lists it prominently on his resume.

The problem? His usage is entirely superficial. He asks single-turn questions, accepts the first output without iteration, never evaluates whether AI responses are contextually appropriate, and has never built a project that demonstrates reasoning or problem-solving. He treats AI as a productivity hack, not a thinking partner. When employers interview him and probe his "AI skills," they quickly discover he cannot evaluate outputs, doesn't understand when AI hallucinates, and struggles to apply AI to novel problems.

Employability outlook: False confidence masking shallow capability. He will likely face rejection in hiring processes that use portfolio assessment or practical auditions. His high access has paradoxically made him complacent, leading him to assume daily usage equals competitive capability.

Intervention needs: Structured evaluation and training, exposure to real problems with local context where superficial AI usage fails, portfolio-building guidance that forces depth over breadth.

The Hustle-Focused Multitasker

Bottom-Right Quadrant: Low Access, High Capability | Population: 10-15% of students

Arjun is a commerce student at a Tier 2 college in Pune. He does not own a laptop, but he shares one with his sister, each getting 2-3 hours per day. He uses free-tier ChatGPT and Gemini (via Jio). Despite these constraints, he is resourceful. He has built a portfolio of marketing case studies by finding local businesses willing to share anonymised data. He uses AI to generate initial analyses,

then cross-checks against business logic, refines outputs, and documents his reasoning. He spends weekends at public hotspots with better Wi-Fi to work on projects.

What drives his success despite constraints? Intrinsic motivation, problem-solving orientation, and hustle. He does not wait for institutional support. He seeks out problems, builds solutions with whatever tools are available, and iterates relentlessly. He treats constraints as puzzles to solve rather than barriers.

Employability outlook: High potential. Employers who use portfolio-based hiring and look for grit and adaptability will value him. But he faces disadvantages in credential-focused hiring (his college lacks brand recognition) and in roles requiring advanced compute (he cannot build models that need GPU training).

Intervention needs: Device access, institutional compute resources, mentorship to refine portfolio presentation, and connections to employers who value capability over credentials. Indian Language models with local datasets that enable contextual exposure to local problems.

The Exam-Focused Traditionalist

Centre Zone: Mixed Access, Low Capability | Population: 40-45% of students - The largest cohort

Lakshmi is a second-year engineering student at a Tier 3 college in Tamil Nadu. She owns a basic laptop with limited RAM. She uses Perplexity (free via Airtel) to get quick explanations of concepts she doesn't understand. She treats AI as "better Google", which is useful for finding answers during exam prep, but not for deep learning. She has never built a project, doesn't know what a portfolio looks like, and is anxious about jobs but unclear on what to do.

Her faculty sends mixed signals: some ignore AI entirely, others threaten grade penalties for "AI-generated assignments," and none provide structured guidance on how to use it wisely. So, she uses AI, secretly copying explanations for assignments, rewriting to avoid detection, and feeling guilty about it. She thinks of AI as "cheating" but does it anyway because "everyone else is doing it."

The tragedy? She is trapped in an institutional vacuum. She is not lazy or uninterested. She is responding rationally to unclear policies, absent mentorship, and faculty who themselves don't know how to teach in an AI era. Her superficial usage is the result of systemic failure, not individual deficit.

Employability outlook: Vulnerable. She will graduate with "AI exposure" on her resume, but no genuine capability. In hiring processes that test for depth (portfolio review, task auditions), she will struggle. Yet she represents 40-45% of India's graduates, which is the cohort at highest risk of the access-capability gap.

Intervention needs: Clear institutional policies on AI usage, faculty training in AI pedagogy, structured project-based assignments that require iteration and evaluation, career clarity programs, and portfolio-building scaffolding, and Indian Language models that break learning barriers.

The Silent Struggler

Bottom-Left Quadrant: Low Access, Low Capability | Population: 20-25% of students

Meera is a humanities student at a rural college in Chhattisgarh. She owns only a low-end smartphone (32GB storage, 2GB RAM) with intermittent 3G connectivity. She has heard about ChatGPT from news articles and YouTube videos, but has never used it. She tried once, but the browser crashed, and she gave up. Her college has no computer lab with adequate internet, no faculty guidance on AI,

and no policy (because the administration hasn't considered the issue).

She does not know what AI could do for her field, like how it could assist research, improve her writing, help analyse qualitative data, or prepare her for communication-intensive roles. She is anxious about the future, sensing that "something big is happening with technology" but feeling entirely excluded from it. She wonders if her degree will be worthless.

What blocks her? Not intelligence, not motivation. Systemic barriers: device poverty, connectivity constraints, institutional silence, absence of local mentors, and a curriculum that treats AI as irrelevant to humanities. She is not "falling behind", rather, she was never given a chance to start.

Employability outlook: At serious risk. She will graduate without any AI exposure in an era when baseline AI literacy is becoming non-negotiable. Even entry-level communication, research, or administrative roles now expect comfort with AI-assisted workflows. She will compete against peers who have that comfort; she will not.

Intervention needs: Urgent and comprehensive. Device access programs (subsidised laptops or community labs), Indian Language AI interfaces (Hindi, Chhattisgarhi), faculty trained to integrate AI into humanities pedagogy, mentorship from professionals in her region, and structured awareness campaigns showing how AI applies to her field.

What the Persona Map Reveals

The distribution of students across these archetypes exposes three structural realities:

First, most Indian students (60-70%) are clustered in low-capability zones despite varying levels of access. The Exam-Focused Traditionalist (40-45%) has moderate access but no depth. The Confident Superficial User (15-20%) has high access but no evaluative skill. The Silent Struggler (20-25%) has neither access nor capability. Only 15-20% of students (The AI-Native and The Hustle-Focused Multitasker) are developing genuine capability that translates to employability.

Second, high access does not guarantee high capability. The existence of the "Confident Superficial User" archetype, comprising of students with laptops, paid tools, and daily usage but shallow capability reveals that infrastructure is necessary but not sufficient. Without mentorship, structured problem exposure, and evaluation training, access alone produces overconfident but underprepared graduates.

Third, capability can emerge despite low access, but it requires extraordinary hustle. The "Hustle-Focused Multitasker" proves that resourceful, self-directed students can build capability with minimal resources. But this should not be the norm. Relying on individual hustle to overcome systemic barriers is neither equitable nor scalable. Institutions and policy must create pathways where capability development does not require heroic effort.

The archetypes also clarify intervention priorities with the need for Indian language models access and exposure to local datasets, applying to all of them:

- **For The AI-Native:** Study what enables their success; work to democratise those conditions (mentorship, problem access, institutional support).
- **For The Confident Superficial User:** Evaluation training, portfolio-building that forces depth, exposure to problems where superficial AI usage fails.
- **For The Hustle-Focused Multitasker:** Remove barriers (device access, compute resources, mentorship networks) so capability can flourish without requiring hustle.

- **For The Exam-Focused Traditionalist:** Clear institutional policies, faculty training, project-based learning, and career guidance to convert latent interest into genuine capability.
- **For The Silent Struggler:** Comprehensive interventions such as device access, faculty training, mentorship, and awareness programs showing how AI applies to their disciplines.

What This Means for Policy and Institutions

The exposure landscape reveals three critical insights:

First, access is no longer the binding constraint for the majority, but capability development is. Most students have heard of AI, have used it at least once, and can access free tools. The challenge is not awareness or basic access. It is moving students from Level 1 (superficial usage) to Level 3 (genuine capability). This requires mentorship, structured experimentation, exposure to real problems, and institutional environments that support rather than police AI usage.

Second, infrastructure inequality, especially devices and compute, is quietly creating a two-tier system. Students with laptops and institutional compute access are building portfolios that will be competitive in hiring. Students with only smartphones and free-tier tools are developing theoretical knowledge, but cannot demonstrate practical capability. This gap will widen unless institutions prioritise device access and shared compute infrastructure.

Third, the anxiety-reality mismatch is driving poor career decisions and eroding confidence. Talented students are abandoning promising fields because of misinformation about AI's impact. Others are pursuing defensive skill-building (chasing credentials, doing generic AI courses) instead of building genuine problem-solving capability. Institutional and policy responses must include disciplined career guidance grounded in labour market data, not hype.

The section that follows examines how the industry is experiencing this landscape. Specifically, what signals are visible in hiring, what capabilities employers actually need, and where the academia-industry alignment breaks down.



2

INDUSTRY LENS - WHAT EMPLOYERS ARE SAYING

2.1 The Ground Truth: How AI Is Actually Reshaping Work

Industry conversations reveal a reality strikingly at odds with both popular hype and prevailing student anxiety. AI is not eliminating jobs broadly. Rather, it is **compressing workflows, raising capability bars, and revealing that the real bottleneck is not AI adoption, but it is organisational uncertainty and misaligned role design.**

Walk through a manufacturing company's planning department, and you see workflow compression in action. Material planning used to take 70 people. Today, AI-optimised forecasting, demand signals, and inventory algorithms handle the work. **Two people remain, but their role is transformed.** They don't process data; they handle exceptions, interpret signals, make judgment calls on supply risk, and override algorithms when context demands. Their work is harder intellectually, not easier.

This pattern repeats across sectors. A retail organisation's finance function shifted from 250 people to 40. A software company that would have hired five or six junior developers now hire three, each more productive because they work with copilots. A consulting firm no longer needs five junior consultants drafting proposals when two can do the work faster and better. A startup with two-person engineering teams produces outputs that would have required five or six persons per team in the pre-Generative AI era.

This is not job elimination. This is workflow compression. And crucially, the roles that remain demand fundamentally different capabilities.

2.2 The Capability Shift: From Execution to Judgment

Every employer interviewed articulated the same core shift: the work that remains after AI handles the routine is work that demands judgment, critical thinking, problem-framing, and contextual understanding. Routine execution that entails following standard operating procedures, processing data based on rules, generating reports, and writing first drafts is increasingly AI-handled.

As one CHRO from a consulting firm observed, *"Functional degrees are losing signalling power. We hire for intellectual agility. Can you ask nested questions? Can you see a problem from multiple angles? We don't care whether you studied finance or literature. We care whether you think well."*

This shift manifests in hiring signals:

Critical Thinking and Problem-Solving are now universally non-negotiable. Yet colleges continue to teach content delivery and procedural execution. A student who can memorise finance theory but struggles to think through ambiguous problems is increasingly non-hireable.

Adaptability and Learning Agility outweigh specific technical depth. Employers recognise that what they are hiring for will change in 18 months. They prefer someone who is comfortable being wrong, iterating, and learning over someone who is an expert in current tools.

Social Intelligence and Communication are now competitive advantages because AI cannot replace persuasion, empathy, cultural sensitivity, and the ability to translate between technical and business contexts. Yet colleges do not systematically teach these skills.

Linguistic Precision has become surprisingly important. Prompt engineering (in the broad sense) requires the ability to ask questions clearly, refine based on responses, and evaluate nuance. English proficiency is now a core employability skill, not a bonus. Non-English-first students face a compounding disadvantage.

Portfolio Evidence increasingly beats degrees. Employers no longer want to see claims of "I know Python" but want to see evidence of "I built X, here's the GitHub, here's how I debugged Y." The shift is from credential signalling to demonstrated capability backed by evidence.

2.3 The Entry-Level Compression: Fewer Doors, Higher Bars

While roles are not disappearing, **entry-level opportunities are compressing**. The traditional skill ladder, where you hire juniors to do basic work and then gradually move them up is being disrupted by AI.

A software engineering manager notes, *"Intern-level work is now automated. Copilots handle boilerplate, scaffolding, and unit testing. We can't give interns the traditional grunt-work progression anymore. They either come in with systems-thinking capability, or they are not hireable."*

This compression follows a pattern:

Year 1 of AI adoption: Efficiency gains in routine work. Workflows streamlined. Headcount reduced in back-office functions. Entry-level hiring pauses because automation eliminates the volume of basic tasks.

Year 2-3: Organisations redesign roles. They realise they can't just remove people; they must redefine what people do. Roles shift from "execute tasks" to "design, oversee, and validate." Hiring cautiously resumes for these redesigned roles, but demands are higher.

Year 3+: New equilibrium emerges. Fewer entry-level positions overall, but they require higher capability. A new category of roles, such as AI orchestrators, systems re-designers, domain-AI hybrids are emerging as AI usage begins to scale.

The risk: **organisations that don't redesign entry-level roles quickly will lose talent pipelines.** A company that continues to hire "data analysts" expecting them to do reports will find that the reports are now AI-generated, and the junior analyst has nothing to do. Employers who explicitly redesign ("Our analysts now validate AI-generated reports and focus on business implications") can successfully hire and develop talent.

Most organisations, however, are caught in ambiguity. They know roles are changing, but don't know how to redefine them. This uncertainty translates to hiring freeze, depressed entry-level recruitment, and mounting anxiety among graduates competing for fewer, ill-defined positions.

2.4 The Organisational Uncertainty: The Real Barrier

Across every interview, a consistent theme emerged: **technology is not the bottleneck. Organisational clarity is.**

A manufacturing company's CHRO reported: *"We know AI will reshape everything. But we don't know what AI-ready talent looks like. Should we hire data engineers? AI specialists? People with strong domain knowledge? We don't even know what skills we need, so we are not aggressively hiring AI talent. We are cautiously hiring people who seem adaptable."*

This uncertainty manifests as:

Vague job descriptions: Companies still post for "AI engineer," "data scientist," "ML specialist" without clarity on whether they want deep AI research or the ability to use AI tools in a domain. Candidates don't know what to prepare for.

Broken hiring processes: Companies maintain legacy formats (coding rounds and theory questions) while needing portfolio evidence and problem-solving capability. The mismatch creates noise and misalignment.

Internal confusion about AI strategy: Many organisations have conducted pilots (AI chatbot, automated reporting and copilot rollout) that succeeded in isolation. But rolling out beyond early adopters fails because the organisation hasn't aligned job redesign, incentives, accountability, and hiring criteria.

Fear and resistance: Employee anxiety about job security, though not always justified, creates resistance that slows adoption more than technology limitations. Managers lack confidence in AI and are uncertain about risks and guardrails.

One leader articulated it plainly: *"The bottleneck isn't building AI systems. It is in getting organisations to actually use them, trust them, and redesign work around them. That's a cultural problem, not a technology problem."*

2.5 The Skills Reality Check: What Employers Actually Want

Across dozens of hiring leaders' statements, a clear hierarchy emerges, and it is not what students typically prepare for.

