**Beyond access: Four dimensions of students' relationship with generative AI, and their stratification in higher education in India**

Brigid Freeman[a] and Nidhi Sabharwal[b1]

[a] Australia India Institute, University of Melbourne, Melbourne, Australia; [b] National Institute of Educational Planning and Administration, New Delhi, India

**Abstract**

India has the second largest higher education system in the world, and generative artificial intelligence is now part of how its students study. Students appear to have first reached generative AI through the free tiers of consumer applications on their own smartphones, and institutional and regulatory policy is developing alongside this use. The emerging research response, in India and internationally, has concentrated on *adoption*, *perception* and *academic integrity*, measured predominantly through surveys with demographic covariates. This article reviews that field alongside research on inequalities in access to, use of and benefits from digital technology and research on equity in higher education in India, and argues that both the existing empirical literature on students and generative AI and the research on digital access and use capture only one dimension of a relationship that has several. It introduces a framework of four dimensions of students' relationship with generative AI: *material* (what AI a student can reach, how securely, and what their own effort is now worth), *affective* (the full spectrum of feeling, distributed across social position), *epistemic* (whether students can see into the systems that judge them, whether they are seen fairly by those systems, and whether their language and knowledge are legible to the AIs at all), and *temporal* (aspiration under acceleration, and the devalued loop of practice). Each dimension is given substantive conceptual treatment: definition, theoretical inheritance, extension claims, assembled evidence and its specific unanswered question in India. The article takes bounded account of other higher education actors, leaders and academics, as shapers of the student relationship, and closes with a discussion drawing the apparatus together as the foreground of a planned empirical study of students in higher education in India.

**Keywords**: generative artificial intelligence, digital divide, higher education, India, stratification, equity, caste, student experience

## 1. Introduction

The higher education system in India is the largest in the world by institutional count and, with 45 million students enrolled in 2023-24, the second largest by enrolment after China, which enrolled 48.5 million in 2024 (Ministry of Education, 2026; Ministry of Education of the People's Republic of China, 2025). The system includes more than 1,250 universities and 48,000 colleges (Ministry of Education, 2026), spanning centrally funded institutions, state universities and a large number of affiliated colleges, which enrol most students. Expansion of the system has been accompanied by growing diversity in the student body. The Scheduled Castes (SCs), Scheduled Tribes (STs) and Other Backward Classes (OBCs) are constitutionally recognised groups that have historically faced social disadvantage, and the National Education Policy 2020 (NEP 2020) includes them among the socially and economically disadvantaged groups, or SEDGs (Ministry of Human Resource Development, 2020). Students from these groups now make up a majority of enrolment (Sabharwal & Malish, 2018;

[1] **CONTACT** Brigid Freeman: brigid.freeman@unimelb.edu.au & Nidhi Sabharwal: nidhis@niepa.ac.in
The views expressed in this paper are those of the authors and should not necessarily be attributed to the organisations where they are employed.

Varghese, 2022; Varghese et al., 2019). This expansion carries consequences at national, state and local levels. India's demographic dividend, the young population whose potential NEP 2020 explicitly commits itself to realising (Ministry of Human Resource Development, 2020), must pass substantially through this system. For many of these students, first-generation entrants among them, a degree represents a significant family investment, and its conversion into occupational outcomes varies by social position and field of study (Tierney et al., 2019; Varghese et al., 2018). Generative AI has arrived in this setting. How any consequential new technology distributes its benefits and burdens across this system is therefore an equity question as well as a technology question, and it bears on the wider participation that expansion has brought.

