



Pratidhwani the Echo

A Peer-Reviewed International Journal of Humanities & Social Science

ISSN: 2278-5264 (Online) 2321-9319 (Print)

Impact Factor: 6.28 (Index Copernicus International)

Volume-XIV, Issue-IV, July 2026, Page No. 125-137

Published by Dept. of Bengali, Karimganj College, Sribhumi, Assam, India

Website: <http://www.thecho.in>

DOI: 10.64031/pratidhwanitheecho.vol.14.issue.04W.124



Generative Artificial Intelligence and Educational Equity in Indian Higher Education

Neetu Chettri, *Asst. Prof, Acharya Brojendra Nath Seal College, West Bengal, India*

Received: 10.07.2026; Accepted: 20.07.2026; Available online: 31.07.2026

©2026 The Author(s). Published by Dept. of Bengali, Karimganj College, Sribhumi. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>)

Abstract

Generative artificial intelligence will not advance educational equity in Indian higher education unless the conditions surrounding its use change in fundamental ways. That is the central claim of this article. A well-resourced university in Bangalore and an under-funded college in rural Chhattisgarh may both have access to the same AI tools, but their students' capacity to use those tools for genuine learning is worlds apart. The differences come down to language, infrastructure, prior schooling, teacher preparedness, and whether anyone has bothered to teach students what AI can actually do – and, just as importantly, what it cannot.

This article traces those fault lines. It looks at what generative AI genuinely offers Indian higher education – personalised learning, multilingual support, broader access to academic resources – and at the structural barriers that prevent those benefits from reaching students equitably. NEP 2020 provides the policy backdrop. Bourdieu, Bernstein, Freire, and Visvanathan supply the theoretical grounding. The argument is that equal access to a technology does not produce equal educational outcomes, and that treating AI as inherently beneficial without attending to infrastructure gaps, linguistic inequality, weak AI literacy, and the absence of clear institutional policies is a recipe for reproducing the very disparities that education policy claims to address.

Keywords: Generative artificial intelligence, educational equity, Higher education, National Education Policy 2020, Multilingual education, India

Generative AI has arrived in Indian higher education whether institutions are prepared or not. Students already use it – for assignment help, exam preparation, language polishing, concept clarification. Some teachers have started experimenting too. The policy conversation shaped partly by NEP 2020, frames technology integration as desirable. But a harder question lingers beneath the enthusiasm: who actually benefits when AI enters the classroom?

Not everyone, and not equally. That much is obvious to anyone familiar with Indian higher education. The system has long been marked by stark inequalities – in infrastructure, in language, in what different institutions can offer their students – and a

new technology does not erase those inequalities. It maps onto them. A student with reliable internet, comfort in English, and some digital fluency will use generative AI in ways that a student from a regional-language background, studying at a college with erratic electricity and no institutional tech support, simply cannot. Same tool, radically different educational experience. The promise of AI — personalised learning, broader access, multilingual assistance — rests entirely on conditions that are distributed unevenly across the higher education landscape, a point that the All-India Survey on Higher Education data continues to underscore (Ministry of Education, 2024).

Language complicates this further, and in ways that deserve more serious attention than they typically receive. Generative AI works better in English. Ask it to explain a concept in English and the output is detailed, well-organised, usually adequate. Ask in Odia, or Assamese, or even Hindi for a specialised topic, and quality drops — sometimes subtly, sometimes dramatically. Large language models perform unevenly across low-resource languages (Kasneci et al., 2023). This means the students who need multilingual support most urgently are precisely the ones getting the weakest version of it. NEP 2020 recognises language as an equity issue — its emphasis on mother-tongue instruction and multilingualism is clear on this point (Ministry of Education, 2020). Whether current AI tools can support that vision is a separate question, and the honest answer is: not yet, or at least not reliably.

What this article argues is that generative AI's value for Indian higher education cannot be assessed independently of the inequalities that structure the system. The technology is not inherently beneficial or harmful — its educational contribution depends on institutional policy, teacher capacity, digital infrastructure, language conditions, and whether students have been equipped to engage with AI critically rather than passively. This is not an argument against using AI in education. Nor is it an uncritical endorsement. It is an attempt to think through what genuinely equitable integration would require, drawing on both policy analysis and theoretical perspectives that have long grappled with questions of educational inequality.