The Skills Map: What Employers Actually Priorities

AI Fluency Matters, but Critical thinking and adaptability leads

	Techsoft	Consulting	Finance/Bank	Retails/Ops	Non-Tech	Ed/Techlearn
AI Tool Fluency	Medium	Medium	Medium	Medium	Medium	Medium
Critical Thinking	High	High	High	Medium	High	High
Social Intel	Medium	High	Medium	Medium	Medium	Medium
Adapt/Learn	High	High	Medium	High	High	High
Domain Knowledge	Medium	Medium	High	Medium	Medium	Medium
Portfolio	Medium	Medium	Medium	Medium	Medium	Medium
Data & Analyst	High	Medium	Medium	Medium	Medium	Medium
System Thinking	Medium	Medium	Medium	Medium	High	Medium
AI Ethics	Medium	Medium	Medium	Low	Medium	Medium
Linguistics	Medium	Medium	Medium	Medium	Low	Medium

Figure 6: Illustrative map showing employer hiring priorities across six segments, revealing that critical thinking and adaptability are universally valued, while domain-specific knowledge importance varies by sector.

The top tier of prioritised skills (across almost all sectors):

- Critical Thinking & Problem-Solving:** Ability to break down ambiguous problems, consider multiple approaches, and to iterate. Employers universally cite this as the #1 gap. Students arrive with answers but not the ability to think through problems.
- Adaptability & Learning Agility:** Comfort with ambiguity, willingness to learn new tools, and

resilience to endure being wrong. This can be a serious differentiator, as an example, a curious person from a Tier 3 college can potentially outcompete an IIT graduate with low curiosity.

3. **Communication & Social Intelligence:** Not just clarity, but the ability to translate between technical and business contexts, persuade, and understand cultural nuance. AI cannot replace this.
4. **Learnability & Ownership Mindset:** Hunger to understand how things work, proactivity, willingness to take on problems without explicit instructions. Surprisingly important.

The second tier (still important, but more role-specific):

1. **AI Tool Fluency:** Ability to use multiple tools strategically, understand their strengths and limitations, and combine them for workflows. This is now baseline across many roles.
2. **Data & Analytics Skills:** In every sector, SQL literacy and hypothesis-driven thinking are becoming mandatory. Yet many graduates graduate without these skills.
3. **Domain Knowledge:** Sector-specific understanding of business, workflows, and customer needs. But much less weight than teams used to give it.

The third tier (important but not the primary filter):

1. **Specific Technical Depth:** Deep machine learning, advanced algorithms, specialised frameworks. Important for research and cutting-edge roles, but not for most entry-level positions.

Crucially, **AI tool fluency sits in the second tier, not the top**. Yet students often prioritise it above critical thinking or communication. A student who completes "30 AI courses" but struggles with creative problem-solving and has weak communication will underperform against a student with strong thinking skills and moderate AI exposure.

2.6 The Role Redesign Challenge: Why Entry-Level Hiring Hasn't Collapsed (Yet)

Despite workflow compression, entry-level hiring has not collapsed. Here's why: organisations that explicitly redesign roles continue to hire.

A consulting firm provides an example: *"We used to hire junior consultants to draft parts of proposals. Now AI handles drafting. So, we redefined: juniors now own small engagements end-to-end, using AI for drafting while focusing on client interaction and strategic thinking. We are still hiring, but we are hiring for different capabilities."*

Companies that haven't redesigned, by contrast, are frozen. They still post job descriptions from the pre-AI world, expecting juniors to do work that AI now handles. When juniors arrive with that expectation, neither party is satisfied.

The research from recruiting firms confirms this: **entry-level hiring will shrink in the short term (as organisations pause to redesign) but will not collapse long-term**. The new entry level, however, will demand higher capability and will expect a faster ramp-up.

2.7 Emerging Roles: The New Opportunity Landscape

One of the most misunderstood aspects of AI's impact on employment is the myth that "new AI jobs" will absorb the disruption. This is both true and false.

It is false if you imagine millions of "AI engineer" or "prompt engineer" roles. Those roles exist but are limited in volume. The AI engineering market is not large enough to absorb all displaced workers.

It is true, however, if you broaden the lens: new roles are emerging, but they are not purely technical. They span:

Table 1: Emerging roles which are an enhancement of existing roles

AI Orchestration Roles	GTM engineers, workflow designers, AI implementation managers. These people trial tools, run POCs, roll out useful AI workflows across teams. They don't need to be AI researchers – they need to be bridge-builders between technology and business problems.
Domain–AI Hybrid Roles	Medical AI specialists who understand both healthcare and how to apply AI to improve outcomes. Finance AI roles focused on risk assessment and compliance. Supply chain AI roles. These are emerging rapidly because they require domain expertise plus AI fluency – a rare combination.
Governance and Ethics Roles	AI model auditors, risk testers, ethics officers, data governance specialists, bias mitigation experts. Financial services, healthcare, and regulated industries are creating these roles at scale.
Data and Infrastructure Roles	As enterprises build private language models, demand for data engineering, data quality, and data governance is surging. This is emerging as the real bottleneck – not model training, but data readiness.
Niche Operational Roles	Token optimisers (in AI-native startups), prompt strategists focused on linguistic and cultural adaptation, learning experience designers for AI-augmented education.

The common thread: **these are not pure AI roles. They are domain roles that have been enhanced or created by AI.** The opportunity is not in becoming an "AI person" but in becoming an expert in your domain who also deeply understands AI applications within that domain.

Yet most students are pursuing the opposite: learning AI in isolation, not domain–AI integration. This misalignment is driving the anxiety-skills mismatch we see.

2.8 The Hiring Process Transformation: What is Changing

Across forward-thinking organisations, hiring processes are shifting in distinct ways:

Resume Screening is Vanishing: A consulting firm noted, "We stopped using resumes for initial screening. AI handles first-pass filtering. What matters is portfolio work and a problem-solving task audition."

Multi-Round Interviews Are Compressing: Instead of 4-5 rounds over 3 months, leading companies now do: a 4-hour task audition (realistic problem), one technical conversation, and one culture conversation. Done in 2 weeks.

Practical Auditions Over Credentials: "We give candidates a real problem we are facing, give them 4 hours, and see how they approach it. The final answer matters less than the process, how they used AI, where they applied judgment, how they explained reasoning."

Portfolio as Primary Evidence: GitHub repositories, project documentation, and case studies showing problem-solving process. Credentials are secondary.

Culture Fit Validation, Not Primary Filter: In-person meetings are now used primarily to verify authenticity and cultural alignment, not to evaluate capability (which is assessed in task auditions).

AI-Supported Interview Panels: Some organisations are experimenting with AI-neutral evaluators to reduce bias in assessments and structured interviewing.

The net effect: hiring is **faster, more merit-based, and less influenced by gatekeepers vetting credentials**. This should advantage Tier 2/3 city/town graduates if they can demonstrate capability through portfolios. But it disadvantages those without access to mentorship and guidance on how to build portfolios.

2.9 The Equity and Access Lens: Who Benefits, Who Loses

Industry conversations on equity revealed three critical insights:

First, access to tools has been democratised but mentorship and context remain concentrated. An AI-native startup founder noted, “Access to Claude Pro matters less now. Free models are competitive. The real divide is between students who have mentors showing them ‘Here’s how we use AI at work’ versus students learning from YouTube and Instagram.”

Second, AI is perceived as a leveller by some employers, a concentrator by others. A leading Global GCC reports, “AI-based screening widens our candidate pool by reducing credential gatekeeping. We find good talent in non-elite schools.” By contrast, some tech companies still rely on LeetCode and CGPA filters, which remain correlated with school prestige. The outcome depends on hiring practice design, not on AI itself.

Third, gender gaps in AI fluency are alarming. Only one-third of women in India⁷ have functional AI literacy. If left unaddressed, AI skilling will widen gender gaps in tech and other male-dominated fields. Deliberate interventions (women-focused skilling programs, mentorship networks, role modelling) are urgent.

The data also reveals: **AI fluency is not evenly distributed. Students in metro areas with industry exposure, urban families, English fluency, and device access are able to accelerate. Students in Tier 2/3 towns without these conditions fall behind.** This geographic divide is sharper than gender or caste in determining AI employability.

2.10 What Industry Needs (But Isn't Getting)

When asked, “What would help us hire better?” employers consistently mentioned three gaps:

- Clearer signals of depth.** “We can’t distinguish between a student who did 10 online AI courses and a student who built one real project. We need evidence of application, not certificate accumulation.”
- Basic judgment capability.** “Most students can use a tool. But they can’t evaluate whether the output is correct. They accept AI responses uncritically. We need graduates who can think critically about AI outputs.”
- Discipline-specific AI knowledge.** “We want engineers who understand what GenAI means in mechanical engineering. Architects who know AI-augmented design workflows. We don’t want generic ‘AI students.’”

⁷ Primary Research Synthesis (Nov–Dec 2025). The figure reflects the qualitative consensus gathered from primary conversations with academic deans, EdTech leaders, and industry practitioners

4. **Communication and reasoning documentation.** "Show us your thinking. How did you approach this? Where did you use AI? Where did you apply human judgment? Process matters more than polish."

These are not new demands. They are simply higher stakes now because AI amplifies good thinking and bad thinking equally.

2.11 The 3-Year Outlook: What's Coming in 2027-2028

Based on industry consensus, the employment landscape in 3 years will look like:

Fewer Entry-Level Openings, Higher Bars: Organisations will have completed role redesign. Entry-level hiring will resume, but at a lower volume. The roles that exist will demand higher capability. A "junior analyst" role will include "uses AI to generate draft analyses, focuses on business interpretation and exception handling."

Portfolio Becomes Baseline: Resumes are optional. Portfolios are required. Candidates without demonstrable project evidence are immediately screened out.

AI Maturity Becomes a Hiring Filter: Not "do you use AI?" but "show us projects where you used AI thoughtfully. How did you evaluate outputs? Where did you apply judgment?" The ability to think about AI, not just use it, is the new minimum bar.

Domain-AI Hybrid Roles Explode: By-function AI skills become common. "Healthcare AI," "Finance AI," "Logistics AI" specialists are in high demand. Generic "AI proficiency" loses value.

Linguistic Fluency Become Non-Negotiable: As more AI interaction is text-based, clear communication and linguistic precision matter more. Non-English-first speakers will face significant barriers with the current dominance of English first LLMs.

Mindset Trumps Credentials: Among equally capable candidates, those with higher curiosity, comfort with ambiguity, and intrinsic motivation advance faster.

2.12 What This Means for Educators

The industry lens reveals a profound disconnect: **employers are designing hiring for a different India than the one that institutions are training for.**

Institutions are still teaching:

- Memorisation and content delivery
- Functional skills (finance theory, IT procedures)
- Credentials as a signalling mechanism
- The assumption that a degree is an employment guarantee

Employers are now hiring for:

- Problem-solving capability and thinking agility
- Judgment and contextual understanding
- Portfolio evidence and demonstrated projects
- The assumption that curiosity and adaptability matter more than pedigree

This mismatch is at the core of the employability crisis. It is not that students can't get jobs. It is that they are being trained for the wrong jobs.

The section that follows examines how academia is (or isn't) responding to this shift.



3

ACADEMIA LENS - HOW UNIVERSITIES ARE RESPONDING

3.1 The Institutional Moment: Caught Between Worlds

Universities across India stand at a peculiar moment in their institutional journey. They face a technology that has already penetrated their classrooms, where students are using AI in every assignment, every study session, and every interview preparation. Yet institutions remain largely unprepared to teach in an AI-augmented world. The response from these universities ranges from denial ("We'll ban AI") to experimentation ("We'll integrate it") to confusion ("We don't know what to do").

This section maps how institutions are actually responding in reality, going beyond their public statements. From their assessment policies, faculty practices, curriculum decisions, and student support systems, a picture emerges of an institution-wide crisis of clarity.

3.2 The Geographic and Institutional Divide

India's AI readiness does not distribute evenly. It concentrates in metros, elite institutions, and engineering programs, leaving vast populations structurally disadvantaged.

Table 2: The AI Readiness Level scores across institutional tiers, with Task Depth and Institutional Environment showing the largest gaps.⁸

Institution Tier → Dimension ↓	Tier 1	Tier 2	Tier 3
Access	Medium	Low	Very Low
Usage frequency	High	Medium	Low
Task depth	Medium	Low	Very Low
Skill transfer	Medium	Very Low	Very Low
Institutional environment	Medium	Very Low	Very Low
Career clarity	Low	Very Low	Very Low
Project readiness	Medium	Very Low	Very Low
AI literacy	Medium	Low	Very Low
Composite score	Medium	Low	Very Low

The AI Readiness Level, developed for this study, scores institutions across eight dimensions: access to devices and tools, usage frequency, task depth, skill transferability, institutional environment (faculty training, clear policies, assessment redesign), career clarity, project readiness, and AI literacy. Scores range from Low (critical gap) to High (high readiness). **Tier 1 institutions averaged a Medium**, placing them in the "moderate readiness" band. **Tier 2 institutions were rated Low**, and **Tier 3 institutions are rated very Low**, with average to severe readiness gaps, respectively.

The table reveals where gaps are sharpest. **Task Depth** (the sophistication of how students actually use AI) shows the largest disparity: Tier 1 institutions score medium, Tier 2 score low, and Tier 3 score very low. This means students in Tier 3 students are using AI, but almost entirely at superficial levels of summarisation, simple queries, and copy-paste workflows. **Institutional Environment** is the second-largest gap: Tier 1 institutions score medium (reflecting some faculty training, clearer

⁸ The AI Readiness Level is a metric developed by AI4India.org. The disparity in AI maturity is a synthesised delta derived from 85+ primary stakeholder interviews. This substantial gap reflects a systemic divergence in AI adoption between Tier 1 and regional institutions across three vectors: (a) Faculty Pedagogy Redesign, (b) High-Compute Infrastructure Access, and (c) Institutional Policy Clarity.

policies, industry partnerships), while Tier 3 institutions score very low (reflecting near-total absence of institutional support, policy clarity, or faculty capability).

Interestingly, **Access** (device ownership, internet, tool availability) is not the widest gap. Tier 1 scores medium, Tier 2 scores low, and Tier 3 scores very low. While Tier 3 access is lower, the gap is smaller than the gaps in depth and institutional support. This reinforces a key finding: **access is necessary but not sufficient**. A student with a smartphone and free ChatGPT has access; what they lack is mentorship, structured learning pathways, real problems to work on, and institutional clarity on how to use AI wisely.

Career Clarity is troublingly low across all tiers, where Tier 1 scores medium, Tier 2 and Tier 3 score very low. Even elite students lack a clear understanding of how AI will reshape their disciplines, what roles will emerge, and what portfolio evidence employers expect. Tier 2 and 3 students face even deeper confusion, often relying on hype-driven narratives from social media rather than trusted guidance.

3.3 The Faculty Readiness Crisis: The Silent Bottleneck

The most significant and least discussed barrier to AI-integrated higher education is **faculty readiness**. Across dozens of interviews, the pattern is consistent: institutions recognise AI is transforming education, but have invested almost nothing in preparing faculty to teach in an AI-enabled classroom.

