

August 15, 2026

## The Crisis of Academic Formation in the Age of AI

By: Aniruddha Jena, Usha Raman

*Generative AI collapses the boundary between assistance & substitution in higher education, simulating rather than summarising thought. While AI can mitigate the inequalities among students, it erodes the sustained, often confusing encounter with ideas that genuine intellectual formation requires.*

A doctoral student submits a literature review. It is fluent, well-organised, and draws on the right sources. The supervisor reads it carefully and senses, without being able to say exactly why, that something is missing. The student, when asked about one of the central arguments, fumbles. The text was there. The understanding was not.

|| We share a concern that has sharpened in recent months—we are beginning to mistake the production of academic output for the production of understanding.

This is not a hypothetical. It is what plays out in research seminars and supervision meetings across Indian universities, and it is becoming harder to ignore. Generative AI tools (ChatGPT, Gemini, Copilot, and their proliferating variants) can now produce summaries of complex texts, synthesise bodies of literature, draft analytic arguments, and write in the register of academic prose. The outputs are often good enough to pass. The question is their cost.

We write this as scholars at two different stages of academic life. One of us is an early career researcher at IIT Mandi, preparing to supervise doctoral students for the first time, and living inside the pressures that shape how young researchers work today. The other is a retired academic with decades of experience forming researchers, and now in a position to reflect from some distance on what that process actually requires.

We do not agree on everything. But we share a concern that has sharpened in recent months—we are beginning to mistake the production of academic output for the production of understanding.

### Shortcut that Does the Work

There is nothing new about students looking for ways to reduce the labour of learning. Generations of undergraduates have relied on guides and summaries to navigate texts they have not read or found too difficult. Study aids, readymade notes, and, in the Indian context, the long tradition of guides printed for competitive exams have always offered shortcuts around the harder work of sustained reading.

These shortcuts were visible, and their limits were clear. A student who relied on a summary of a difficult book could get through an exam, but usually could not discuss the book in any depth. One might say that the reliance on such tools, which served the competitive exam-prep market well, normalised taking shortcuts to reach the putative goal of higher education—to gain a degree. It was a perfect context for the entry of the ultimate shortcut machine—generative AI.

However, generative AI goes beyond summary or simplification. It simulates the process of thinking. It produces text that is syntactically original, contextually responsive, and difficult to distinguish from work produced by mental/intellectual labour. A student can feed a difficult theoretical text into a large language model and receive, within seconds, a lucid—even if not elegant—exposition of its argument, a critique of its limitations, and a comparison with related work. The student may never have read the text.

|| PhD programmes operate under strict timelines. Supervisors are evaluated partly on how many doctorates they produce. Students are under pressure to show visible progress at regular intervals.

This collapses a boundary that educators have long relied on—the boundary between assistance and substitution. When a student uses a guide to supplement their reading, there is still, in most cases, some reading happening, even if only at the level of summary. When a student uses AI to replace that reading, the substitution is nearly invisible. The supervisor cannot easily tell the difference. And most often, neither can students—because they have never engaged with the original text.

This matters because, especially in the social sciences and humanities, understanding is not produced by acquiring information. It is produced through a particular kind of sustained encounter with ideas, one that involves incomprehension, confusion, rereading, the slow revision of one's assumptions, and the gradual formation of a critical perspective.

For example, to read B.R. Ambedkar seriously is not to extract a set of positions. It is to encounter a mode of thinking about caste, rights, and democratic possibility that requires the reader to reorganise what they already know. That reorganisation takes time, and it is not pleasant. It cannot be outsourced. You cannot arrive at that destination without having undertaken, and experienced, the journey.

### **Pressure to Produce**

To understand why AI is being adopted so rapidly and so uncritically in Indian universities, it helps to look honestly at the institutional environment in which doctoral students are working.