The article moves through several stages. It examines the changing role of AI in Indian higher education, then considers the opportunities it presents and — separately, because the two should not be conflated — the challenges it raises for educational equity. Four theoretical perspectives (Bourdieu, Bernstein, Freire, Visvanathan) are brought in to explain why these challenges are structural, not incidental. The final sections address implications for policy and practice.

Generative Artificial Intelligence and the Changing Landscape of Indian Higher Education:

Generative AI is not just another digital tool added to an existing toolkit. It is qualitatively different. Earlier educational technologies stored, organised, or delivered content — learning management systems, digital libraries, online lecture recordings. Generative AI produces content. It writes explanations, drafts lesson plans, generates quiz questions, summarises research papers, composes essays. That productive capacity is what makes it both promising and problematic, and it is what has forced universities to pay attention in ways that previous educational technologies did not demand (UNESCO, 2023). Sullivan et al. (2023) observe that institutions are moving beyond treating generative AI as a novelty

and starting to develop policies around its use – though “starting to” carries a lot of weight in that sentence, and the pace varies enormously.

The Indian context adds specific layers of complexity. NEP 2020 positions technology as a means to improve quality, flexibility, and access across higher education. The policy proposed a National Educational Technology Forum (NETF) to guide evidence-based decisions about adopting educational technologies (Ministry of Education, 2020). These are sound ambitions. The difficulty is that institutional capacity to act on them is wildly uneven. A central university in Delhi with a dedicated IT department, stable funding, and faculty who attend international conferences is in a completely different position from a state-affiliated college in a semi-rural district with two overworked computer labs and intermittent internet. Both are part of “Indian higher education.” Their relationship to AI could not be more different. The OECD’s Digital Education Outlook (2023) confirms what observers of Indian education have known for some time: institutional readiness determines whether AI contributes to inclusion or entrenches existing advantages.

For teachers, the practical reality of AI is less dramatic than either the hype or the alarm suggests. AI can draft lesson plans. It can generate assessment questions. It can prepare classroom activities and offer preliminary feedback on student submissions. These are genuinely useful functions – they save time on tasks that are necessary but repetitive, and that time can theoretically be redirected towards actual teaching, mentoring, and engaging with students individually. The word “theoretically” matters here. Whether teachers actually benefit depends on whether they can evaluate AI output critically. Generative AI is not reliable by default – it produces plausible-sounding text that may contain factual errors, oversimplifications, culturally inappropriate examples, or explanations that are simply wrong in ways that are not immediately obvious. Without pedagogical knowledge and disciplinary expertise, a teacher cannot distinguish between AI output that is useful and output that would mislead students. International guidance from both UNESCO (2023) and the OECD (2023) makes this point repeatedly: AI should support teachers’ professional judgement, not substitute for it.

Students are not waiting for their institutions to develop formal positions on AI. Many are already using it – to understand concepts they struggled with in lectures, to condense lengthy readings into manageable summaries, to improve the English of their assignments, to prepare for exams. Some of this use is educationally productive. It amounts to self-directed learning supported by a responsive tool. But some of it is not. Students who accept AI-generated answers at face value, without verifying accuracy or engaging with the underlying reasoning, are not learning more effectively – they are learning to outsource thinking (Sullivan et al., 2023). The educational question is not whether students will use AI. They already do, and they will continue to. The question is whether anyone is teaching them to use it in ways that develop understanding, encourage critical engagement, and maintain academic rigour (OECD, 2024; UNESCO, 2023).

The distribution of AI’s benefits across the higher education system is deeply unequal – and this inequality deserves more direct acknowledgement than it typically receives in policy discussions. Metropolitan universities tend to have better connectivity, more devices, stronger institutional support, and greater willingness to experiment. Institutions serving rural or economically marginalised communities often lack basic digital infrastructure, let alone the capacity to integrate emerging technologies meaningfully. This

is not only about hardware. It shapes whether teachers can access professional development, whether students develop digital confidence, and whether institutions can respond to new technological possibilities at all (Ministry of Education, 2024). The OECD (2023) identifies institutional readiness as a decisive factor in whether AI-supported learning becomes more inclusive or simply widens existing gaps.