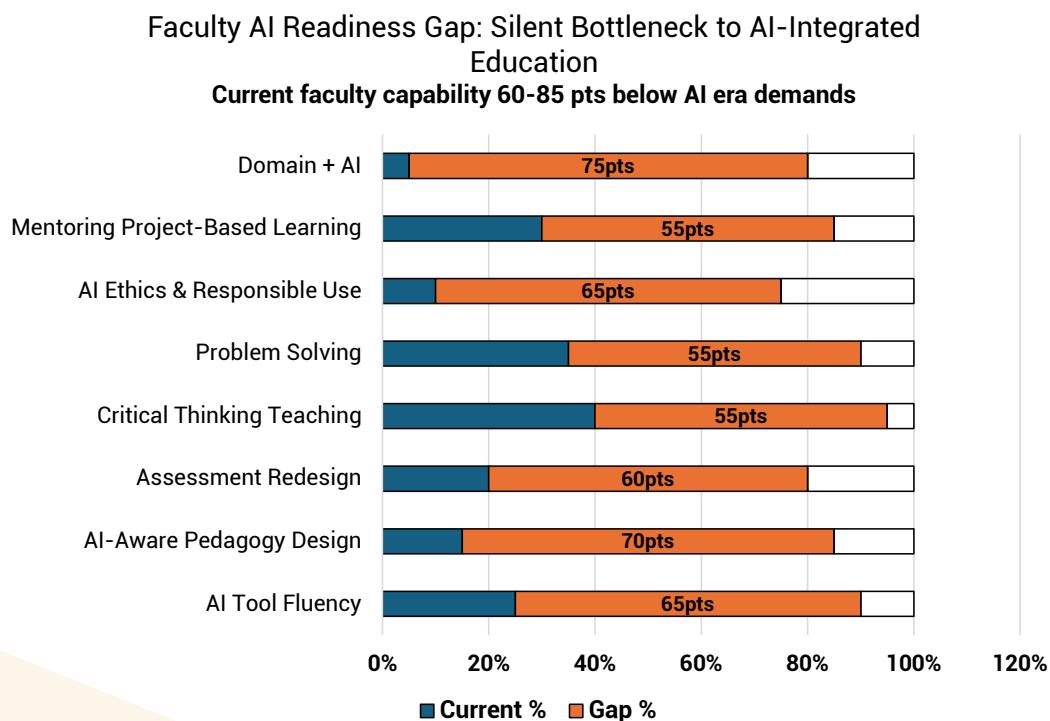


Figure 7: Faculty Readiness Chasm: An analysis of current versus required capability across eight dimensions. Data is synthesised from primary stakeholder conversations, revealing a systemic deficit in the mentorship layers needed to guide the student population.⁹

⁹ Primary Research Synthesis (Nov-Dec 2025). These metrics are qualitative deductions synthesized from 85+ primary stakeholder conversations, including in-depth interviews with CHROs, deans, and faculty members. "Required Capability" represents industry-benchmarked expectations for the 2026 talent market, while "Current Capability" reflects the consensus of faculty self-assessments and administrative observations gathered during field research.

Consider the capability gaps across eight core dimensions:

AI Tool Fluency: Current faculty capability stands at 25%. Required capability for the AI era: 90%. A gap of 65 percentage points. Most faculty have not used ChatGPT meaningfully. They've heard of it, perhaps attended a one-hour workshop, but lack hands-on experience with how to use AI for teaching, research, or even personal productivity.

AI-Aware Pedagogy Design: Current: 15%. Required: 85%. Faculty understand traditional pedagogy (lectures, case studies, problem sets). They do not understand how to redesign assignments for an AI-enabled world. How do you structure an essay assignment when students can generate essays? How do you teach coding when copilots handle syntax? These are not yet-answered questions in most institutions.

Assessment Redesign: Current: 20%. Required: 80%. Exams and assignments designed for a pre-AI world are broken in an AI-enabled world. Yet most faculty continue using them unchanged. They cannot simultaneously assume "students are writing original essays" and "students can use ChatGPT." Assessment redesign requires entirely rethinking what you're evaluating.

Critical Thinking Teaching: Current: 40%. Required: 95%. Paradoxically, this is one of the few dimensions where faculty have some baseline capability (because good teaching has always emphasised critical thinking). Yet the gap remains large because faculty are not systematically trained in how to teach critical evaluation, specifically of AI outputs.

AI Ethics & Responsible Use: Current: 10%. Required: 75%. Most faculty lack training in bias in LLMs, hallucination patterns, responsible AI use, and ethical frameworks. When they encounter these issues in student work, they lack scaffolding to teach responsible usage.

Mentoring Project-based Learning: Current: 30%. Required: 85%. Most faculty are trained to deliver content and grade assignments. Mentoring students through real, ambiguous, project-based work, especially work that integrates AI, is outside the comfort zone of many.

The result: **faculty are overwhelmed, confused, and often default to the safest option: policing AI usage.** They cannot teach what they don't understand, so many choose to ban it instead.

3.4 The Policy Response Spectrum: From Bans to Integration

Institutions are responding to student AI usage across a spectrum of approaches. The distribution reveals the degree of institutional unpreparedness:

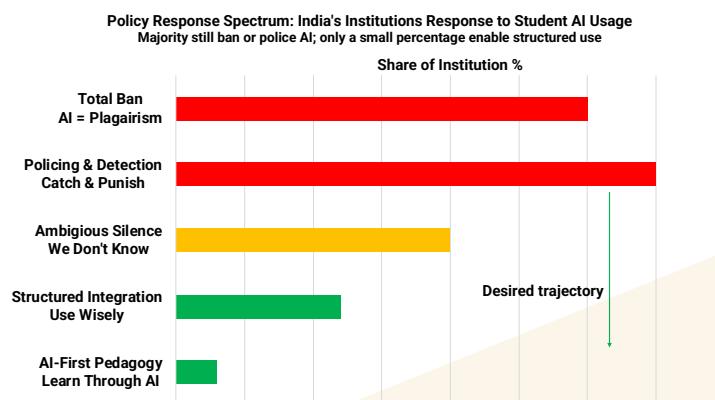


Figure 8: Illustrative visualisation of policy response spectrum across five institutional approaches to student AI usage, showing that the lion's share remains restrictive while a small percentage have moved to integration or AI-first models.

Total Ban | AI ≠ Plagiarism

These institutions treat AI as plagiarism. Policies prohibit the use of all AI tools in assignments. Students who use AI face grade penalties. Detection tools like Turnitin are configured with AI-detection thresholds where any assignment detected to contain AI-generated content above these thresholds is automatically rejected or penalised.

Consequences: Students don't stop using AI; they just hide it. Assessment becomes unreliable because faculty can't distinguish between a student who used AI as a learning tool (legitimate) and a student who used AI as a shortcut (problematic). Underground usage grows.

Policing & Detection | Catch & Punish

These institutions allow AI but heavily monitor it. AI detection tools are deployed; faculty use them as gatekeepers. The assumption is: catch problematic usage and punish it.

Consequences: False positives (such as human writing flagged as AI), student anxiety, inconsistent enforcement. A student gets penalised by one faculty member but not another. Uncertainty drives students toward underground usage.

Ambiguous Silence | We don't know

These institutions have no clear policy. Each faculty member decides independently whether to allow the use of AI. A student might encounter one professor who bans the use of AI and another who encourages it in the same semester.

Consequences: Students adopt risk-averse behaviour ("Better not use AI, I don't know what this professor thinks"). No institutional learning about AI-integrated teaching. No consistency.

Structured Integration | Use Wisely

These institutions (typically elite colleges) have explicit policies: AI is allowed, but with transparency and reflection. Students must disclose their AI usage (prompts used, outputs generated). Assignments are redesigned to evaluate the thinking process, not just the product. Grading rubrics explicitly include "prompt strategy," "critical evaluation of AI outputs," and "human judgment applied."

Consequences: Students begin to distinguish between AI as a learning tool and AI as a shortcut. Underground usage decreases because there's no fear. Genuine AI literacy develops.

AI-First Pedagogy | Learn through AI

A tiny fraction of institutions (primarily IIIT-Delhi's pioneering work) is redesigning pedagogy around AI. Assignments are built assuming students will use AI. The skill being taught is not "do this task" but "think about how to use AI to do this task and justify your approach." Faculty are trained in AI-aware teaching. Assessment directly evaluates critical thinking about AI, not adherence to rules.

Consequences: Students develop sophisticated AI literacy. They learn to evaluate, iterate, and apply judgment. They graduate AI-ready.

The critical insight: **The majority of Indian institutions are still banning or policing AI, while only a few have moved toward integration or AI-first approaches.** This distribution explains why graduates enter the workplace with superficial AI exposure but no genuine capability.

3.5 The Underground Learning Pathway: How Students Actually Learn

One of the most striking findings across student interviews is the disconnect between what institutions are trying to teach and what students are *actually* learning.

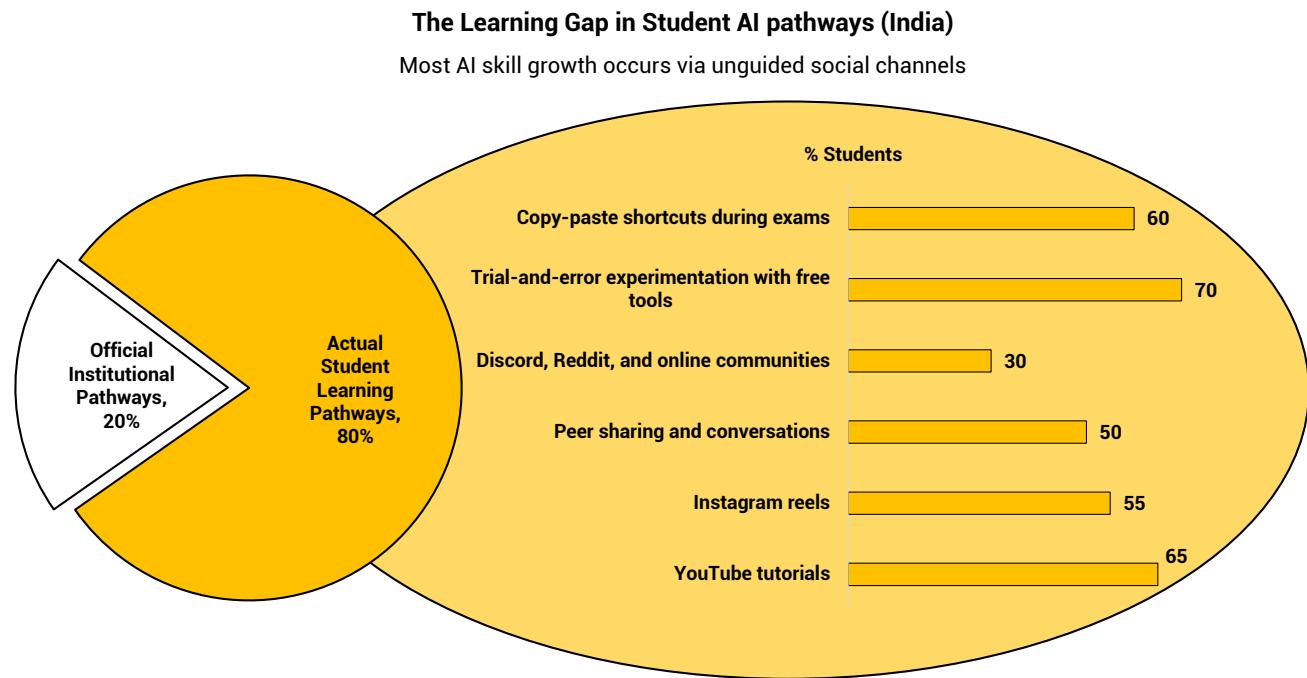


Figure 9: Diagram showing the stark contrast between official college AI learning pathways (20% engagement) and actual student learning through YouTube, social media, and peer sharing (80% engagement).⁴

Official Institutional Pathways (20% of student learning):

- College AI electives or modules
- Faculty guidance on projects
- Structured assignments
- Assessment and feedback

Actual Student Learning Pathways (80% of student learning):

- YouTube tutorials (65% of students)
- Instagram reels and TikTok (55% of students)
- Peer sharing and conversations (50% of students)
- Discord, Reddit, and online communities (30% of students)
- Trial-and-error experimentation with free tools (70% of students)
- Copy-paste shortcuts during exams (60% of students during crunch)

The pathways diverge dramatically in quality. Official pathways are designed to develop depth, critical thinking, and capability. Actual pathways are fragmented, trend-driven, and focused on shortcuts.

A student described her learning journey: *"I watched a YouTube video on how to use ChatGPT for essays. Then I saw an Instagram post about prompting. Then my friend told me how to use Perplexity with Airtel. I tried all of them, copied answers sometimes, and rewrote sometimes. But I never really learned how to evaluate if the output is correct. I just used whatever worked that day."*

This pattern has consequences:

1. **Shallow knowledge:** Students learn tools, not principles. They know "how to use ChatGPT" but not "when AI is appropriate" or "how to evaluate AI outputs."
2. **No evaluation skills:** Students accept AI outputs at face value. They don't systematically check correctness, relevance, or bias. When asked, "How do you know if that answer is right?" most respond with "I assume it is" or "It sounds reasonable."
3. **Underground usage continues:** Because faculty policies are inconsistent, students remain uncertain about what's allowed. Many hide their AI usage even when it would be legitimate.
4. **High anxiety:** Fear of plagiarism accusations, AI detection, or grade penalties drives stress. Students are using AI but feeling guilty about it.
5. **No mentorship or guidance:** Learning happens in isolation. There's no mentor helping a student think through whether a particular AI use is helping or hurting their learning.

What's missing: Institutions could bridge this gap by making unofficial pathways official. Instead of banning YouTube learning, teach students how to evaluate YouTube content. Instead of policing peer sharing, facilitate structured peer learning communities. Instead of hiding trial-and-error, create safe spaces for experimentation with AI.

3.6 The Curriculum Lag: Teaching Yesterday's Skills

Most institutional curricula have not been meaningfully redesigned for an AI era. The changes made are typically superficial: adding an "AI module" to existing courses, offering an AI elective, or integrating AI case studies into existing lessons.

Substantive redesign is rare. It would require:

- **Rethinking assessment:** From "can you memorise and regurgitate?" to "can you think critically and apply judgment?"
- **Redesigning assignments:** From "write an essay" to "use AI to draft an essay, evaluate the output, improve it, and reflect on your process".
- **Changing reading lists:** From "here's the textbook" to "here's how to find, synthesise, and critically evaluate information with AI".
- **Reframing projects:** From "solve this defined problem" to "identify your own problem, use AI to explore solutions, iterate, and learn".

This requires faculty to think differently about their subject. And most faculty haven't been trained to do this.

One statistics professor at a leading institution captured the shift: *"I used to teach students how to compute statistical analyses by hand. Now AI can do that instantly. So, I've shifted to teaching them how to interpret AI-generated analyses, how to evaluate whether AI chose the right statistical test, and how to catch mistakes or biases in AI's output. I'm teaching judgment, not computation."*

This is a profound pedagogical shift. But it is happening in isolation. Most institutions continue teaching computation, assuming AI isn't a factor.