Research in India-as perhaps elsewhere-is now heavily governed by metrics. The University Grants Commission, the National Assessment and Accreditation Council, and institutional ranking systems have all contributed to a culture in which the value of academic work is increasingly measured by outputs-papers published, citations accumulated, and projects completed. PhD programmes operate under strict timelines. Supervisors are evaluated partly on how many doctorates they produce. Students are under pressure to show visible progress at regular intervals.

In this environment, the appeal of tools that accelerate output is not a failure of character. It is a rational response to structural incentives. If the system rewards the appearance of productivity, and AI can generate that appearance faster, then students and researchers who use AI are not being lazy. They are simply being strategic in the terms that the system has set. The problem is that those terms are in serious tension with what research actually requires.

They describe something real about how understanding develops. The difficulty is not incidental to the learning. It is the mechanism through which learning occurs.

One of us, teaching in a research methods course, has watched doctoral students produce in an afternoon what would previously have taken a week of careful reading and writing. The students are not ashamed of this; many are proud of how efficiently they have managed their time.

What is harder to see, from inside that efficiency, is what has not happened. There are the three days of confusion and frustration while trying to understand what Michel Foucault was actually doing with the concept of discourse. There is the puzzling through cross-references while grappling with an unfamiliar idea, and the going back to the supervisor or instructor seeking clarity. There is the conversation with a classmate that shifted how a problem was framed, and the marginal note written at midnight that turned out to be the seed of a dissertation chapter.

These are not romantic anecdotes about the hardships of academic life. All of these steps towards understanding can be deeply fraught, many times more so when working in an alien language. But they describe something real about how understanding develops. The difficulty is not incidental to the learning. It is the mechanism through which learning occurs.

### **Question of Access and Fairness**

Any honest account of AI in Indian higher education has to grapple with a more uncomfortable reality. The Indian academy is shot through with inequalities that AI, in some respects, partially mitigates. Students from regional-medium schooling, from non-English-speaking backgrounds, and from institutions without strong library resources face disadvantages that are systemic and deep. For many of these students, AI tools that help with language, with finding sources, and with understanding the conventions of academic writing can reduce a gap that would otherwise limit their participation.

English continues to function as the gatekeeping language of global scholarship. A student from a rural Odisha background, working in a field where the canonical literature is entirely in English, faces a burden that their counterpart at a metropolitan institution with English-medium schooling does not. AI tools can help bridge that gap, and to deny this would be both intellectually dishonest and unfair.

For many students, AI tools that help with language, with finding sources, and with understanding the conventions of academic writing can reduce a gap that would otherwise limit their participation.

But this enabling function is not without its own costs. If AI becomes the primary means through which non-dominant-language scholars access and participate in global scholarly conversations, it may also standardise their expression in ways that flatten what is distinctive about their intellectual formation. The risk is not that they will participate less, but that they will participate on terms that are not their own, and without the resources to know that something has been smoothed away. Even where they surmise this, they lack the linguistic tools to effectively combat it.

This is not an argument against AI as an accessibility tool. It is an argument for being clear-eyed about what it does and does not do, and for not allowing the real benefits to obscure the real costs. It is also an argument for a deeper look at the colonial legacies of higher education in contexts like India.

### What is Lost

One concept from the history of science is worth introducing here, not as jargon, but because it names something important-serendipity. In the sociology of science, serendipity refers to the unexpected findings that emerge when a researcher is genuinely immersed in a problem: following an unfamiliar thread, sitting with a contradiction, noticing something anomalous. These are not lucky accidents. They are the products of sustained, open-ended engagement with a question.

In qualitative and interpretive research-the kind of work that characterises much of the social sciences and humanities in India-serendipity is not peripheral. It is where the most important insights come from. Consider the fieldworker who spends two months in a village before anything interesting seems to happen, and then one afternoon witnesses something that reorganises everything. Or the theorist who reads a book from an adjacent discipline and finds, unexpectedly, a concept that unlocks their own argument. Or the doctoral student who follows a footnote into an entirely different literature and discovers that the question they have been asking is the wrong one.