Generative artificial intelligence arrived into this system in a manner that is notable analytically. Where learning management systems arrived through top-down institutional mandates and state-sanctioned provision (Choudhury & Kumar, 2024), generative AI typically arrived from below, through free tiers of consumer applications reaching students directly on their smartphones, ahead of institutional policy and carried mainly by mobile data (Telecom Regulatory Authority of India, 2026). The arrival is at once vast and uneven. Industry survey data count 958 million active internet users in 2025, 57 per cent rural and growing at nearly four times the urban rate (Internet and Mobile Association of India [IAMAI] & Kantar, 2026). The proportion accessing the internet for e-learning is low (3 per cent), according to this report. The same survey reports AI-enabled feature use within the past year by 44 per cent of all users and 57 per cent of those aged 15 to 24. Furthermore, 18 per cent of users in 2025 reach the internet on someone else's device, nearly 80 per cent of them rural (IAMAI & Kantar, 2026). Connection statistics also leave out factors that shape usable access, including reliable electricity, network quality and a data budget that extends to study, which vary across states and households. For the tertiary-age cohort, contact with AI-enabled features has already happened at scale; the conditions of that contact differ by position. What has not happened, in India or anywhere, is research adequate to the question this conjunction forces: How is students' relationship with generative AI structured by social and institutional position?

This article prepares the ground for that question. It states the review's method, the framework's origin and status, and the article's mode of production (Section 2). It reviews the relevant fields, namely the higher education system in India, its equity scholarship, and its COVID-era digital baseline (Section 3), the empirical literature on students and generative AI in India and internationally (Section 4), and the extension to AI of research on access to, use of and benefits from digital technology (Section 5). It provides the first conceptual work on a framework developed to explore this question, giving substantive treatment to four dimensions of a student's relationship with generative AI: *material*, *affective*, *epistemic* and *temporal* (Section 6, this article's centrepiece). The closing discussion draws the four dimensions together as one apparatus, states what remains unknown, and positions the empirical study this review foregrounds (Section 7).

## 2. Methodology

This review assembled its evidence through structured literature scans across seven strands: the genealogy of research on digital access and use; the Indian equity and stratification corpus; empirical research on students and generative AI in India and internationally; decolonial, postdigital and posthuman scholarship; Global South comparators; theoretical anchors for the material, affective, epistemic and temporal dimensions; and methodological literature. The literature review identified peer-reviewed scholarship, grey literature and preprints. All references were checked against their original sources and verification-graded before use, with preprints and grey literature carried only at the evidential weight appropriate to their status.

The four dimensions at the article's centre were first codified by one of the authors within a prior program concerned with the human experience of AI emergence. That program involved multi-model AI-assisted literature scans, synthesis and conceptual development undertaken for a separate project. The present article develops one subset of the terrain revealed through that earlier work and warrants it directly through the literatures reviewed here, including literatures on digital access and use; risk, anxiety and social amplification literatures; opacity, datafication, decolonial AI and recognition scholarship; and work on aspiration, deskilling and cognitive offloading. The dimensions therefore have a sensitising status, in Blumer's (1954) sense, in that they provide "directions along which to look" (p. 7) rather than prescriptions of what the article must find. They organise the review's conceptual argument without being treated as fixed categories, validated constructs or empirical findings. This stance sits within postdigital scholarship's refusal of a clean separation between digital and human educational life, and the posthumanist scholarship on which that literature draws (Jandrić et al., 2018). Through this lens, students' relationships with generative AI are treated as part of contemporary study itself rather than as accessory uses of a tool.

This article was prepared with generative AI assistance, reported here as part of the method in line with editorial guidance from the Committee on Publication Ethics (COPE) and the AI-use policies now common across major publishers, which require disclosure of the tools used, the tasks those tools performed and the authors' responsibility for the result. Anthropic's Claude supported literature scanning, synthesis and drafting under author direction. Other systems, including OpenAI's ChatGPT, Google's Gemini and NotebookLM, contributed to earlier stages of the broader conceptual program from which the framework derives, as described above. No AI system meets authorship criteria and none is credited as an author. All sources were retrieved and verified against their original records by the authors, every citation was checked against a maintained verification log before inclusion, and all interpretive, methodological and editorial judgements are the authors' own. The authors accept full responsibility for the accuracy and integrity of the article. The role of AI systems in both the prior conceptual development and the preparation of this article is disclosed as a substantive feature of the method rather than an incidental influence, a stance consistent with the article's own argument that AI entanglement in academic work should be visible and examinable.