The result: **graduates are trained on yesterday's skill frameworks**. They can compute, memorise, and follow procedures, essentially all the things that AI can now handle. They often cannot think critically, evaluate ambiguity, or apply judgment, which are the things that will matter in their careers going forward.

3.7 Assessment in Crisis: The Evaluation System Breaks Down

When AI can generate essays, write code, solve math problems, and create presentations, traditional assessment becomes unreliable. An essay assignment no longer evaluates writing ability; it evaluates willingness to follow rules about AI usage. An exam evaluates whether a student can generate outputs that look smart, not whether they can think.

This has created a cascading crisis:

Problem 1: Detection Tool Unreliability

AI-writing detection systems are still highly unreliable. They frequently mark human-written work as AI-generated, fail to catch a lot of AI-written text, and behave inconsistently across subjects and formats. As an example, using tools as a high-stakes gatekeeper that automatically penalises or fails students once a report crosses an arbitrary percentage threshold. This can be very hard to defend, both educationally and in terms of due process. Yet many institutions continue to rely on them this way instead of treating them as low-stakes, advisory signals that must always be interpreted by humans.

Problem 2: The Evaluation Mismatch

Faculty are trying to evaluate whether students generated outputs themselves, when the real question should be: "Can this student think, evaluate, iterate, and produce work of value with the tools available?"

A student might use ChatGPT to draft an essay, rewrite it substantially, fact-check the claims, refine the argument, and produce something genuinely good. This involves thinking, judgment, and capability. Yet traditional assessment penalises this. Conversely, a student might hand-write an essay that is mediocre but "original." Traditional assessment rewards this.

Problem 3: Misalignment with Employer Needs

Employers don't care whether a candidate writes code from scratch or with a copilot. They care whether the candidate can solve problems, evaluate trade-offs, and produce good outcomes. But universities are still evaluating "originality of output", while employers are evaluating "quality of thinking."

Problem 4: Student Confusion and Anxiety

Because assessment policies are unclear and inconsistent, students are anxious and confused. They want guidance on how much AI usage is acceptable. They get contradictory messages: "AI is the future, use it!" combined with "We'll penalise AI usage if we catch it." Students internalise this as "AI is good, but risky."

3.8 Leading-Edge Responses: A Model Emerges

A small number of institutions are moving beyond the crisis toward genuine innovation.

IIIT-Delhi's Prompt-Based Evaluation Model is the clearest example. The institution has redesigned core courses around a principle: **allow AI but require transparency and demonstrate thinking.**¹⁰

Students must submit:

1. The prompts they used (exact wording)
2. The AI outputs received
3. Their evaluation of those outputs
4. Changes they made and why
5. Their reasoning for the choices

Grading rubrics explicitly evaluate:

- **Prompt quality:** How well did the student structure their request to AI?
- **Critical evaluation:** Did they assess whether the output was correct, biased, or limited?
- **Judgment and refinement:** Where did they apply human thinking rather than accepting AI outputs?
- **Final output quality:** The actual work product, in the context of all the above

This shifts evaluation from "Did you generate this yourself?" to "Did you think about this?"

One institute's structured AI policy emphasises disclosure and reflection, not bans or silent use. AI usage is explicitly permitted; faculty design prompt guides so students learn to use AI as a thinking partner rather than a shortcut. Assignments increasingly ask questions such as: "How would you approach this using AI? How would you solve it without AI? In which steps are AI genuinely helpful, and where does it become a limitation?"

Result: Students are more honest about AI usage because there's no fear. Underground usage decreases. Genuine AI literacy develops.

These models provide a template that other institutions can adopt and improve upon.

3.9 Student Impact: What the Underground Reveals

Student interviews reveal the profound impact of institutional ambiguity:

From Lakshmi (Tier 3 engineering college):

"My faculty doesn't explicitly forbid AI, but we are all scared. I use Perplexity to understand concepts, but I rewrite everything in my own words so I don't get accused of plagiarism. I know I'm using AI, but I don't think about whether I'm learning from it or just cheating myself. I just want to pass and not get caught."

From Priya (Tier 1 college):

10 Jha, S. (2025, December 1). AI in exams: IIIT-Delhi pilots new model requiring students to submit prompts. The Times of India.

"I use Claude and Perplexity openly in my projects. My professors know and expect it. I document my thinking process, like what I asked AI, how I evaluated the response, and where I applied my own reasoning. I'm learning how to think with AI, not just getting answers. I feel confident about my capability."

The difference in learning outcomes is stark. One student is learning in hiding, with no feedback or mentorship. The other is learning openly, with institutional scaffolding.

3.10 The Faculty Barrier: Resistance and Readiness

Faculty resistance to AI-integrated teaching has several roots:

Fear of irrelevance: *"If students can generate content with AI, what is the point of my teaching?"*

Protective instinct about rigour: *"AI will make students lazy. They'll stop thinking."*

Lack of confidence: *"I don't understand AI well enough to teach it or evaluate its usage."*

Workload concerns: *"Redesigning all my assessments will take months."*

Administrative friction: *"Our university hasn't approved new tools. I can't assign work using paid AI platforms."*

These concerns are not irrational. They reflect real challenges. But they can be addressed with institutional support:

- **Faculty training programs** in AI tool use, AI-aware pedagogy, and assessment redesign
- **Peer learning communities** where faculty share what they are experimenting with
- **Explicit institutional permission** and budget for tool access (paid AI platforms for teaching)
- **Recognition and incentives** for faculty who redesign courses
- **Reduction in course load** during transition periods to allow curriculum redesign
- **Clear guidance** on what AI-integrated teaching should look like

Most institutions have barely done a few of these. The result: faculty remain stuck, and innovation happens at the margins.

3.11 The Tier 2/3 Challenge: Structural Disadvantage

Tier 2 and 3 institutions face compounding challenges:

1. **Affiliated college constraints:** Many are affiliated with universities, so they cannot freely redesign curricula. Changes must go through departmental committees and the university senate, processes that take 2-5 years.
2. **Faculty capability gaps:** Tier 2/3 colleges often have faculty who lack industry exposure or advanced degrees. Training them in AI pedagogy requires significant institutional investment.
3. **Limited infrastructure:** Shared computer labs with limited seats and internet bandwidth. No access to paid AI platforms for teaching.
4. **Student population:** More first-generation students, fewer resources at home. These students need more mentoring, not less. Yet institutions have fewer mentors.

5. **Reputation pressures:** Tier 2/3 colleges compete on placement outcomes. Quick wins (teach interview prep, provide placement coaching) are prioritised over curriculum redesign.

One Tier 2 college faculty member described the challenge: *"I want to integrate AI into our curriculum. But our faculty don't have the skills. Our students don't have laptops. Our infrastructure can't support it. And the university's approval process takes forever. So, we do what we can: offer AI electives, partner with EdTech for upskilling, and encourage students to use free tools on their own time. But it is not systemic."*

This is the typical Tier 2/3 story: sincere intent, structural constraints, piecemeal progress.

3.12 EdTech's Incomplete Promise: Filling and Widening Gaps

EdTech platforms and skilling providers have become de facto curriculum designers for AI. They are offering AI modules, running bootcamps, and partnering with colleges. In theory, this democratises access. In practice, it often widens gaps.

What's working:	What's not working:
<ul style="list-style-type: none"> • Making AI tools and frameworks accessible at scale • Providing structured, project-based learning (vs university's theory-heavy approach) • Offering job-linked programs with placement support • Reaching working professionals who can't do full-time re-education 	<ul style="list-style-type: none"> • Over-teaching generic AI content that becomes outdated every 6 months • Offering credentials without genuine capability building • Skipping foundational problem-solving in favour of tool tutorials • Concentrating offerings in tech/coding roles, ignoring domain–AI integration • Creating a two-tier education: expensive paid programs for those who can afford it, free but low-quality content for others

One EdTech leader was blunt: *"Most AI skilling today is cosmetic. We are teaching ChatGPT and Gemini usage because it is easy to package and market. Real transformation that entails teaching students to think about AI in the context of their domain, to evaluate outputs, and to build judgment is harder, slower, and less marketable. So, we don't do it."*

3.13 What's Missing: The Elements Needed

Across all institutional conversations, a consistent set of needs emerged:

1. **Institutional Clarity on What Ai-Integrated Learning Should Look Like** Colleges don't know whether the goal is "teach students to use AI tools", or "teach students to think critically about AI's role in their domain", or "teach students to build AI systems." These lead to very different curricula.
2. **Faculty Capability Building at Scale** Not one-off workshops. Sustained, peer-led, practice-heavy training. Faculty learning communities where professors experiment together, share what works, and iterate.

3. **Assessment Frameworks Aligned With Ai-Era Skills** Clear rubrics and methods for evaluating thinking, judgment, and evaluation capability and not just output quality. Models like IIIT-Delhi's prompt-based evaluation need to spread and be adapted.
4. **Clear Policies on Responsible AI Use** Not "ban AI" or "police AI," but "here's when and how to use AI as a learning tool; here's when it crosses into shortcuts; here's how to think about it." Transparency and reflection, not secrecy and punishment.
5. **Access to Mentorship** Most importantly, students need mentors (faculty or industry) who can guide them through real problems, help them think about when AI is appropriate, and provide feedback on their learning. This cannot be scaled through platforms or courses. It requires human relationships.
6. **Curricular Time and Mental Bandwidth** Faculty need explicit permission and time to redesign courses. Not "redesign in your spare time," but "here is release time; here is a process; here is institutional support."

3.14 The 3-Year Outlook: Where This Is Heading

Based on institutional trends and leadership signals, here's what is likely by 2027-2028:

Institutional Responses:

- Tier 1 institutions will have integrated AI across multiple courses and refined assessment methods
- Tier 2 institutions will still be experimenting with scattered AI modules
- Tier 3 institutions will have minimal change in most programs

Assessment Evolution:

- AI detection as a primary enforcement tool will fade (tools are unreliable; approach is pedagogically unsound)
- Process-based evaluation (showing thinking) will become more common in better institutions
- Portfolio-based assessment will expand, particularly in tech and applied fields

Curriculum:

- More domain-AI integration (rather than generic "AI courses")
- Critical thinking and evaluation skills will be more explicitly taught
- Foundational skills (writing, problem-solving) will remain important

Faculty:

- Minimal change in most institutions without mandates
- Leading institutions will have redesigned key courses
- Peer learning communities will have emerged in pockets
- Faculty anxiety will persist; training will remain insufficient

Student Outcomes:

- Tier 1 graduates will have sophisticated AI literacy and capability

- Tier 2 graduates will have tool familiarity but limited depth
- Tier 3 graduates will have self-taught, informal AI exposure but no institutional scaffolding
- The capability gap between tiers will widen unless targeted interventions happen

3.15 What This Means for Employability

The institutional response landscape directly impacts graduate employability. Employers hiring from Tier 1 institutions with AI-integrated pedagogy will find graduates who can think critically about AI. Employers hiring from Tier 2/3 institutions will find graduates with tool awareness but limited depth.

This creates a reinforcing cycle: Tier 1 graduates get premium roles at better companies. Tier 2/3 graduates compete for generic roles. The gap widens.

Breaking this cycle requires institutional change, but not in all colleges at once, but in enough of them that the models spread and become normalised. **The section that follows examines what policy and institutional levers could accelerate this change.**

4

POLICY PATHWAYS & SUGGESTED RECOMMENDATIONS

4.1 The Choice Before India: Three Futures

India's higher education system stands at a fork in the road. The choices made in the next 18 months will determine whether the nation develops a generation of AI-ready graduates or deepens existing inequalities.

Future 1: Drift - Institutions continue as they are. Some elite colleges integrate AI. Tier 2/3 institutions ban or police AI usage. Students continue learning informally through YouTube and social media. The capability gap between tiers widens. Graduates from Tier 1 colleges compete globally; graduates from Tier 2/3 colleges struggle domestically.

Future 2: Fragmentation - Different institutional responses proliferate without coordination. Some adopt an AI-first pedagogy. Others ban AI. Some colleges partner with EdTech; others go it alone. The result is incoherent, with no national baseline for AI-era learning.

Future 3: Coordinated Transformation - Institutions receive clear policy guidance, resource support, and time to redesign. Faculty is systematically trained. Assessments are redesigned. Students learn with institutional scaffolding. Tier 2/3 colleges are supported to close gaps. Graduates across all institutions develop genuine AI capability.

This section suggests pathways toward Future 3.

4.2 Core Principles for Policy and Institutional Action

Core principles that may guide potential interventions:

1. Enable, Don't Police

AI is not going away. Students will use it. The question is whether they learn to use it wisely within institutions or hide usage and develop poor habits. Policies should enable responsible usage with transparency and reflection, not police and punish.

Clear institutional policies are the first step. Rather than "AI is prohibited," policies should read: "AI is allowed for learning, with these expectations: (1) Disclose your AI usage, (2) Show your thinking process, (3) Evaluate AI outputs critically." This shifts the focus from enforcement to learning.

2. Address Structural Barriers, Not Just Tool Access

Providing free ChatGPT access to students without addressing mentorship, curriculum redesign, and faculty readiness is like handing a powerful tool to someone without instructions. Structural barriers include:

- Lack of faculty capable of teaching AI-integrated courses
- Assessment systems that reward memorisation, not judgment
- Absence of real, contextualised problems for students to work on
- Geographic isolation and limited industry exposure

All of these must be addressed alongside tool access.

3. Differentiate by Tier (One-Size-Fits-All Fails)

Tier 1 institutions need depth and sophistication. Tier 2 institutions need structure and faculty

development. Tier 3 institutions need foundational access and awareness. A single national policy cannot serve all three. Support mechanisms must be tiered.

4. Invest in Faculty First

The bottleneck is faculty readiness, not student access. Every policy recommendation should include faculty the capability building. This is not one-off workshops but sustained, peer-led, practice-focused development.

5. Measure Impact on Employability, Not Adoption

Success is not "percentage of students exposed to AI" or "number of AI courses offered." Success is "percentage of graduates with capability to use AI effectively in their roles" and "employer satisfaction with AI-augmented capabilities." Policies should tie incentives to outcomes.

4.3 Suggested Recommendation Category 1: Immediate (Next 6 Months)

These interventions can be implemented quickly with modest resources. They can help create a foundation and establish momentum for longer-term work.