Language, again, is central. India's higher education system is linguistically diverse in a way that AI systems were not built to handle. Students arrive at university from schooling in regional languages and encounter English-dominant academic environments. AI applications increasingly claim multilingual capability, but the reality is mixed. Explanations in Hindi or Tamil or Bengali may omit local expressions, miss culturally specific references, or use subject terminology inaccurately – deficiencies that English-language outputs rarely display. Language is not a transparent medium. It shapes comprehension, shapes academic confidence, shapes whether a student feels the system is speaking to them or past them. The multilingual vision of NEP 2020 remains profoundly relevant here (Ministry of Education, 2020; UNESCO, 2023), and generative AI's ability to serve that vision is, at present, incomplete.

What is changing is not just the tools available to teachers and students. It is the relationship between them and the knowledge they work with. Whether that shift proves productive depends on institutional willingness to plan, invest, and confront the uneven conditions under which AI is being adopted. The question that matters is not what AI can do in principle – it is what AI can do in the specific, stratified reality of Indian higher education.

Opportunities of Generative Artificial Intelligence for Indian Higher Education:

Acknowledging the problems should not mean ignoring the possibilities. When conditions are favourable – when infrastructure works, when teachers know how to evaluate AI output, when students have been taught to engage critically – generative AI can genuinely strengthen teaching and learning. NEP 2020 envisions technology as something that complements classroom practice rather than replacing it, and AI fits that vision when deployed thoughtfully (Ministry of Education, 2020; UNESCO, 2023).

Personalised learning is the most frequently cited opportunity, and for good reason. Indian university classrooms are often large – two hundred, three hundred students in a single lecture are not unusual. Teachers cannot individually tailor explanations for each student, cannot know which concepts each person found difficult, cannot offer feedback on every piece of written work at the pace students need it. AI tools can fill part of that gap. They generate explanations at different levels of complexity, produce practice exercises, summarise dense academic texts, and respond to questions at any hour. Students can revisit material they struggled with, at their own pace, without waiting for the next office hour or tutorial session. Research suggests that this kind of scaffolded support can encourage self-regulated learning and engagement with course material beyond the classroom (UNESCO, 2023; Sullivan et al., 2023).

The qualification is important, though. Personalised learning works when AI supplements human teaching, not when it replaces the interaction altogether. A student confused about a concept often needs something AI cannot provide – a conversation, a contextually grounded example, a teacher who notices the specific nature of the misunderstanding. AI can rephrase an explanation. It cannot read a student's face and

realise the problem is not comprehension but confidence. Teachers remain indispensable for cultivating critical thinking, for helping students weigh competing interpretations, for connecting abstract concepts to disciplinary practice.

Teachers themselves benefit in practical, unglamorous ways. Writing lesson plans takes time. Designing assessment questions takes time. Preparing instructional materials, drafting feedback comments, organising classroom activities – all of it requires hours that could be spent on higher-value pedagogical work. AI can produce rough drafts of these materials. Teachers then revise, adapt, and contextualise them. The OECD (2023) and UNESCO (2023) both emphasise this division of labour: let AI handle preparatory and administrative tasks while keeping pedagogical responsibility and professional decision-making firmly with the teacher. The caveat, which bears repeating, is that AI-generated materials can be wrong – factually, culturally, pedagogically. The teacher who uses AI output without reviewing it carefully is not saving time. They are introducing risk into the classroom.

Access to learning resources is where the equity argument becomes most compelling. Not every college has a good library. Not every student can afford textbooks or access academic databases. First-generation learners often lack the tacit knowledge that more privileged students absorb from family and social networks – knowing how to approach academic writing, how to structure an argument, which sources are credible. AI can partially fill these gaps: explaining unfamiliar terms, suggesting further reading, helping students organise ideas for an essay. The word “partially” is deliberate. AI is not a replacement for institutional investment in libraries, tutoring programmes, or academic mentoring. But for students who currently have very little support, even partial assistance makes a material difference (Ministry of Education, 2020; UNESCO, 2021; Sullivan et al., 2023).

Multilingual learning is a more complicated opportunity. AI-generated content in Indian languages has improved, but remains uneven. Translation, summarisation, and bilingual learning support are available in several languages now, and when teachers review and contextualise that output, students can use it as a bridge between their home language and the English-medium academic environment. The phrase “when teachers review” carries significant weight. In specialised fields – law, medicine, engineering – translation errors are not just inconvenient; they can be educationally harmful. AI can support the multilingual vision of NEP 2020 (Ministry of Education, 2020; UNESCO, 2023), but it cannot deliver it unassisted.