**3. The system and its stratifications**

Higher education in India has expanded considerably since the early 2000s. Enrolment increased from 8.8 million in 2001-02 (Varghese & Sabharwal, 2022) to 45.0 million in 2023-24, and the gross enrolment ratio (for the 18-23 age cohort) rose from 23.7 per cent in 2014-15 to 30.0 per cent in 2023-24 (Ministry of Education, 2026). Favourable educational policies and affirmative action have contributed to this growth and to a more diverse student body, as an expanding system is more inclusive than elite higher education. Education for All has produced more secondary school graduates, and reservation, relaxed admission criteria and financial support have encouraged access for disadvantaged groups (Varghese et al., 2019). Affirmative action policy has widened access across all segments of the population. Students from the Scheduled Castes, Scheduled Tribes and Other Backward Classes now make up a majority of those enrolled. In 2023-24 they together accounted for 62.0 per cent of total enrolment, with Scheduled Castes at 15.5 per cent, Scheduled Tribes at 6.4 per cent and Other Backward Classes at 40.1 per cent (Ministry of Education, 2026).

Research on equity in higher education in India shows that students' experience of the system varies with their social position. This scholarship treats access as having several dimensions: who enrols, which institutions and fields students enter, and what students experience once they are there. Caste, class, gender, language and region each shape these dimensions (Borooah & Sabharwal, 2021; Sabharwal, 2021; Varghese et al., 2019). Access has widened (Ministry of Education, 2026), and the experience of access varies by social group, institution type and field of study (Sabharwal, 2021).

Varghese et al. (2019) reported that growth proceeded at different rates across states, that regional differences in gross enrolment ratios had grown, that differences in access by socio-economic position continued, and that students from more advantaged backgrounds were over-represented in selective institutions and in fields such as science and engineering.

Within institutions, mixed-methods case studies across six states have examined students' experiences. They reported differences by social group in interaction with teachers and peers, including reports of unequal access and treatment of students from socially and economically disadvantaged groups (SEDGs) (Sabharwal & Malish, 2017, 2018). Institutional mechanisms, from remedial provision to grievance structures, mediate whether diversity becomes inclusion, and the affirmative action architecture is itself an object of contested student attitudes that institutions must manage (Sabharwal, 2024). Parallel work maps the axes along which experience varies: caste and class jointly (Thorat & Newman, 2010; Tierney et al., 2019), gender and its intersections with caste in academic life (Sabharwal et al., 2020; Varghese & Sabharwal, 2024), the nature and quality of access as well as its extent (Sabharwal, 2021), and English as medium of instruction, which functions in India simultaneously as pedagogy, gatekeeper and class marker (Borooah & Sabharwal, 2021). This scholarship also draws on a methodological approach to studying social exclusion, including its conceptualisation, indicators and measurement (Thorat & Sabharwal, 2015). Any study of how a new technology distributes its benefits and burdens across students in India inherits this map of positions. This article therefore places the proposed study within that equity landscape of higher education, and it takes this landscape as its starting point for asking how generative AI reaches students across the system.

Policy frames the moment in several directions at once, within a system undergoing considerable change. The National Education Policy 2020's implementation is redrawing curricula, credit architecture, institutional categories and the focus on equity simultaneously (Ministry of Human Resource Development, 2020). NEP 2020 positions adaptive technologies as core to educational modernisation and, through its Indian Knowledge Systems agenda, aspires to carry Indian epistemologies into technological development. This agenda is now in active implementation, with the University Grants Commission (UGC) issuing dedicated guidelines for the training and orientation of faculty in Indian Knowledge Systems (UGC, 2023). The state is also building AI infrastructure on both fronts this article will theorise. The IndiaAI Mission includes national compute capacity and open dataset platforms (Ministry of Electronics and Information Technology, 2024), the Bhashini National Language Translation Mission is developing hundreds of AI models across more than thirty-six Indian languages precisely to close the language legibility gap (Bhashini, 2026), and official communications catalogue AI's insertion across curricula and platforms (Press Information Bureau, 2026a). The provision question, meanwhile, is already under consideration. The national technical education regulator has acknowledged the sustainability of student AI access beyond expiring free pilot licences as a live policy question (The Secretariat, 2025). Commentary within India also draws attention to the importance of infrastructural equity in AI integration (Niazi, 2025; Thakur, 2026). All of these directions appear, at present, empirically untested against student experience.