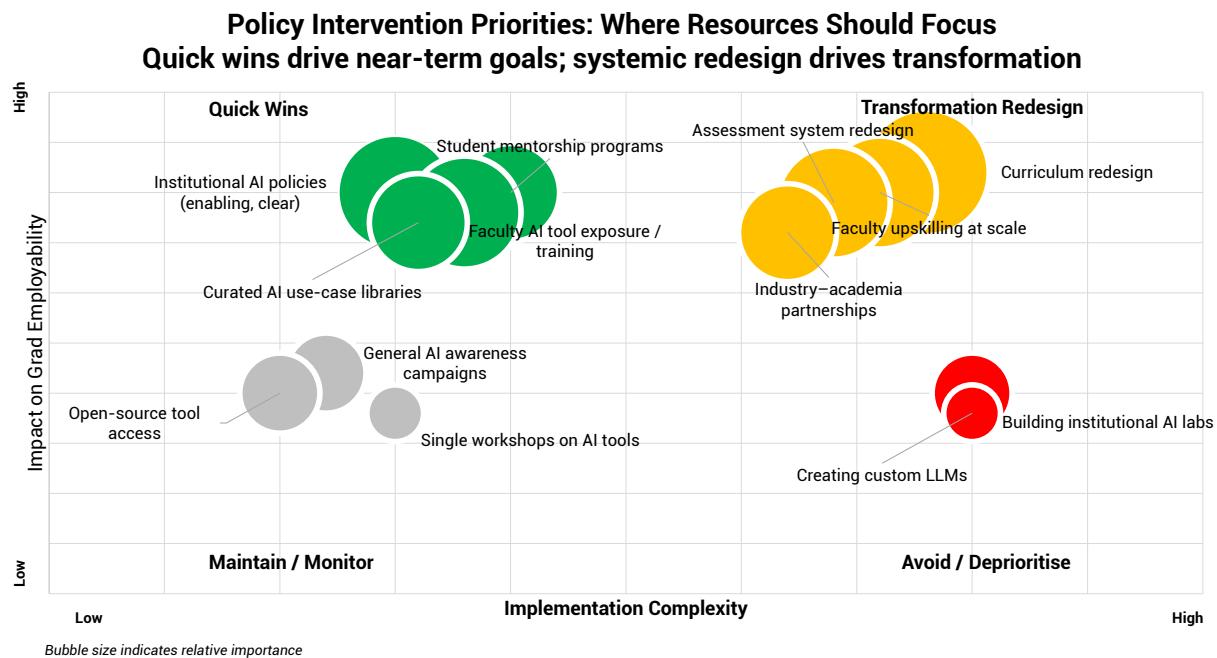


Figure 10: Illustrative - Impact-complexity matrix for AI policy interventions, showing clear institutional policies and mentorship as high-impact, low-complexity quick wins, while curriculum and assessment redesign are high-impact but resource-intensive strategic priorities.

4.3.1 National and Institutional AI Usage Policies

What is needed: Clear, written policies at national, university, and institutional levels that replace the current ambiguity.

National Level (UGC/Ministry of Education):

- **Publish guidance document:** "AI in Higher Education: Assessment, Usage, and Responsible Integration"

- Acknowledge that students will use AI; provide a framework for institutions to allow it responsibly
- Distinguish between legitimate AI use (learning support, exploration, evaluation practice) and problematic use (complete outsourcing of thinking)
- Recommend assessment redesign principles rather than detection-based enforcement
- **Make clear:** institutions banning AI are inhibiting learning, not protecting rigour

Institutional Level:

Each college should adopt an explicit policy such as:

AI Usage Policy (Sample)

Allowed: Using AI tools to clarify concepts, explore ideas, draft documents, check work, and learn from varied explanations.

Expected: Disclosure of AI usage in assignments (e.g., "I used ChatGPT to understand this concept" or "I used Cursor to write this code"). Show your thinking process. Evaluate AI outputs for accuracy and relevance. Apply your own judgment.

Not Allowed: Submitting AI-generated work as your own without disclosure. Relying entirely on AI without critical evaluation. Using AI to avoid learning.

Assessment Principle: We evaluate your thinking, judgment, and ability to work with AI and not whether you generated words yourself.

Suggested Ownership: Ministry of Education / regulatory bodies (national), university/board (for affiliated colleges), institutional leadership (implementation)

Likely Impact: Reduces student anxiety. Ends underground usage. Creates a foundation for redesign.

4.3.2 Faculty Access to AI Tools (Subscriptions)

What is needed: Every faculty member should have access to at least one high-quality generative AI tool and, where relevant, an AI-augmented coding / IDE environment, so they can experiment, prepare materials, and design AI-aware assignments from direct experience.

How Institutions Can Approach This:

- Allocate a modest annual AI tools budget per faculty member (for example, in the range of ₹5,000–10,000) that departments can use flexibly for subscriptions or credits.
- Prioritise multi-purpose tools faculty can use for lesson planning, content generation, feedback on assignments, and research support, plus developer-friendly tools for those teaching programming, data, and engineering.

Where budgets are tight, institutions can explore:

- Campus-wide or departmental licences and volume discounts.
- Consortia or state-level negotiations so that multiple institutions share a contract.
- Time-bound pilots before committing to longer contracts.

Examples of Tools (Illustrative, Not Exhaustive)

- General-purpose chat/writing/reasoning tools: institution or personal licences for systems like ChatGPT Plus, Claude, Gemini, or Perplexity.
- AI augmented IDEs and coding platforms: tools in the family of Cursor, Replit, Bolt.new, or GitHub Copilot can help faculty who teach programming or data science design realistic, AI-aware coding tasks.

- Model building and experimentation environments: browser-based studios and notebooks (for example, products like Google AI Studio, Azure AI Studio, etc.) can support faculty who want to show students basic prompting, fine-tuning, or API workflows.

These are examples only; the right stack will vary by discipline, budget, and institutional policy. The key recommendation is that every teaching faculty member has at least one reliable AI environment they can use regularly, rather than prescribing any single product.

Academic Pricing and Access

- Most major AI providers now publish education or research programmes, for example: Discounted or free credits for verified university domains, Special pricing for labs, classrooms, and hackathons, Grants or fellowship schemes that bundle tool access with training.
- Institutions should systematically scan these schemes, nominate a central point of contact (often IT or a teaching and learning centre), and consolidate applications so faculty are not doing this individually. This often brings costs down substantially compared to ad hoc personal subscriptions.

Suggested Ownership and Likely Impact Ownership

- Academic leadership with IT/finance and the teaching and learning unit jointly defining an approved tool list, support, and data governance guardrails.
- **Impact:** Faculty gain hands-on fluency, can design far more realistic assignments, and are more confident discussing both the strengths and limits of AI with students.

4.3.3 Student Mentorship Program Design and Launch

What is needed: Structured mentorship that connects students, especially those without strong professional networks, to people who use AI in real work contexts. Mentors help students understand how AI is actually used in roles, what skills matter, and how to navigate careers in an AI-augmented economy.

Mentor Sources:

- **Alumni and industry partners:** Graduates and local professionals who can commit to a semester of light-touch mentoring (for example, 1–2 hours every two weeks).
- **Professors of Practice:** Under existing higher education norms, Professors of Practice appointed from industry can be treated as anchor mentors, with a portion of their workload explicitly earmarked for student mentoring, project clinics, and small group guidance rather than only classroom teaching.
- **National / state level mentor pools:** Institutions can also draw on a higher education equivalent of initiatives like the "Mentor of Change" model in school innovation missions, where professionals register on a central platform, receive basic orientation, and are matched to colleges for structured, time-bound engagements.

Scale and Participation (Illustrative):

- Aim for **at least 5–10% of students** to be in a formal mentorship relationship in the first phase, with a mix of one-to-one and small-group formats.
- Use a **flexible mentor-to-student ratio**, such as 25–100 mentors per 1,000 students, depending on local capacity, and then adjust upward as recruitment pipelines, platforms, and coordination processes mature.

The medium-term goal can be to expand coverage over time, with particular emphasis on first-generation learners and campuses that currently have limited exposure to industry.

Suggested Program Design:

- **Duration:** Typically aligned with an academic semester (4–6 months), with clear expectations on contact hours and activities.
- **Focus of sessions:** How AI is used in the mentor's role; what tools and workflows matter; how to evaluate AI outputs; and how students can build portfolios that demonstrate real capability, not just course completion.
- **Support and coordination:** A central cell (for example, the placement office or a teaching-and-learning unit) could handle mentor onboarding, matching, basic orientation, and feedback loops, so that individual faculty and students are not left to coordinate informally.

Likely Impact: When designed this way, mentorship becomes a systematic way to:

- Translate abstract AI skills into concrete career pathways.
- Reduce anxiety and misinformation about AI's impact on jobs.
- Give students, especially in resource-constrained institutions, regular, guided exposure to how professionals are actually working with AI in their fields.

4.3.4 Assessment Redesign Pilot (2-3 Courses per Institution)

What is needed: Select 2-3 courses and redesign them to evaluate thinking and problem-solving, not just output generation.

Redesign Principles:

1. Prompt Transparency: Students must show the prompts used and explain why?
2. Process Documentation: How did you evaluate the AI output? What did you change? Why?
3. Rubric Redesign: Add criteria for "critical evaluation of AI output" and "appropriate use of AI"
4. Assignment Redesign: Structure assignments so AI handles routine; students handle judgment

Example Redesign (Statistics Course):

Old assignment: "Using this dataset, compute the mean, median, and run a regression. Show your work."

New assignment: "Using this dataset, use an AI tool to generate an initial statistical analysis. Show the prompts you used. Evaluate: Is the AI analysis correct? Are the assumptions valid? What is the AI missing? Run your own check on the most important finding. Document your reasoning."

Suggested Ownership: Individual faculty and department heads

Likely Impact: Demonstrates that the redesigned assessment is feasible. Faculty confidence increases. Student capability visibly improves.

4.3.5 Communication Campaign: "What Matters Now"

What is needed: Consistent messaging to students about what's actually valued in hiring and careers.

Key Messages:

- "AI skills matter, but critical thinking matters more."

- "We are not banning AI. We are teaching you to use it wisely."
- "Your portfolio matters more than your resume."
- "We want to see your thinking process, not just your output."
- "Curiosity and adaptability are your competitive advantages."

Channels:

- Institutional website, Instagram, LinkedIn
- Faculty communications in class
- Placement cell messaging
- Mentorship orientation
- Alumni spotlights (showing how AI is used in real roles)

Suggested Ownership: Institutional communications and placement cell

Likely Impact: Shifts student mindset from "hide AI usage" to "use AI wisely." Reduces anxiety. Increases engagement.

4.4 Recommendation Category 2: Mid-Term (6-12 Months)

These initiatives require more sustained effort and resource allocation. They build on the foundation created in the first 6 months.

4.4.1 Sustained Faculty Reskilling Programs

What is needed: Structured, peer-led training for faculty in AI tool use, AI-aware pedagogy, and assessment redesign.

Model:

- **Cohort size:** 30-50 faculty per cohort
- **Duration:** 8-12 weeks (2-3 sessions per week)
- **Format:** 50% online (self-paced learning and discussions), 50% in-person (hands-on practice, peer learning)
- **Content:**
 - AI tool fundamentals (ChatGPT, Claude, domain-specific tools)
 - Hands-on experimentation (try tools on your course content)
 - Pedagogy redesign (how to structure assignments for thinking)
 - Assessment methods (rubrics, process evaluation)
 - Ethical considerations and responsible AI
 - Peer teaching (each participant designs one redesigned assignment)

Facilitators:

- **Tier 1:** Internal faculty experts and external consultants
- **Tier 2/3:** Master trainers from Tier 1 institutions and online facilitators

Rollout:

- **Tier 1:** Cohort 1 (months 6-9), targeting 30% of faculty
- **Tier 2:** Cohort 1 (months 8-11), targeting 20% of faculty (with government subsidy for trainer costs)
- **Tier 3:** Cohort 1 (months 10-12 or early next year), targeting 15% of faculty (full government support)

Suggested Ownership: Institutional faculty development offices, with government and EdTech partnership support

Likely Impact: Faculty move from "I don't understand AI" to "I can teach with AI." 30-50% of faculty are significantly more capable within 12 months.

4.4.2 Curriculum Mapping and Domain–AI Integration Planning

What is needed: For each program (engineering, commerce, humanities, etc.), map where AI naturally integrates and design learning outcomes that blend domain and AI.

Process:

1. **Curriculum audit:** What are the core learning outcomes for this program?
2. **AI application mapping:** Where does AI enhance or transform each outcome?
3. **Redesign planning:** How do we redesign courses to teach domain–AI integration?
4. **Sequencing:** Which courses should be redesigned first?

Illustrative Examples:

Program	Core Outcomes	AI Integration	Redesigned Focus
Finance	Analyse financial data, forecast trends, and assess risk	AI does routine analysis and forecasting; humans do contextual judgment	"Use AI to analyse investment risk, evaluate assumptions, and apply business judgment"
Engineering	Design systems, solve problems, build prototypes	AI handles research and initial design; humans do trade-off analysis and validation	"Use AI to research design solutions, critique proposals, and justify final design"
Data Science	Build models, extract insights	AI handles model building; humans do interpretation and business application	"Use AI to build initial models, interpret results in business context, validate assumptions"
Humanities	Interpret texts, construct arguments, and understand context	AI assists with research and synthesis; humans do critical analysis	"Use AI to gather sources and context, analyse critically, construct original argument"

Suggested Ownership: Department heads and curriculum committees, with institutional coordination

Likely Impact: By the end of month 12, all programs have a clear plan for AI integration. Redesign work can begin immediately after.

4.4.3 Student Portfolio Guidance System

What is needed: Structured guidance helping students build portfolios that demonstrate capability to employers.

Components:

1. **Portfolio platform:** Centralised space for students to document projects (GitHub, LinkedIn, or institutional)
2. **Guidelines:** What should be in a portfolio? How to document your thinking?
3. **Feedback mechanism:** Mentors/faculty review portfolios and give guidance
4. **Employer signalling:** Clear communication to employers about what portfolio evidence means

Portfolio requirements by year:

- **Year 1:** 1-2 small projects documenting the problem-solving process
- **Year 2:** 2-3 medium projects showing iteration and critical thinking
- **Year 3:** 1-2 significant projects with documented reasoning and AI integration where relevant

Suggested Ownership: Placement cell and academic advisors

Likely Impact: By the end of the year, 50%+ of students have documented portfolio evidence. Employers can evaluate capability, not credentials.

4.4.4 Industry Problem Integration (Pilot Projects)

What is needed: Real, contextualised problems from industry embedded into courses.

Model:

- **Partner with 2-3 local companies** per Tier 1 institution, 1-2 per Tier 2, and at least 1 per Tier 3
- **Scope:** Problem should be real but not critical (companies willing to share)
- **Integration:** 2-4-week problem embedded in a course
- **Student work:** Teams of 3-4 work on a problem, with mentorship from an industry sponsor
- **Outcome:** Company provides feedback; students get real-world exposure

Illustrative Examples:

- **Logistics company:** "Optimise delivery routes in our region using AI-assisted analysis"
- **Healthcare startup:** "Build a dataset for patient feedback analysis using AI"
- **Fintech:** "Assess credit risk for a customer segment using AI tools"
- **EdTech:** "Design a learning workflow using AI for a specific skill"

Suggested Ownership: Industry partnerships and faculty

Likely Impact: Students work on real problems. See immediate application. Employers identify talent early.

4.5 Suggested Recommendation Category 3: Long-Term (12-18 Months)

These suggested initiatives can transform the system by building on the work of the first 12 months.