Online and blended learning environments make AI support more relevant still. Higher education extends well beyond the physical campus now. Students learn from home, from workplaces, from rural areas with intermittent connectivity. AI tutoring, automated formative feedback, and interactive exercises can enrich these flexible arrangements – particularly for students managing work, family, or distance alongside their studies. But the benefits remain contingent on infrastructure. Blended learning enhanced by AI means nothing to a student who cannot reliably connect to the internet. Without stable connectivity, affordable devices, and institutional commitment to supporting distance learners, the opportunities remain promises rather than realities for many (OECD, 2023; Ministry of Education, 2024).

These opportunities are real. They are also conditional – conditional on infrastructure, on teacher preparation, on institutional policies that treat equitable access as a serious commitment rather than a line in a mission statement. Whether existing conditions allow these benefits to reach broadly, or whether they remain concentrated among students who were already better served has been discussed further.

Challenges to Educational Equity in the Use of Generative Artificial Intelligence:

The opportunities outlined above exist within a system that is already stratified. Generative AI does not enter a level playing field – it enters Indian higher education as it actually is, with all its inequalities of language, infrastructure, institutional capacity, and preparedness intact. Those inequalities determine who benefits from AI and who does not. Four areas warrant close examination.

Language Diversity and Multilingual Learning:

Language is not a peripheral concern in Indian higher education – it is a defining one. English dominates instruction in many universities, yet a substantial proportion of students completed their schooling in Hindi, Tamil, Telugu, Bengali, Marathi, Kannada, or one of dozens of other languages. They enter English-medium classrooms carrying the linguistic resources of their schooling, which may or may not map onto the academic register expected of them. NEP 2020 explicitly promotes multilingual education and the use of Indian languages as pathways to improved access and learning outcomes (Ministry of Education, 2020). This is a recognition, overdue in some respects, that language determines who can participate fully in higher education and who is forced to participate at a disadvantage.

Generative AI could, in theory, help bridge this gap – translating materials, explaining concepts across languages, assisting students with academic writing in English when that is not their strongest language. In practice, performance varies drastically. An explanation generated in English tends to be detailed, coherent, and reasonably accurate. The same request in Odia or Assamese often produces something thinner, less nuanced, occasionally misleading. This is well-documented in the literature on large language models and low-resource languages (Kasneci et al., 2023; UNESCO, 2023). The irony is bitter: students who most need multilingual support receive the weakest version of what AI can offer.

Translation, in any case, is not the whole story. Learning is not simply a matter of converting information from one language to another. Concepts in law, economics, or sociology carry specific resonances that emerge through classroom discussion, local examples, and lived cultural context. An automated translation that is technically correct may still strip away the contextual meaning that makes a concept intelligible to a specific group of learners. AI should function as a supplementary resource in multilingual education, one that is guided and verified by teachers who understand the linguistic and cultural realities of their students. It should not be treated as a substitute for that understanding.

The Rural-Urban Digital Divide:

India has expanded digital connectivity significantly. That is true, and worth acknowledging. It is equally true that the expansion has been deeply uneven. In most cities, reliable internet is a background assumption. In large parts of rural India, connectivity is intermittent, devices are shared among family members or unavailable

altogether, and electricity supply — the most basic prerequisite — cannot always be counted on. These are not minor obstacles. They are structural barriers that determine whether a student can participate in AI-supported learning at all (Ministry of Education, 2024).

The COVID-19 pandemic illustrated this with painful clarity. Well-resourced urban institutions moved to online learning with relative speed. Students in rural and economically marginalised areas frequently could not join — or joined under conditions so constrained that meaningful learning was impossible (UNESCO, 2023). The pandemic has receded, but the digital infrastructure gap it exposed has not. The OECD (2023) identifies unequal digital infrastructure as one of the most significant barriers to equitable AI adoption in higher education globally, and the pattern holds with particular force in the Indian context.