There is a documented precedent for how a digitally mediated disruption distributes itself across this system, and it is recent. The COVID-19 emergency put higher education in India online and made the distribution of digital capacity suddenly consequential. Choudhury and Kumar (2024), analysing nationally representative household data alongside student survey evidence, found that only around 40 per cent of college-going students had home access to digital devices and the internet, with variation by caste and class. This study also reported that over 80 per cent of surveyed students expressed dissatisfaction with online learning. Where the emergency transition in industrialised systems could presume laptops and home broadband, India's context disproportionately involved smartphones and mobile data, under the constraints that mobile-first access carries in India: shared

devices, limited data, uneven network quality and unreliable electricity, constraints that connection statistics do not register (IAMAI & Kantar, 2025; Telecom Regulatory Authority of India, 2026). The same constrained mobile-first infrastructure would later carry generative AI's free tiers. Choudhury and Kumar's (2024) central analytical move, from inequality of access to inequity of experience, is the move any serious study of generative AI and Indian students must repeat.

The COVID literature also contributed a warning directly relevant to this article's framework. Observers cautioned that the no-cost educational technology offered to Indian students during the emergency should not be mistaken for durable provision. Elsewhere, such offers had proved short-lived, such that a digital-by-default system risked leaving low-income families unable to afford access once the free period ended (Gilbertson et al., 2021). Ethnographic work on low-income families' use of proprietary platforms has since examined how these arrangements play out in practice (Bhatia et al., 2025). Free access that can be withdrawn, that warning implies, is a different object from continuous, guaranteed access. Subsequent sections build that distinction into the framework's material dimension as precarity.

The COVID baseline gives this review three points to carry forward: differences in students' digital experience in higher education in India are documented and recent; the factors associated with those differences are the ones that research on equity in higher education maps; and the system has already experienced one episode in which a technology introduced as emergency provision was followed by differences in students' experience. Generative AI is a second such episode, arriving more quickly, closer to academic work itself and with less institutional mediation.

**4. The field as it stands: adoption, perception, integrity**

The empirical literature on students and generative AI has grown at remarkable speed since the release of ChatGPT by OpenAI in November 2022. Its characteristic shape is consistent enough, in India and internationally, to be described as a type: single-institution or single-program surveys measuring adoption, attitudes, perceived usefulness and integrity concerns, frequently framed through technology acceptance models (TAM, UTAUT and their extensions),[2] with responses compared across demographic covariates such as gender, age, discipline and level of study. Having said that, there are emerging cross-country studies, such as the Digital Education Council's (2024) Global AI Student Survey. Their survey found that 86 per cent of 3,839 students across sixteen countries used AI in their studies and 72 per cent expected their universities to provide AI training. Similarly, Borzanović et al.'s (2026) cross-regional study of 788 university students in Europe and Australia found that large majorities valued generative AI for rapid information access (72.5 per cent) and personalised learning support (66.4 per cent). However, their findings revealed that while students acknowledged ethical risks such as data privacy, algorithmic bias and overreliance, these concerns rarely deterred actual adoption, reflecting a pragmatic trade-off of risk for academic productivity gains.

The Indian instances are representative. Menon and Shilpa (2023) offered a UTAUT-framed qualitative account of ChatGPT adoption, and identified performance and effort expectancy, social influence and privacy concerns. Raman et al. (2024) analysed university students' ChatGPT adoption through Diffusion of Innovations (DOI) theory examining relative advantage, compatibility and

[2] The Technology Acceptance Model (TAM; Davis, 1989) evaluates adoption through perceived usefulness and perceived ease of use. The Unified Theory of Acceptance and Use of Technology (UTAUT; Venkatesh et al., 2003) unifies prior adoption frameworks around four core constructs: performance expectancy, effort expectancy, social influence and facilitating conditions. Extended applications in generative AI research frequently incorporate domain-specific variables such as data privacy concerns, trust, academic integrity anxieties and perceived interactivity (e.g., Menon & Shilpa, 2023; Strzelecki, 2024).

complexity/ease of use, while single-institution perception surveys reported broadly positive attitudes alongside integrity anxieties. The pattern remains consistent at systematic-review level. Mateen et al. (2025), reviewing generative AI in Indian medical education through Preferred Reporting Items for Systematic reviews and Meta-Analyses (PRISMA) methods, screened 19,777 records to find eleven eligible Indian studies. The evidence those studies supply primarily concerns awareness, perception and readiness.