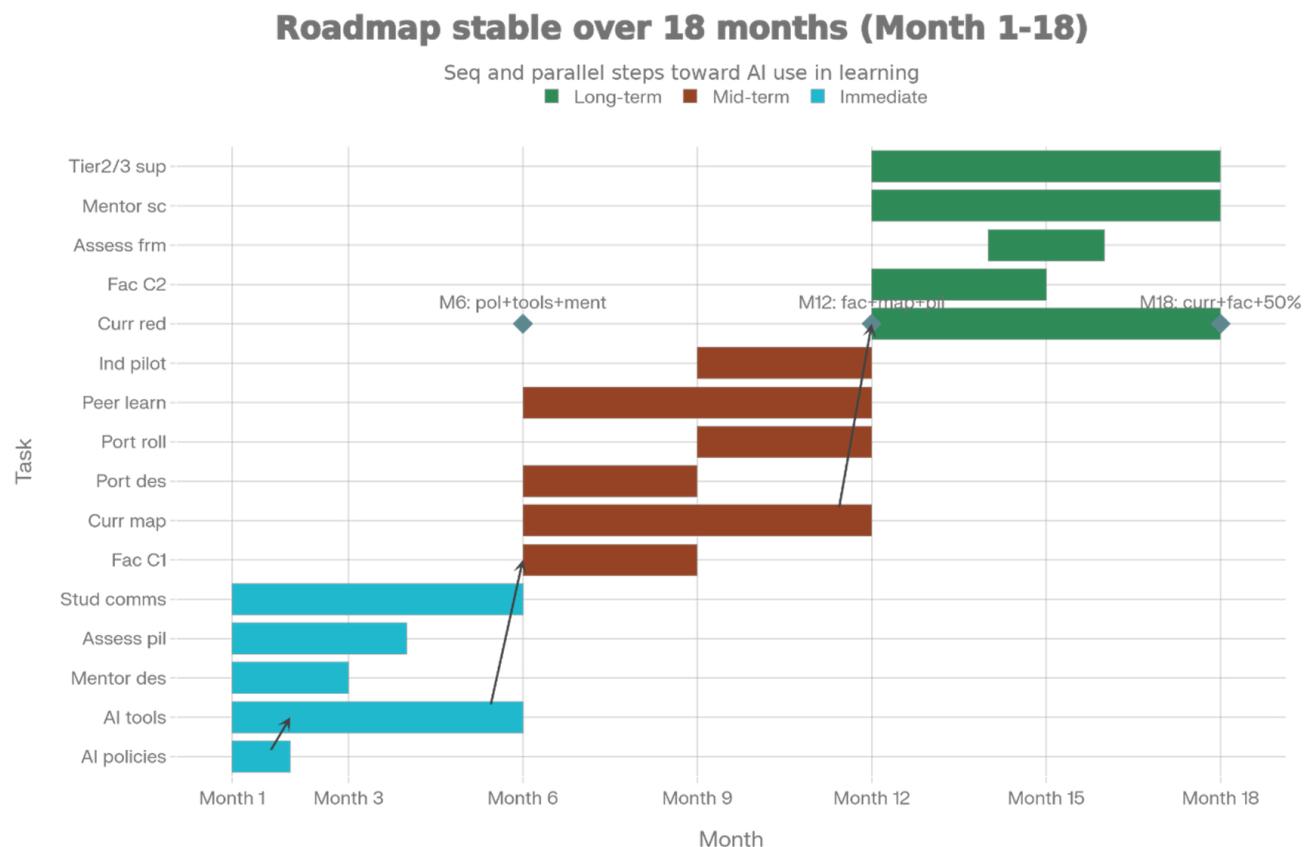


Figure 11: 18-month Gantt timeline showing phased implementation of AI-integrated learning across policy, faculty training, curriculum, and assessment redesign, with quick wins in months 1-6 enabling longer-term transformation.

4.5.1 Full Curriculum Redesign (Phase 1 Complete for Core Programs)

What is needed: By month 18, all core courses in core programs have been redesigned to integrate AI meaningfully.

Phase 1 Focus:

- Core courses across all undergraduate and postgraduate programs
- Affects 60-80% of student learning hours

Redesign Checklist:

- Learning outcomes include "use AI appropriately in this domain"
- At least 40% of assignments are redesigned for the AI era
- Assessments evaluate thinking, not just output
- Faculty trained in new pedagogy
- Students have clear guidance on AI usage expectations

Suggested Ownership: Faculty, department heads, and curriculum committees

Likely Impact: By the end of 18 months, all students experience AI-integrated learning. The capability baseline improves.

4.5.2 Assessment Framework Implementation at Scale

What is needed: Shift from AI detection to evaluation of thinking across all courses.

New Assessment Framework:

1. **Process evaluation:** Show your thinking. Document your prompts, your evaluation of outputs, and your reasoning.
2. **Rubric redesign:** All rubrics include criteria for "critical evaluation" and "appropriate AI usage"
3. **Portfolio integration:** Final assessment includes portfolio evidence
4. **No AI detection tools:** Remove AI detection flags (keep plagiarism detection, drop AI detection)

Implementation:

- **Months 12-14:** Train all faculty on new rubrics and evaluation approaches
- **Months 14-16:** Implement across all courses
- **Months 16-18:** Collect data on impact (student outcomes, faculty comfort, employer feedback)

Suggested Ownership: Faculty and assessment committees

Likely Impact: Assessment becomes reliable again. Evaluation focuses on what matters (thinking), not compliance (tool usage).

4.5.3 Mentorship at Scale

What is needed: Scale the mentorship program from 5-10% of students (phase 1) to 30-50% of students.

Expansion:

- **Recruitment:** Industry partnerships, alumni networks, and volunteers
- **Training:** Mentors get brief training on how to guide AI-era learning (not just career advice)
- **Platform:** Scale mentorship platform to support 100s of relationships simultaneously
- **Incentives:** Recognise mentors (alumni spotlight, corporate partnerships)

Suggested Ownership: Career/placement office and alumni association

Likely Impact: By month 18, 30-50% of students have industry mentorship. Dramatically increases real-world exposure and guidance.

4.5.4 Tier 2/3 Institutional Support Programs

What is needed: Targeted programs to help Tier 2/3 institutions catch up.

For Tier 2 Institutions:

- **Master trainer deployment:** Tier 1 faculty spends 2-3 months at Tier 2 institution training faculty
- **Curriculum partnership:** Co-develop curriculum with Tier 1 institution

- **Student exchange:** Some Tier 2 students work on projects with Tier 1 mentorship
- **Shared resources:** Access to GPU labs, datasets, and industry partnerships

For Tier 3 Institutions:

- **Device subsidy:** Government funding for subsidised laptops (~₹20,000 per device)
- **Community labs:** Regional hubs with GPU clusters (state/TBI supported)
- **Indian Language support:** AI interfaces in Hindi and other Indian languages
- **Master trainer and on-campus coordinator:** Someone physically present, helping with implementation
- **Simplified curriculum:** Focus on foundational AI literacy rather than advanced topics

Suggested Ownership: Government (Ministry of Education / regulatory bodies), state education departments, central agencies

Likely Impact: By month 18, Tier 2/3 institutions have closed 30-40% of the AI readiness gap.

4.6 Suggested Recommendation Category 4: Governance and Accountability

4.6.1 Establish National-level Coordination Mechanisms

What is needed: A coordination mechanism to align national and state-level imperatives with attention to local needs.

Composition:

- Ministry of Education / regulatory bodies representatives
- State education department representatives
- Faculty leaders (Tier 1, 2, 3 institutions)
- Experts in AI / emerging technologies
- Employer representatives
- Student representatives
- EdTech sector representatives

Mandate:

- Monitor implementation progress
- Identify and solve coordination barriers
- Collect data on outcomes
- Adjust recommendations based on learning

Frequency: Monthly meetings, quarterly reports

Suggested Ownership: Ministry of Education / regulatory bodies

4.6.2 Outcome-Based Accountability

What is needed: Clear metrics tied to institutional funding/recognition.

Metrics (by 2028):

Metric	Tier 1 Target	Tier 2 Target	Tier 3 Target
% of core courses with AI-integrated assessments	80%	50%	30%
% of faculty trained in AI-aware pedagogy	70%	50%	30%
% of students with documented portfolio projects	70%	40%	20%
% of faculty with tool access (subscriptions)	100%	80%	50%
Student self-reported AI capability (survey)	7/10	5/10	5/10
Employer satisfaction with AI capability	7/10	5/10	5/10

Incentives:

- Tier 1 institutions meeting 80%+ targets: Recognition and modest funding boost
- Tier 2 institutions meeting 50%+ targets: Government support for next phase
- Tier 3 institutions meeting 20%+ targets: Continued support and progression support

Suggested Ownership: Ministry of Education / regulatory bodies and state education departments

4.6.3 Research and Evidence Generation

What is needed: Rigorous evaluation of what works.

Studies to commission:

- **Impact evaluation:** Do redesigned courses improve learning outcomes?
- **Employer perception study:** Do graduates from AI-integrated programs perform better?
- **Equity study:** Are students in Tier 2/3 institutions closing gaps?
- **Faculty adoption study:** What enables vs blocks faculty change?

Suggested Ownership: Ministry of Education / regulatory bodies and research institutions

Likely Impact: Data-driven adjustments to policy; credibility with stakeholders.

4.7 Closing the Tier Gap: Differentiated Support

Tier AI support gap narrows (2025-2027)

Depth for Tier 1, Structure for Tier 2 and Access for Tier 3 show different paths

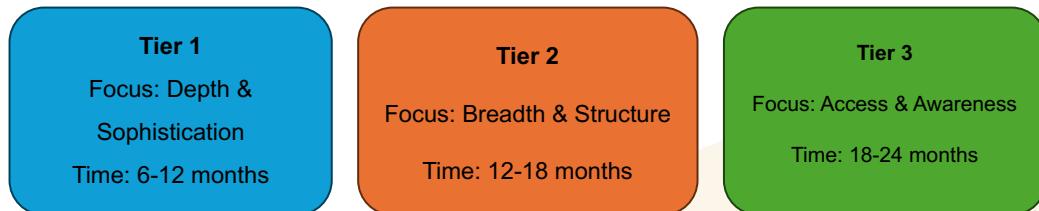


Figure 12: Three-column differentiated support framework showing tailored AI integration strategies for elite, mid-tier, and lower-tier institutions, with distinct interventions, timelines, and success metrics for each segment.

Across all recommendations, a critical principle emerges: **one-size-fits-all policies do not work in a stratified system.** Tier 1, Tier 2, and Tier 3 institutions need different support.

The Non-Negotiables

Across all recommendations, these elements cannot be compromised:

- Faculty capability building is non-negotiable:** Without it, everything else fails.
- Tier 3 institutional support is non-negotiable:** If only Tier 1 institutions benefit, inequality widens dramatically.
- Student agency and choice are non-negotiable:** Top-down mandates on how to use AI will be resisted. Institutions must involve students.
- Employer engagement is non-negotiable:** The goal is employability. Design with employer input throughout.
- Measurable outcomes are non-negotiable:** Policy must be tied to data. What's not measured isn't managed.

What Success could look Like by 2028

If these suggested recommendations are implemented:

For Students:	For Faculty:
<ul style="list-style-type: none"> Clear policies and no fear of punishment for legitimate AI usage Access to mentors who guide wise AI usage Curriculum that teaches both domains and AI integration Assessment that evaluates thinking, not tool usage Portfolio guidance that helps them demonstrate capability Realistic career clarity tied to the actual job market 	<ul style="list-style-type: none"> Tool access to experiment and learn Peer learning communities for support and knowledge sharing Curriculum redesign guidance and time Clear pedagogy frameworks for AI-era teaching Professional recognition for innovation
For Institutions:	For Employers
<ul style="list-style-type: none"> Coherent AI strategy aligned with employer needs Measurable outcomes showing improvement Competitive advantage, attracting students and employers Reputation for preparing graduates for AI-era roles 	<ul style="list-style-type: none"> Graduates with genuine AI capability, not just exposure Diverse talent pipeline from Tier 1, 2, and 3 institutions Clear portfolio evidence of thinking and judgment Reduced need for reskilling on entry-level hires

For India:

- More equitable distribution of AI capability across geography and institutions
- Graduates that are globally competitive in AI-augmented roles
- Reduced technology-induced job anxiety
- National advantage in human-centric AI roles (where judgment and context matter)
- Stronger talent pipeline for AI innovation in every sector

The Next 18 Months Matter

The decisions institutions, government, and employers make in the next 6-18 months will determine whether India's higher education system adapts successfully to the AI era or falls further behind.

The pathway exists. The recommendations are actionable. The resources are available at scale. What is required is coordination, commitment, and clarity.

The final section synthesises the various learnings and findings to offer a possible roadmap for next steps.



5

CONCLUSION: FROM AWARENESS TO ACTION

5.1 The Moment We are In

India stands at an inflection point in higher education. The arrival of generative AI has exposed a fundamental misalignment between how institutions teach and what graduates actually need to thrive in the workforce. This is not a technology crisis. It is an institutional readiness crisis.

The evidence is clear: Students are using AI extensively, almost 53.5% daily. Yet institutional responses range from denial (banning AI) to confusion (inconsistent policies) to scattered experimentation (a few courses in a few colleges). The result is that students are learning AI in fragmented, unguided, informal channels while institutions pretend the technology doesn't exist in official contexts.

The AI Revolution in Higher Education: Eight Critical Findings

The Usage-Capability Gap 53.5% use AI daily Only 10-15% at Level 3+ Usage is widespread; deep capability is rare.	The Tier Divide Tier 1: 70.2 AI Index Tier 3: 26.3 AI Index 2.7x difference in AI readiness	The Faculty Crisis 65 pt capability gap 25% current v/s 90% needed Faculty readiness is bottleneck	The Policy Paradox Most ban or police AI A few structured integration Policing drive underground usage
The Underground Learning Path 80% learn outside campus YouTube 65%, IG 55%, peers 50% Institutions miss where learning actually happens.	The Employability Signal Critical thinking, adaptability AI fluency is 2 nd tier Capability over credentials matter	The Anxiety-Risk Mismatch 45.9% fear reduced readiness Risk varies by field Misinformation drives career choices	The Coordination Imperative Policies, faculty reskilling, student mentorship, redesign assessment, device access Transformation change is achievable

These findings point toward a clear policy imperative: coordinate, fund, and support institutional transformation in the next 18 months.

Figure 13: Eight key research findings summarised as cards showing the usage-capability gap, institutional tier divide, faculty readiness crisis, policy paradox, underground learning patterns, employability signals, anxiety-risk mismatch, and the coordination imp.

This report documents what is really happening. Through conversations with 40+ industry leaders, faculty members, EdTech providers, and in-depth interviews with 170+ students across India, we have mapped the landscape on how higher education could respond, and also the reality of how it is currently responding.