Educational equity in AI-supported learning requires more than distributing access to applications. It demands sustained investment in the enabling conditions — reliable connectivity, appropriate devices, institutional capacity — and focused attention on communities that have been chronically underserved. Without these investments, technology does not narrow inequality. It entrenches it, sometimes in ways that are harder to see than the inequalities that came before.

AI Literacy among Teachers and Students:

Access and literacy are different things, and conflating them produces bad policy. A student who can open ChatGPT but does not know how to evaluate its output, recognise its biases, or understand its limitations is not meaningfully literate in AI use. A teacher who can prompt a language model but cannot assess whether its response is pedagogically appropriate or factually sound is in a similar position. AI literacy — encompassing technical competence, critical evaluation, ethical awareness, and capacity for informed decision-making — is not something most teachers or students in Indian higher education have been formally taught (UNESCO, 2023). Ng et al. (2024) propose a framework for AI literacy that integrates these dimensions, and it is a useful conceptual tool. But as a description of current practice in most Indian institutions, it describes an aspiration, not a reality.

What actually happens is improvisation. Teachers experiment independently, without structured training or institutional guidance, drawing on whatever they can find online or learn from colleagues. Students figure out AI through trial and error. Some develop productive practices — using AI to test their understanding, to generate alternative explanations, to prepare for discussions. Others develop habits that are educationally corrosive, like pasting questions into ChatGPT and submitting whatever comes back. The variation is enormous, and largely invisible to institutions that have not built systems for monitoring or supporting AI use.

Professional development is the obvious response, but it needs to be the right kind. A single workshop on “how to use ChatGPT” accomplishes very little. What teachers need is sustained, discipline-contextualised support — learning to evaluate AI output for accuracy and cultural relevance, to design activities that use AI without reducing intellectual demand, to model responsible engagement for students. Students, for their part, need to be taught — explicitly, not incidentally — that AI is a tool with serious limitations, not an

authority to be trusted uncritically. None of this is quick. All of it requires institutional investment that most universities have not yet committed to.

Academic Integrity and Ethical Considerations:

AI has complicated an already difficult conversation about academic integrity. Students can now generate essays, reports, and assignment responses with substantial AI assistance. Some uses are legitimate and pedagogically defensible – organising ideas, refining language, exploring different framings of a problem. Others amount to submitting AI-generated work as one's own, with little intellectual engagement and no acknowledgement of AI's role (OECD, 2024; UNESCO, 2023). The line between these categories is not always clear, which is part of what makes the conversation difficult.

The institutional response should not be primarily technological. AI-detection tools exist, but they are unreliable. They generate false positives. They create an adversarial relationship between students and institutions that undermines the educational environment. A better approach – and this is where the literature is increasingly converging – is to redesign assessment. Tasks that require critical analysis, disciplinary reasoning, personal reflection, and the application of knowledge to specific, contextualised problems are inherently more resistant to AI misuse than tasks that ask students to reproduce information or produce generic text (Gallent-Torres et al., 2023). This is not a new insight; good assessment design has always prioritised higher-order thinking. AI simply makes the case more urgent.

Ethical considerations extend well beyond plagiarism. Data privacy – what happens to student data entered into commercial AI platforms. Algorithmic bias – whose perspectives are overrepresented, whose are invisible. Transparency – whether students and teachers understand how AI systems work and what their limitations are. Fairness – whether the deployment of AI within an institution advantages some students at the expense of others. These are questions that require clear, publicly stated institutional policies, developed through genuine consultation and revised as the technology and its uses evolve. The goal is a culture of responsible use, not a culture of surveillance (OECD, 2023).

Theoretical Perspectives on Generative Artificial Intelligence and Educational Equity:

The patterns described above – unequal access, uneven benefits, structural barriers – are not new problems created by AI. They reflect longstanding dynamics in how educational systems distribute opportunity, reward certain kinds of knowledge, and respond (or fail to respond) to social inequality. Understanding why AI benefits some students and not others require theoretical tools that go beyond technology itself. Four thinkers, working from quite different traditions, are particularly illuminating: Pierre Bourdieu, Basil Bernstein, Paulo Freire, and Shiv Visvanathan.