These studies provide an important empirical baseline, with the most nuanced among them revealing tensions that extend well beyond standard technology acceptance. Baek et al. (2024), surveying 1,001 United States college students, found gender, age and major influencing ChatGPT use, but so did institution type and institutional policy. Higher-income students viewed AI more positively than lower-income students, indicating that institutional and economic gradients are already visible. Ravšelj et al. (2025), spanning a large dataset of 23,218 students across 109 countries including India, found perceptions patterned by the country's economic context while gender and urban-rural location showed only minor differences. This is the signature one would expect if position, not demography as such, were doing the structuring. A United Kingdom national survey (J. Freeman, 2025) found use increasing from 66 to 92 per cent in a year. Women were reportedly more deterred than men by accusation risk and by fear of false results. Wealthier students were more enthusiastic and better able to buy premium access, and only 36 per cent received institutional skills support. The report concluded that the divide was widening. Comparative Global South research includes a study of student perceptions in programming education across Sub-Saharan Africa framed explicitly around a pronounced digital divide (Oyelere & Aruleba, 2025), as well as a validated Ghanaian instrument for student ChatGPT experience (Baidoo-Anu et al., 2024). The institutional layer has been documented in analyses of university policies and guidelines (Crompton et al., 2026; Jin et al., 2025; LaFrance, 2026). Concerns regarding academic integrity, assessment and cognitive offloading are pervasive, while policy lags practice (Chan & Hu, 2023; Lodge & Loble, 2026; Maleki, 2026).

The limits are equally consistent. The unit of analysis is frequently the single institution, which structurally prevents the question of how institutional position shapes experience from arising. The analytical treatment of social difference is frequently the demographic covariate. Gender, age and discipline enter as grouping variables for attitude scores, not typically as positions within a stratified system whose structure the analysis theorises. The object measured is typically perception and adoption rather than the texture of the relationship itself: what students can reach and how securely, what they feel across the full range of feeling, what they can contest, what they imagine. The covariate approach, moreover, appears to underdetermine the phenomenon. Reviews of demographic differences report inconsistent findings across contexts, with gender effects appearing, disappearing and reversing between samples (Maxwell et al., 2025), while one of the field's largest single datasets found economic context patterning perceptions where gender barely did (Ravšelj et al., 2025). Such inconsistency is what one would expect if shallow demographic categories were proxying, unevenly, for underlying positional structures the studies do not measure (e.g., access security, language background, institutional resources, prior digital formation). That possibility cannot be tested from inside the adoption-perception paradigm. It requires a framework built for stratification, which is what research on access to, use of and benefits from digital technology, reviewed next, begins to supply and what subsequent sections extend.

A final feature of the field must be registered, because the study this article grounds will make bounded use of it. Students are not the only higher education actors whose relationship with generative AI is being documented and theorised. Academics have their own stratified relationship with the technology, differentiated by their understandings of what the university is for (Schoon et al., 2026), by precarity and position in the prestige economy (Watermeyer et al., 2025) and by self-assessed AI proficiency (Digital Education Council, 2025). Institutions act through policy, while

leadership-facing surveys report students' expectations of institutions directly (Digital Education Council, 2026; Stephenson & Armstrong, 2026). In the context of this article, these actors are important for one key reason: their decisions, anxieties and policies constitute much of the environment in which the student relationship is formed, provisioned, policed and signalled. The framework developed next therefore treats system and institutional leaders, academics and professional staff strictly as shapers of the student relationship, informing its institutional-shaping question. The object of study is, and remains throughout, the student.