The findings are simultaneously encouraging and alarming:

Encouraging: Students are capable, curious, and adaptable. Many are building genuine AI capability despite institutional barriers. Employers recognise that AI-augmented skills matter and are beginning to signal what they need. A small number of institutions have designed pedagogies that develop real AI literacy alongside domain expertise. These models prove that transformation is possible.

Alarming: 65% of institutions are still banning or policing AI usage. Faculty capability gaps are 65 percentage points wide across critical dimensions. The gap between institutional readiness at Tier 1

and Tier 3 colleges is substantial. Students are developing evaluation skills through trial-and-error, not mentorship. Geographic and socioeconomic inequality in AI access and opportunity is widening. And most troublingly, institutions are making decisions in the next 6-12 months that will lock in these inequalities for a generation.

5.2 What This Report Reveals

The Core Finding: Capability ≠ Usage

The first and most critical insight is that high usage does not equal high capability. A student who uses ChatGPT daily to summarise readings and rewrite assignments is not developing the same capability as a student who uses AI iteratively to solve novel problems, evaluates outputs critically, and understands when AI is appropriate versus when human judgment is essential.

This distinction matters because employers are not hiring for "people who have used AI." They are hiring for "people who can think with AI." These are different skill sets. The former requires tool familiarity. The latter requires judgment, critical thinking, and problem-solving capability.

The consequence: Students with high exposure but low depth will graduate with resume credentials ("I know ChatGPT") but lack portfolio evidence of capability. They will compete against graduates from structured programs with genuine depth. In hiring processes that use portfolio-based assessment (increasingly common), the gap will be decisive.

The Structural Reality: Tier Divide Is Real

India's higher education system is deeply stratified. Tier 1 institutions (IITs, top private colleges) have infrastructure, faculty expertise, industry connections, and resources. Tier 2 institutions (state universities, regional private colleges) have partial resources and mixed capabilities. Tier 3 institutions (rural colleges, small towns) have severe constraints across all dimensions.

This stratification directly maps to AI readiness:

- **Tier 1:** Some courses integrated with AI; faculty have experimented; students have mentorship access; device access is universal
- **Tier 2:** Scattered AI modules; faculty inconsistently prepared; limited mentorship; device access is mixed
- **Tier 3:** Minimal AI integration; faculty lack capacity; no mentorship; device access is a barrier

If nothing changes, this gap will widen. Tier 1 graduates will have genuine AI literacy and global competitiveness. Tier 2 graduates will have a partial understanding. Tier 3 graduates will have theoretical knowledge and self-taught informal skills with no institutional support. The result: a two-tier job market where geography and institutional affiliation determine opportunity more than capability.

This is unacceptable and preventable. But only with targeted, resource-backed policy interventions.

The Hidden Reality: The Underground Economy of Learning

One of the most significant findings is that **80% of student AI learning happens outside formal institutions.** Students are learning from YouTube tutorials, Instagram reels, peer sharing, Discord communities, and trial-and-error experimentation. They are not learning from faculty, structured courses, or mentored problem-solving.

This underground learning has consequences:

1. **No quality assurance:** YouTube tutorials vary from excellent to deeply misleading. Students have no filter.
2. **No evaluation training:** Self-taught learning rarely includes the critical meta-skill of evaluating AI outputs.
3. **No mentorship:** Learning happens in isolation without guidance on where AI is appropriate.
4. **No accountability:** If a student learns AI incorrectly (accepting hallucinations, misunderstanding limitations), there's no correction mechanism.
5. **Reinforces inequality:** Students in metros with tech-savvy peer groups learn differently from students in small towns with no local expertise.

Institutions could convert this underground educational economy into official learning pathways. They could validate YouTube resources, create curated learning paths, facilitate peer learning communities, and provide mentorship. Instead, many institutions are acting as gatekeepers, trying to prevent students from accessing the tools students are already using.

This is strategically backwards. The energy students have for AI learning may be channelled, not blocked.

The Faculty Bottleneck: Not a Technology Problem

Every conversation with institutional leaders and employers arrived at the same diagnosis: the bottleneck is not technology or student access, but faculty readiness.

Faculty are overwhelmed because they are being asked to teach in an AI-integrated world using teaching methods from the pre-AI world. A statistics professor can no longer assume that students will hand-compute regression analyses, as AI does this instantly. So, what should she teach? How to evaluate AI-generated analyses. How to catch mistakes in AI's reasoning. How to apply business judgment to AI's outputs. This is a fundamentally different pedagogical challenge. It requires faculty to think differently about what they are teaching and how to assess it.

Yet most institutions have invested almost nothing in helping faculty make this transition. There have been few sustained faculty development programs. There is minimal guidance on AI-aware pedagogy. There is no systematic way for faculty who have figured out how to teach with AI to share their approach with peers.

The result: Many faculty are falling back on the safest option by banning AI. This allows them to teach courses unchanged. This has the negative consequence that students don't develop genuine capability.

The solution is clear: Invest in faculty capability through sustained, peer-led, practice-focused training rather than organising one-off workshops or issuing top-down mandates. Further develop learning communities where faculty experiment, share, and iterate. Incentivise and recognise innovation while dedicating adequate time and resources to curriculum redesign.

The Employability Signal: Critical Thinking Over Credentials

One of the clearest insights from employer conversations is a shift in what signals capability:

Before: Credentials (degree from IIT, score of 85%, certificate in "AI"), technical depth (knows advanced algorithms, can code in three languages), resume (list of skills)

Now: Critical thinking, adaptability, portfolio evidence, demonstrated problem-solving, ability to learn fast, cultural fit and communication

In transition: AI tool fluency (important, but not the primary filter)

Employers are realising that specific tool knowledge becomes obsolete in 18 months. While it may be ChatGPT today, it could be some other new tool tomorrow. What matters is the ability to pick up new tools and apply them wisely. This requires critical thinking, curiosity, and the ability to evaluate trade-offs.

Yet institutions are not preparing students adequately for these changes. Students memorise content, list credentials, and hope for job offers. This mismatch is driving student anxiety and poor career choices. A humanities student with strong critical thinking and communication skills is often more employable than an engineering student with weak problem-solving, despite knowing more technical content.

Bridging the gap: Employers need to engage deeply with Academic Institutions to promote skills that matter most for employability, such as critical thinking, communication, and problem-solving, explicitly. Mentorship that counsels students through portfolios and task-based auditions, not exams and credential accumulation, could go a long way in improving employability. They should be helping students understand that while AI tool fluency matters, it is the above critical skills that will be the key differentiator in preparing them for the job market in the AI era.

The Policy Imperative: A Narrow Window

The findings converge on a single policy imperative: **India has an 18-month window to coordinate institutional transformation to improve student employability.**:

1. **Students are using AI right now:** They will graduate in 2025-2026, having been shaped by current institutional responses. If those responses are incoherent (some faculty ban, others encourage, and most are silent), students graduate underprepared for the job market.
2. **Faculty decisions made now will shape pedagogy for the next 3-5 years:** A faculty member who says "I'll never integrate AI" in 2025 will teach the same way in 2028. A faculty member who begins experimenting with AI-aware assessment in 2025 will refine their practice and improve outcomes across hundreds of students over the years to dramatically impact their employability.
3. **Institutional culture shifts take time:** Starting now with policy redesign, curriculum mapping, and faculty development with the objective of improving employability means substantive change will likely be visible twelve to eighteen months from now.
4. **Tier 2/3 institutions need external support to move:** Without targeted government funding and master trainer deployment, these institutions will fall further behind. This support requires policy design, resource mobilisation, and institutional partnerships, all of which can take up to 6-9 months to set up.
5. **The employer signalling is shifting now:** Forward-thinking companies are starting to hire based on portfolio evidence and capability assessment rather than credentials. This advantage will widen if institutions lag. By 2028, if graduates from Tier 1 institutions have AI-integrated training and those from Tier 2/3 don't, the hiring gap will be substantial.

5.3 What Needs to Happen (Next 18 Months)

For the Ministry of Education and Regulatory Bodies

1. Publish AI in Higher Education Guidance

- Clear guidance on what "responsible AI usage" looks like in academic contexts
- Framework for institutions to move from bans to structured integration
- Recommendation that detection-based enforcement (AI flagging) be deprioritised in favour of transparency-based assessment

2. Announce Resource Commitment

- Yearly investment for faculty training, device access, and infrastructure
- Funding mechanism: National fund for Tier 2/3 institutions

3. Commission Research and Evidence

- Rigorous evaluation studies: Do redesigned courses improve learning? Do AI-integrated programs produce more employable graduates?
- Equity studies: Are Tier 2/3 institutions closing gaps?
- Faculty adoption studies: What enables vs blocks institutional change?

For Institutions

1. Publish Clear AI Usage Policies

- Move from implicit (ambiguous) to explicit (transparent)
- Frame AI as a tool for learning, not cheating
- Require disclosure and process documentation, not policing

2. Assessment Redesign Pilots

- Redesign 2-3 core courses per institution
- Shift from detection-based to process-based evaluation
- Rubrics that evaluate thinking, not originality
- Collect data on student outcomes

3. Faculty Tool Access and Training

- Budget allocation
- Structured training: 8-12-week peer-led cohorts
- Focus: AI tool fluency, pedagogy redesign, assessment methods

4. Industry Problem Integration

- Partner with 2-3 local companies
- Embed real, contextualised problems in courses
- Student teams work with industry mentorship

5. Become Model Institutions

- Publish findings, share pedagogy frameworks, and host faculty from lesser endowed institutions
- Deploy master trainers to help other lesser-endowed institutions
- Contribute to the national evidence base

Additionally, lesser endowed Institutions may undertake

1. Clarity Campaigns and Curriculum Mapping

- Simple, clear messaging: "AI is a tool, learn to use it wisely"
- Identify AI champions among faculty
- Start awareness conversations
- For each program: where does AI naturally integrate?
- Design domain–AI learning outcomes
- Sequence of courses for redesign

2. Student Mentorship Program

- Recruit mentors from alumni, industry, and professionals
- Pair with students (5-10% participation in phase 1)
- Focus: Career guidance and AI application in roles

3. EdTech Partnerships and Device Access Programs

- If needed, partner with skilling platforms for curriculum and project access
- Vet partnerships; prioritise quality over breadth
- Government subsidy for laptops
- Community labs with internet access

5.4 The Equity Imperative

These suggested recommendations will fail if the underlying inequality is not addressed. Therefore, equity must become an integral element of every policy:

Device Access: Yearly subsidy for laptops and community labs ensures that laptop poverty doesn't prevent participation. This will likely be the single largest cost item in the budgets, and this expenditure is justified because device access is a prerequisite for genuine capability development.

Tier 2/3 Institutional Support: Targeted investments to close gaps such as regional GPU labs and master trainers, will ensure students in these institutions have equal access to opportunities.

Mentorship Access: Mentorship platform and coordination ensure that students in Tier 2/3 colleges without industry proximity can still access professional guidance. This is critical for reducing the advantage of geography and family networks.

Faculty Development: Government subsidies or CSR funding for training ensure that faculty capability is not limited by institutional budget.

Indian Language Support: AI interfaces in multiple Indian languages and using AI to generate

simplified learning materials in these languages will ensure that English fluency is not a barrier. In fact, a goal that ought to be set for sovereign foundational models being funded by the IndiaAI Mission should be to enable code generation using AI in any of the multiple Indian languages.

With these suggested equity commitments, a level playing field can be created in the entry-level job market.

5.5 What Success Looks Like (By 2028)

If these recommendations are implemented systemically:

For Students:	For Faculty:
<ul style="list-style-type: none"> Clear policies and no fear of punishment for legitimate AI usage Access to mentors who guide wise AI usage Curriculum that teaches both domains and AI integration Assessment that evaluates thinking, not tool usage Portfolio guidance that helps them demonstrate capability Realistic career clarity tied to the actual job market 	<ul style="list-style-type: none"> Tool access to experiment and learn Peer learning communities for support and knowledge sharing Curriculum redesign guidance and time Clear pedagogy frameworks for AI-era teaching Professional recognition for innovation
For Institutions:	For Employers
<ul style="list-style-type: none"> Coherent AI strategy aligned with employer needs Measurable outcomes showing improvement Competitive advantage, attracting students and employers Reputation for preparing graduates for AI-era roles 	<ul style="list-style-type: none"> Graduates with genuine AI capability, not just exposure Diverse talent pipeline from Tier 1, 2, and 3 institutions Clear portfolio evidence of thinking and judgment Reduced need for reskilling on entry-level hires
For India:	
<ul style="list-style-type: none"> More equitable distribution of AI capability across geography and institutions Graduates that are globally competitive in AI-augmented roles Reduced technology-induced job anxiety National advantage in human-centric AI roles (where judgment and context matter) Stronger talent pipeline for AI innovation in every sector 	

5.6 The Non-Negotiable Commitments

If implemented, this plan requires unwavering commitment to five principles:

1. Faculty Capability Is First

No other investment will succeed without it. If we build platforms, curricula, and policies but leave faculty unprepared, students will experience fragmented instruction and mixed signals. Faculty development must be prioritised by provisioning for adequate time, resources, alongside peer support and incentives for recognition.

2. Tier 2/3 Institutions Cannot Be Left Behind

If only Tier 1 benefits, inequality widens structurally. The recommendations are designed so that Tier 2/3 institutions can achieve 50-70% of Tier 1 outcomes within 18 months with appropriate support. This requires targeted resources and master trainer deployment, but it is achievable.

3. Outcomes Matter More Than Adoption

Success is not "percentage of students exposed to AI" or "number of AI courses." Success is "percentage of graduates with genuine capability to use AI effectively in their roles" and "employer satisfaction with AI-augmented capabilities." Every policy must be evaluated on these metrics.

4. Student Voice Must Be Heard

Students have shown remarkable agency in learning AI despite institutional barriers. Institutions should trust this agency, enable it, and guide it and not suppress it. Policies should be co-created with students, not imposed on them.

5. Coordination Prevents Fragmentation

Without a coordination mechanism, different institutions will make different choices, creating incoherence. National-level coordination and regular data collection are essential for learning what works and scaling it.

5.7 The Call to Action

This report documents both crisis and opportunity. The crisis is real: institutions are largely unprepared; faculty lack capability; inequality is widening. But the opportunity is equally real: clear recommendations exist; resources are available at an affordable scale; leading institutions have demonstrated workable models; student motivation is high.

The choice is not whether AI will transform higher education. It will. The choice is whether India's institutions will lead that transformation or be left behind by it.

For Government:

National level commitment to the goal of creating a level playing field in the entry-level job market through a coordinated action plan backed by resources and guidance.

For Institutions:

Publish policies. Invest in faculty. Redesign assessments. Partner with industry. Start now. Don't wait for policy signals that may be slow in coming.

For Employers:

Signal what you actually need. Hire for capability, not credentials. Support institutions through mentorship, problem statements, and hiring clarity.

For Students:

Advocate for clarity. Push institutions to teach you what matters. Build portfolios that demonstrate capability. Don't accept ambiguous policies or outdated pedagogy.