Pierre Bourdieu: Cultural Capital and Educational Inequality:

Bourdieu's concept of cultural capital (1986) explains, with uncomfortable directness, why providing everyone with the same tool does not produce the same outcome. Students enter university carrying unequal quantities and kinds of cultural knowledge – familiarity with academic norms, comfort with particular communication styles, exposure to digital technologies, confidence in navigating institutional environments. Educational systems

tend to reward what already-advantaged students already possess. The system does not create equality; it processes and certifies pre-existing inequality.

Apply this framework to AI and the implications are immediate. A student who grew up with digital devices at home, who reads and writes fluently in English, who has been encouraged to question information sources — that student will extract far more educational value from generative AI than a student lacking those forms of capital. The tool is identical. The educational outcome is not. Bourdieu would find this entirely predictable. The educational system is doing what it has always done: converting unequal starting positions into unequal results while appearing to offer equal opportunity.

Basil Bernstein: Language, Communication, and Knowledge:

Bernstein (1971) narrows the focus to something specific and, for the Indian context, particularly consequential: the role of language in how knowledge is structured, communicated, and understood. His argument is that the linguistic codes students bring to education — shaped by their social and cultural backgrounds — determine how they engage with learning, participate in academic activities, and demonstrate competence.

In a country where higher education operates primarily in English but draws students from dozens of language backgrounds, this framework has obvious relevance. When AI generates content in English, it operates within a communicative register that some students command fluently and others do not. When it generates content in Indian languages, the quality often deteriorates — not just linguistically but conceptually, because the disciplinary vocabulary and contextual nuances that matter for genuine understanding may be poorly represented. Bernstein's point is that access to information is not the same as access to knowledge. Knowledge depends on how information is communicated, and AI systems that treat language as a simple translation variable misunderstand what actually happens when students learn.

Paulo Freire: Critical Pedagogy and Human Agency:

Freire (1970/2005) insisted — the word is appropriate, because he was unyielding on this — that education is not the transfer of pre-formed knowledge from teacher to student. It is a dialogical process involving questioning, interpretation, disagreement, and the active construction of understanding. He called the alternative the “banking model” of education: teacher deposits information, student stores it, nobody thinks.

Generative AI, when used uncritically, fits the banking model with unsettling precision. Student enters question. AI returns answer. Student accepts it. There is no dialogue. No productive struggle with ambiguity. No moment where the student's own reasoning is tested or refined. Freire would argue that this diminishes the educational process irrespective of whether the AI's answer happens to be accurate. Education's value lies in the process of thinking, not in the product of having been given an answer. Used differently, though — as a provocation, a starting point for critical discussion, a source of claims that students are then required to evaluate — AI could support exactly the kind of inquiry Freire championed. But that outcome requires deliberate pedagogical design. It does not happen automatically, and it will not happen at all if teachers are not prepared for it.

Shiv Visvanathan: Knowledge Diversity and Social Context:

Visvanathan (2009) raises a fundamentally different concern. His work challenges the assumption – widespread in technology discourse – that scientific and technological knowledge is neutral, universal, and adequate on its own. Different communities, he argues, produce knowledge that is shaped by their specific histories, cultures, and circumstances. Treating any single knowledge system as comprehensive is both intellectually impoverishing and, in a country as diverse as India, politically dangerous.

This critique applies to AI with considerable force. Generative AI systems are trained on massive datasets that disproportionately represent certain languages, certain academic traditions, and certain cultural perspectives. What they produce is not neutral knowledge – it is knowledge shaped by the biases and gaps in their training data. For Indian higher education, this has concrete implications. AI-generated content may systematically underrepresent regional knowledge traditions, local educational practices, indigenous perspectives, and the kinds of contextual understanding that are essential for meaningful education in a diverse society. Visvanathan's perspective requires educators to treat AI as one knowledge source among many – useful, but partial, and always in need of critical supplementation.

These four thinkers share something despite their different starting points: a refusal to treat educational inequality as a technical problem with a technical solution. Bourdieu reveals the unequal resources students carry with them. Bernstein shows that the medium of instruction is never neutral. Freire insists that education must remain dialogical or it ceases to be education. Visvanathan challenges the adequacy of any single knowledge system, AI training data included. Collectively, they suggest that the work of integrating AI equitably into Indian higher education is not primarily technological. It is pedagogical, institutional, and – there is no way around this word – political.

Implications for Policy and Practice:

If the preceding analysis holds – that AI's educational value is determined by conditions, not capabilities – then the policy implications follow logically, even if implementing them does not. Five areas demand coordinated attention.