Generative AI is not structured to help someone inhabit a question, to stay productively confused, to follow an unmarked detour.

Generative AI, by design, compresses-even erases-this. It is optimised to produce efficient, coherent responses to queries. It is very good at arriving at answers. It is not structured to help someone inhabit a question, to stay productively confused, to follow an unmarked detour. When AI becomes the default mode of engaging with ideas, the exploratory space that serendipity requires begins to close.

This loss is already visible in supervision. Doctoral students who rely heavily on AI-generated literature reviews often know the terrain but cannot draw a pathway through it. They can identify what the major arguments are, but they cannot say why one is more persuasive than another, or what it would mean for their own research if they accepted it. The landscape has been mapped for them. But they have never walked through it.

### Better Incentive Structure

The temptation, when faced with a problem like this, is to reach for a reactive policy. Ban AI in assessments. Require handwritten exams. Introduce AI detection software. These are inadequate responses that tend to push the problem underground rather than address it.

The deeper challenge is institutional. Indian universities, like universities everywhere, have built systems that reward output and are largely indifferent to process. The doctoral student who spends three years in genuine intellectual formation and produces a dissertation that advances knowledge in modest but real ways is evaluated by the same metrics as the student who produces four papers in two years, regardless of how those papers came to exist. The incentive structure does not distinguish between understanding and its simulation.

There is still space to ask what universities are for, and to insist that the answer includes something more than the efficient production of academic outputs, which veers dangerously into producing academic automatons.

Changing this requires more than better pedagogy, though better pedagogy matters. It requires institutional willingness to protect time, the one thing that genuine research requires and that the present system seems increasingly reluctant to allow. It requires rethinking what doctoral training is for, and being willing to say, clearly, that it is for the formation of people capable of producing knowledge, not merely for the production of papers.

Supervisors can do some of this individually. They can design conversations and assessments that cannot be completed through AI substitution. They can ask students to walk them through their reasoning, not just their conclusions. They can create space for productive confusion rather than demanding constant visible progress. They can be welcoming of uncertainty and confusion. They can be patient with those who are grappling to gain not only conceptual understanding but also linguistic competence.

But supervisors working against the grain of an institutional culture that rewards output will always be swimming upstream. Individual pedagogical commitment cannot substitute for structural change.

### What are Universities For?

[Minas Karamanis's blog post](#) "The Machines Are Fine. I'm Worried about Us" made a point that has stayed with us and prompted the original essay (published in a scholarly journal and referenced here): the machines are fine. They are doing exactly what they were designed to do. The problem is not AI. The problem is what we are designing our universities to do.

Indian higher education is at a moment when this question is still genuinely open. The expansion of doctoral education, the growing pressure to produce internationally visible research, the increasing role of metrics in institutional governance: these are all forces that are currently in play, and their direction is not fixed. There is still space to ask what universities are for, and to insist that the answer includes something more than the efficient production of academic outputs, which veers dangerously into producing academic automatons.

The slow, frustrating, occasionally exhilarating process through which a person becomes capable of thinking independently is not a romantic ideal. It is what research education is supposed to do. When doctoral students arrive at answers without inhabiting the questions that gave rise to them, something has failed. Not the students, and not even the tools. The conditions that make genuine intellectual formation possible have been allowed to erode, and AI has made that erosion harder to ignore.

The question is not whether to use AI. It is whether we are willing to insist that the purpose of education is not the production of outputs, but the formation of people who can think. And whether our institutions are prepared to build the conditions that make that possible.

*Aniruddha Jena is an Assistant Professor at the School of Management, Indian Institute of Technology Mandi; Usha Raman taught media studies at the University of Hyderabad.*

**Acknowledgement:** This essay draws on "[Shortcuts and Scholarship: A Cross-generational Reflection on AI and Academic Formation](#)" by the authors that was published in *Media, Culture & Society* on 29 July 2026. The article is available at <https://doi.org/10.1177/01634437261472373>.