For Faculty:

Experiment. Share what works. Build peer learning communities. Don't retreat to bans and policing. The future of education is in your hands.

5.8 Conclusion

India's AI-era higher education revolution is happening now. The question is not whether institutions will transform, but whether they will transform deliberately, equitably, and in coordination with each other or chaotically, unequally, and fragmented.

This report provides a possible roadmap based on 40+ stakeholder interviews, conversations with 45 students, documented case studies from leading institutions, and extensive analysis of what actually happens in classrooms across India.

The pathway to transformation is clear. The resources are available. The time to act is now.

The next 18 months will determine whether India's graduates enter the workforce prepared for AI-augmented roles or saddled with credential anxiety and shallow capability. Leaders who act now will shape this outcome.

Appendices

Appendix A: Quick Reference for Policymakers

If you implement only 3 things:

1. Clear institutional AI policies (enable, don't police)
2. Faculty tool access and training cohorts (addresses bottleneck)
3. Student mentorship program (provides guidance, reduces anxiety)

Timeline: Months 1-6 | **Impact:** Transforms immediate experience for 80% of students

If you have 18 months for full transformation:

Follow the phased roadmap in Section 5:

- **Months 1-6:** Immediate actions (policies, access, mentorship, pilots)
- **Months 6-12:** Mid-term initiatives (faculty training at scale, curriculum mapping, portfolio systems)
- **Months 12-18:** Long-term transformation (curriculum redesign, assessment overhaul, mentorship scale-up, Tier 2/3 support)

Impact: Systematic transformation across all tiers

Measurement of Success (by Month 18):

- 100% of institutions have clear AI policies
- 50% of faculty trained in AI-aware pedagogy
- 40% of students with documented portfolio projects
- 80% of faculty with AI tool subscriptions
- Employer satisfaction surveys show measurable improvement

This is India's moment to lead AI-augmented higher education transformation. The evidence is documented. The pathway is clear. What's required now is commitment and coordination.

Appendix B: Name of contributors who shared their perspectives

Bibin Shivas Vice President - Operations Ati Motors	Aryan Yadav Co-founder Neosapiens	Amit Garg CEO MXV Consulting	Hariharan Subramanian Head HR India Iris software
Shreya Krishnan CEO AnitaB.org	Roydon Gonsalvez CHRO Northern Trust	Manmeet Sandhu CHRO PhonePe	Neeraj Sharma CEO V18hub

Amit Gujar CEO Purva Infotech	Krishna Kumar CEO Green Pepper	Sangeetha Vijay Head of HR Gnani.ai	Selvam George Chairman 5E Serpraise
Smitha Murthy HR Leader	Bharadwaj R Tech Startup Leader	Ganesh S CHRO More Retail	Tan Moorthy CEO Revature
Ishan Kapur CEO Hoping Minds	Deepshikha Thakur Chief People Officer Bikaji Foods International Ltd	Vivek Sapre Group CHRO Veranda Learning	Deepak Chawla Founder Hidevs
Ashish Kulkarni Faculty Learnforeverybody	Viplav Baxi CEO AmplifyU	Bharath Reddy Faculty Thakshala Institute	Rishabh Cecil Antony Director NSHM Knowledge Campus
Mino Thomas Sr Director Talent & Global Head Talent Operations Adobe	Ankit Aggarwal CEO Unstop	Smitha Sircar Deputy CEO Labournet Foundation	Arvind Singhal Independent Consultant
Prof. Arkapravo Sarkar Assistant Professor Suffolk University	Subha Mahata Vice President Karnataka Digital Economy Mission (KDEM)	Sarita Digumarti Consulting Partner Dono Consulting	Srikanth Iyengar CEO UpGrad Enterprise
Satvik Paramkusham Founder Build fast with AI	Vijayanti Margassery HR Leader	Swati Rustagi HR Leader	Anand Kumar Professor - Economics Azim Premji University
Dr Sandeep Das Senior Vice President	45 higher education students via various focus group discussions		

Appendix C: Purpose, Methodology, and Research Foundation

Purpose and Motivation

This research report emerges from a simple but urgent question: **What are students actually doing with AI, and how is higher education preparing them for the future of work?**

The answer, we discovered, is more complex and more critical than it might first appear. While generative AI has captured headlines and sparked anxiety across society, the story playing out in classrooms, dormitories, and career conversations is far more nuanced. Students are not waiting for institutional guidance; they are actively learning AI, building capability, and navigating a future that educational institutions have barely begun to acknowledge.

This report documents what is really happening. Through 85+ conversations with industry leaders, faculty members, students and EdTech providers, and weeks of analysis of real-world practices, we map the gap between institutional readiness and student agency. We identify where the bottlenecks are, where transformation is already occurring, and what policy pathways can enable systemic change.

This is not a report about what AI should do for education. It is a report about what education must do to stay relevant in an AI-augmented world.

Appendix D: Research Scope and Partners

This research was conducted through partnerships with multiple stakeholder groups:

- Industry & Corporate Partners
- EdTech & Skilling Ecosystem Partners
- Higher Education Institution Partners
- Government & Policy Partners
- Researcher & Analysis Partners

Contributors to research design, primary conversations, data analysis, and synthesis:

- Primary researcher: Adarsh Lathika

Appendix E: Methodology Overview

This research triangulates data from four primary sources:

1. One-on-One Stakeholder Interviews (40+ Interviews)

Sample: 40+ leaders across industry, EdTech, academia, and government

Interview Types & Duration:

- **Industry interviews:** 60-90 minutes, semi-structured
 - **Focus:** AI adoption workflows, job role evolution, skill expectations, hiring practices
 - **Sample:** 19 industry leaders (CHROs, CEOs, HR heads, founders)

- **EdTech interviews:** 45-60 minutes, semi-structured
 - **Focus:** Learning effectiveness, curriculum design, student skill gaps, institutional partnerships
 - **Sample:** 10 EdTech leaders and providers
- **Academic interviews:** 45-60 minutes, semi-structured
 - **Focus:** Pedagogical challenges, assessment redesign, faculty readiness, policy needs
 - **Sample:** 5 faculty members and institutional leaders from Tier 1, 2, and specialised institutions
- **Government & Policy interviews:** 60 minutes, structured
 - **Focus:** Policy direction, resource availability, state-level implementation challenges
 - **Sample:** 1 state government representative (KDEM)

Key Topics Covered:

- Current AI adoption in workflows and decision-making
- Impact on job roles, skill requirements, and hiring criteria
- Assessment methods for AI-era capability
- Institutional barriers and enabling factors
- Education and training gaps
- Recommendations for higher education transformation
- Equity and access considerations
- Future workforce requirements (two-five-year horizon)

Geographic Coverage: Bengaluru, Delhi, Mumbai, and national/regional scope

Data Management:

- All interviews were recorded with participant consent
- Transcribed and coded for thematic analysis
- Anonymised in reporting (titles and organisations preserved, names protected where appropriate)

2. Focus Group Discussions with Students (45 Conversations)

Format: Semi-structured focus groups and individual conversations

Duration: 20-40 minutes per session

Locations: Bengaluru, Delhi, Kolkata

Participant Type: Mix of undergraduate and postgraduate students from different tiers

Focus Areas:

- Lived experiences with AI tools in academic work
- Barriers and enablers to effective AI use
- Institutional policy experiences (supportive vs restrictive)
- Career concerns and aspirations
- Support needs and preferred guidance channels
- Peer learning and knowledge-sharing practices

- Informal learning pathways (YouTube, communities, peer sharing)

Data Approach: Thematic analysis of discussion transcripts, pattern identification across locations and student cohorts

3. Secondary Benchmarking and Literature Review

Sources:

- Global research on AI task exposure and job market impact (academic papers, industry reports)
- Institutional policy documents and curriculum frameworks
- Job postings and hiring requirement trends
- EdTech industry reports and market analysis
- Government education policy documents (NEP 2020, etc.)

Focus: Understanding global context and comparative perspective on AI adoption in higher education across countries

Data Analysis Approach

Qualitative Analysis

- Thematic coding of interview transcripts (open coding focused coding theme development)
- Pattern identification across stakeholder groups (industry, EdTech, academia)
- Case study synthesis from institutional observations
- Triangulation with survey data for validation

Synthesis & Framework Development

- Cross-case analysis comparing perspectives across stakeholder groups
- Framework development (AI Readiness Index, Student Maturity Model, Policy Spectrum, etc.)
- Recommendation generation grounded in evidence
- Validation through expert review with research partners

Research Quality & Limitations

Strengths

- **Multiple data sources:** Triangulation across interviews, surveys, focus groups, observations
- **Diverse stakeholder representation:** 40+ interviews spanning industry, EdTech, academia, government
- **Large student sample:** 45 conversations, across multiple tiers and fields
- **Geographic breadth:** Metro cities, regional centres, small towns; international perspective
- **Rigorous analysis:** Thematic coding, pattern identification, framework validation
- **Practical focus:** Emphasis on actionable findings grounded in real-world contexts

Limitations

- **Self-reported data:** Survey and interview responses reflect participant perceptions, not necessarily the objective truth
- **Potential bias:** Respondents willing to participate may differ from those who declined
- **Geography:** While India is the primary focus, global perspectives are included; generalisability may vary

- **Institutional selection:** Not a random sample of all institutions; partnership-based selection
- Response burden: Qualitative interviews reflect volunteer participants

Mitigation Strategies

- Multiple data sources reduce reliance on a single perspective
- Large student sample size increases confidence in patterns
- Thematic analysis codes for alternative interpretations
- Clear attribution of findings to specific research sources
- Transparent acknowledgement of limitations in findings

Appendix F : Use of AI in Report Preparation

This research report documents higher education's readiness for the AI-era learning. Appropriately, we used AI tools throughout the research and writing process. Full transparency about these tools is essential for credibility.

AI Tools Used in the Research & Writing Process

1. ChatGPT (OpenAI)- Interview Analysis & Synthesis

Purpose: Analysing and synthesising qualitative interview data

Specific Use Cases:

- Thematic coding of interview transcripts (identifying patterns, themes, categories)
- Generation of initial theme summaries from coded interviews
- Cross-interview pattern identification (comparing responses across stakeholders)
- Draft synthesis of findings from raw interview notes
- Validation of coding patterns against survey data

Process:

- Raw interview transcripts (anonymised) were provided to ChatGPT
- Output was reviewed, validated, and refined by human researchers
- Final synthesis was verified against original transcripts

Transparency Note: All ChatGPT output was treated as draft analysis, not final findings. Human researchers reviewed all outputs, validated against source data, and made final decisions on inclusion/exclusion. This is standard practice in qualitative research (using software to assist analysis, not replace human judgment).

2. Google NotebookLM (Google)- Literature & Secondary Source Summarisation

Purpose: Summarising YouTube educational content and other online resources

Specific Use Cases:

- Summarising YouTube educational videos on AI and higher education
- Extracting key points from online articles and blog posts

Process:

- Educational videos and secondary sources were selected by researchers

- Google Notebook was used to create structured summaries
- Key insights were extracted and organised
- Summaries were reviewed for accuracy and relevance
- Only validated points were carried forward to the final analysis

Transparency Note: Google Notebook was used as a research organisation tool, not a primary source. Secondary sources enhanced context but were clearly distinguished from primary research findings.

3. Perplexity (Perplexity.com) and Google Gemini - Preliminary Report Drafting & Visualisation

Purpose: Drafting report sections and generating visual concepts

Specific Use Cases:

- Drafting preliminary versions of report sections based on research findings
- Generating visualisation concepts and chart descriptions
- Writing explanatory text for complex findings
- Creating bulleted summaries and framework descriptions
- Organising research findings into a coherent narrative structure

Process:

- Research findings were organised into structured inputs
- Perplexity was provided with explicit instructions and research data
- Draft outputs were reviewed, fact-checked, and refined by researchers
- All claims were validated against primary research data
- Final text was substantially revised, reorganised, and fact-checked

Transparency Note: Perplexity output served as a starting point for drafting, not final content. All final report sections were reviewed for accuracy, alignment with research data, and completeness. Researchers made all final decisions on what to include, emphasise, and exclude.

What AI Did NOT Do

To be clear about the boundaries of AI use:

- **AI did not replace human analysis.** All major interpretations, conclusions, and recommendations emerged from human researchers working with the data
- **AI did not determine what to include in the report.** All editorial decisions, including what findings matter, how to organise the narrative, and what recommendations to prioritise were made by researchers
- **AI did not generate primary data.** All survey, interview, and focus group data came from direct human engagement with stakeholders
- **AI did not validate findings.** All validation against source data and expert review was conducted by human researchers
- **AI did not make policy recommendations.** While AI helped draft explanatory text, the underlying logic and recommendations came from research findings and expert judgment

Why Transparency Matters

In a report about preparing students for AI-era work, it would be hypocritical not to transparently disclose our own AI use. This transparency serves multiple purposes:

1. **Credibility:** Readers can understand the research process and assess strengths/limitations
2. **Accountability:** We model responsible AI use (transparency about tools, validation of outputs, human judgment on decisions)
3. **Practical Learning:** This research documents best practices for responsible AI use in professional work
4. **Alignment with Findings:** The report recommends that students learn to use AI with transparency and critical evaluation. We practised this ourselves

Conclusion on AI Use

AI tools significantly accelerated our research and writing process. They helped us organise complex qualitative data, synthesise diverse sources, and draft initial versions of findings. But they did not replace human judgment, critical thinking, or responsibility for the final product. This is precisely the model we recommend for AI use in higher education: tools that augment human capability, with humans retaining decision-making authority and accountability.

Appendix G: Acknowledgements & Thank You

This research would not have been possible without the generosity of our 40+ stakeholder partners who gave their time, shared candid perspectives, and engaged deeply with our questions. Their insights, which were often grounded in their experience managing real consequences of AI disruption, have formed the foundation of this analysis.

We thank Centre of Policy Research and Governance (CPRG), Ashoka University, St. Joseph's College of Commerce, and others who opened their classrooms and shared their pedagogical experiments. Their willingness to share both successes and struggles provides concrete proof that transformation is possible.

Finally, we acknowledge the critical importance of regional coordinators in Bengaluru, Delhi, and Kolkata who facilitated focus group discussions and provided on-ground context that enriched our understanding.

This research is offered in the service of institutional transformation and more equitable access to AI-era learning across India's higher education system.

Note on Accessibility & Updates

This report was published in January 2026. Given the rapid pace of AI development and change in higher education, we expect that the findings, recommendations, and tools mentioned will evolve. We welcome feedback, suggestions, and updated information from practitioners implementing these recommendations.

The report's core findings about institutional readiness, faculty capability gaps, and policy pathways are likely to remain relevant for two-three years. More specific recommendations about tools, platforms, and timelines may require periodic updates.

We encourage institutional leaders, faculty, and policymakers to treat this report not as a fixed blueprint but as a living framework to adapt to local context.



Scan the QR code to access the full, detailed version of the report.

Join the Transformation



AI4India

AI4India.org | January 2026