Strengthening Digital Infrastructure:

This is the most fundamental requirement and, in practical terms, perhaps the hardest to fulfil. Without reliable internet, functional devices, and adequate institutional technology, generative AI is irrelevant to a student's educational experience. The infrastructure gap between well-resourced and under-resourced institutions is well documented, and it is not closing fast enough.

Investment must be targeted rather than generalised. National connectivity announcements mean little if they do not reach the specific institutions and communities that currently lack adequate infrastructure. Colleges in rural areas, institutions serving economically disadvantaged populations, universities enrolling large numbers of first-generation learners – these should receive priority. The argument here is not specific to AI. It is about the digital foundation that any form of technology-supported learning requires.

Professional Development for Teachers:

Whether AI helps or harms educational quality depends on teachers. The technology does not teach. Teachers do, and their professional judgement determines whether AI-generated material enters the classroom in a form that supports learning or undermines it.

Professional development cannot be a one-day workshop on prompt engineering. It needs to be sustained, discipline-embedded, and pedagogically oriented. Teachers need to evaluate AI outputs for accuracy, design learning tasks that incorporate AI without reducing intellectual demand, adapt AI-generated materials for local contexts, and guide students towards responsible use. This requires institutional commitment and – something rarely mentioned – time, which is already desperately scarce for most educators in Indian higher education.

Promoting AI Literacy Among Students:

Students need to be taught, explicitly and repeatedly, that AI is a tool with significant limitations. Most have not received this instruction. They need the capacity to evaluate AI-generated information, verify factual claims, identify when outputs are biased or incomplete, and use AI in ways consistent with academic integrity.

Bolting a standalone AI literacy module onto an already crowded curriculum is unlikely to work. AI literacy is most effective when woven into disciplinary teaching – when students learn to assess AI's usefulness and shortcomings within the specific intellectual context of their field. The objective is not to make students afraid of AI. It is to make them discerning.

Supporting Multilingual and Inclusive Education:

India's linguistic diversity is a structural fact that any serious approach to AI in education must reckon with rather than work around. Policy should encourage the development and improvement of multilingual AI tools, while honestly acknowledging their current limitations. AI-generated content in many Indian languages is not yet dependable enough to serve as a primary educational resource without teacher mediation and contextual review.

Inclusion means more than language, of course. It means recognising that students enter university from very different starting points – different educational histories, different degrees of digital familiarity, different relationships to the institutional culture of higher education. AI can contribute to inclusion when it is deployed with genuine attention to these differences. Deployed on the assumption that one tool serves all students equally, it will reproduce exclusion under a different name.

Developing Ethical Institutional Frameworks:

Universities need clear, publicly stated policies on AI use – and most do not currently have them. Students and teachers should know what is expected: which forms of AI use are acceptable in academic work, how AI assistance should be acknowledged, what the institution regards as a breach of academic integrity. These policies should be developed consultatively, not imposed from above, and they must be revisited as the technology evolves.

Surveillance is not the answer. AI-detection software is unreliable, and a regime built around policing student behaviour creates precisely the wrong educational culture. What works better – and what the evidence supports (OECD, 2023) – is building a culture of

responsible use through awareness, guidance, open discussion, and assessment practices that test genuine understanding rather than the ability to generate polished text.

These recommendations are interdependent. Infrastructure investment without teacher preparation accomplishes nothing. Ethical policies without AI literacy programmes remain aspirational. Multilingual support without attention to how knowledge is communicated across linguistic contexts misses the point. The challenge is coordinated, sustained action across all five areas simultaneously – and that demands a level of institutional seriousness that goes well beyond purchasing subscriptions to AI platforms. What generative AI ultimately contributes to Indian higher education will be determined not by the technology's capabilities but by the commitment of institutions and policy makers to addressing the inequalities that shape how it is used.

Conclusion:

The question running through this article is not whether generative AI belongs in Indian higher education – it is already embedded in the daily practices of students and, increasingly, of teachers. The question is whether its presence will deepen the inequalities that have long characterised the system or whether, through sustained and deliberate effort, it can be made to work against them.

The analysis presented here points towards cautious realism. Generative AI offers genuine benefits – personalised learning support, broader access to academic resources, multilingual assistance at a scale that was previously unavailable. These are not trivial gains, and dismissing them would be intellectually dishonest. But they are gains that accrue unevenly, shaped by the same factors – infrastructure, language, institutional capacity, teacher readiness – that have always determined who benefits from educational innovation and who gets left behind. A student who uses AI to generate an assignment without engaging with the material has not learned anything. A teacher who adopts AI-produced lesson plans without contextualising them for local conditions has not improved their teaching. Technology without intentionality is noise.

Bourdieu, Bernstein, Freire, and Visvanathan – approaching from very different intellectual traditions – converge on a shared insistence: educational outcomes cannot be understood apart from the social conditions that produce them. Bourdieu exposes how cultural and digital capital determines who profits from ostensibly equal opportunities. Bernstein demonstrates that language is never a neutral vehicle for knowledge. Freire refuses any educational model that reduces learners to passive consumers. Visvanathan demands recognition that the knowledge embedded in large-scale AI systems is neither comprehensive nor representative. These are not marginal academic concerns. They are central to understanding why a technology that serves some students well may simultaneously fail others – and why treating AI as a universally beneficial innovation is, at best, naive.

The practical demands are clear enough: invest in infrastructure, develop teachers' professional capacity, build students' AI literacy, strengthen multilingual AI capability, establish transparent ethical governance. These are not items on a checklist to be completed and forgotten. They constitute the substance of whether generative AI becomes a force for educational equity or another mechanism through which existing advantage compounds itself. The decisions being made now – by policy makers, by university administrators, by individual teachers – will shape the trajectory of AI in Indian higher

education for years to come. The window for getting this right is open, but it will not remain so indefinitely.

References:

1. Bernstein, B. (1971). *Class, codes and control: Vol. 1. Theoretical studies towards a sociology of language*. Routledge & Kegan Paul.
2. Bourdieu, P. (1986). The forms of capital. In J. G. Richardson (Ed.), *Handbook of theory and research for the sociology of education* (pp. 241–258). Greenwood Press.
3. Freire, P. (2005). *Pedagogy of the oppressed* (30th anniversary ed., M. B. Ramos, Trans.). Continuum. (Original work published 1970)
4. Gallent-Torres, C., Zapata-González, A., & Ortego-Hernando, J. L. (2023). The impact of generative artificial intelligence in higher education: A focus on ethics and academic integrity. *RELIEVE*, 29(2), Article M5. <https://doi.org/10.30827/relieve.v29i2.29134>
5. Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günemann, S., Hüllermeier, E., Kruber, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Saber, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, Article 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
6. Ministry of Education. (2020). *National education policy 2020*. Government of India. https://www.education.gov.in/sites/upload_files/mhrd/files/NEP_Final_English_0.pdf
7. Ministry of Education. (2024). *All India Survey on Higher Education (AISHE) 2021–2022: Final report*. Government of India. <https://aishe.gov.in/document/aishe-final-report-2021-22/>
8. Ng, D. T. K., Leung, J. K. L., Su, J., Ng, R. C. W., & Chu, S. K. W. (2024). Design and validation of the AI literacy questionnaire: The affective, behavioural, cognitive and ethical approach. *British Journal of Educational Technology*, 55(3), 1082–1104. <https://doi.org/10.1111/bjet.13411>
9. Organisation for Economic Co-operation and Development. (2023). *OECD digital education outlook 2023: Towards an effective digital education ecosystem*. OECD Publishing. <https://doi.org/10.1787/c74f03de-en>
10. Organisation for Economic Co-operation and Development. (2024). *Education policy outlook 2024: Shaping responsive and resilient education systems*. OECD Publishing. <https://doi.org/10.1787/dd5140e4-en>
11. Sullivan, M., Kelly, A., & McLaughlan, P. (2023). ChatGPT in higher education: Considerations for academic integrity and student learning. *Journal of Applied Learning & Teaching*, 6(1), 31–40. <https://doi.org/10.37074/jalt.2023.6.1.17>
12. UNESCO. (2021). *AI and education: Guidance for policy-makers*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf00000376709>
13. UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf00000386693>
14. Visvanathan, S. (2009). The search for cognitive justice. *Seminar*, 597. https://www.india-seminar.com/2009/597/597_shiv_visvanathan.htm