**5. Digital access and use, and their extension to AI**

Research on inequalities in access to, use of and benefits from digital technology, which began as work on the digital divide, is the one body of research that has consistently treated technology's distribution as a question about social structure rather than individual attitude. In its developed form, it operates as a multi-level framework spanning physical access (first level), skills and usage patterns (second level) and tangible offline outcomes (third level). In relation to the first level, van Dijk (2005, 2020) positions motivational and physical access as the foundational moments of technological engagement. As physical access spreads, inequality shifts to the second level, where Hargittai (2002) established the skills divide among the already connected, while van Deursen and van Dijk (2013) showed usage gaps deepening across different user groups. Moving to the third level, van Deursen and Helsper (2015) formalised how equivalent connection yields differential tangible offline outcomes. Scheerder et al. (2017), reviewing determinants across the second and third levels, reported education as the most consistently linked determinant in the literature, showing that the divide runs through, not merely alongside, educational stratification. Robinson et al. (2015) stated what is at stake in general form, arguing that digital inequality deserves standing alongside race, class and gender because it increasingly shapes life chances. The literature's central argument is that as each lower level saturates, inequality migrates upward rather than disappearing. Its normative point is that digital inequality reinforces social inequality (Ragnedda, 2017; van Dijk, 2020).

Application of this apparatus to artificial intelligence began before the generative moment, with Carter et al.'s (2020) review carrying the three-level structure into AI contexts (an early application, though it coined no settled term). Kitsara (2022) framed an AI divide from an innovation perspective, while Lythreatis et al. (2022) systematised the wider literature as generative AI arrived. International agencies now frame AI's distribution as a global divide question (International Labour Organization & United Nations, 2024) and measure generative AI exposure in developing economies with adoption-adjusted methods (Gmyrek et al., 2024). Furthermore, an interdisciplinary consensus statement acknowledged that generative AI can either exacerbate or ameliorate socioeconomic inequality, with policy decisive between the two (Capraro et al., 2024).

Empirical mapping has followed. Surveying 1,088 Dutch adults, Wang et al. (2025) derived five latent user classes and found the most vulnerable were older, less educated and less skilled, mapped explicitly onto the divide's levels. In higher education, the closest existing study to the one this article grounds is Beckman et al. (2025). This mixed-methods analysis of 194 Australian university students, grounding generative AI inequality explicitly in research on digital access and use, identified novice, cautious and enthusiastic student profiles, and warned that the critical diffusion stage was unfolding in a policy vacuum in which advantage compounds. Two features of that study particularly resonate. Its own analysis of the institutional context was that a risk-oriented policy discourse acts differentially. Students already high in digital and AI literacy leveraged their capabilities despite restrictive policy, while cautious and novice students absorbed the risk discourse and held back. In such circumstances, the institution's protective posture itself compounds disadvantage. As Beckman et al. (2025) did not collect demographic data, they were unable to connect student profiles to social position even in principle.

The Indian telephone data suggest one instructive complication. With 958 million active users in 2025, rural India supplying 57 per cent of them, and 38 per cent of the population still offline but shrinking annually (IAMAI & Kantar, 2026), connectivity for the college-age cohort has broadened dramatically, and AI-feature contact has followed, now reaching a majority of 15 to 24-year-olds. The previous edition of the same survey reported women at 47 per cent of users, with 98 per cent of all users accessing the internet in Indian languages (IAMAI & Kantar, 2025). The same edition listed the unavailability of local-language content among the barriers keeping the remaining offline population offline, behind unawareness of the internet and lack of access. The complication is that the levels do not saturate neatly in sequence here, as official statistics still record lower rural internet subscriptions when compared with urban India (Telecom Regulatory Authority of India, 2026), even as higher-level gaps in device sharing, skills, tier, language and outcomes open above them. India displays the divide's levels simultaneously rather than sequentially. An analysis of Indian students must be built multi-level from the start, because a binary access measure applied to this cohort would record high connectivity among enrolled students and miss differences at the higher levels.

Yet research on digital access and use, even in its AI extension, has a boundary this review must acknowledge. Its object is resources: connection, skills, usage and the tangible outcomes they yield. But a student's relationship with generative AI goes beyond the question of resources. The same technology that is unevenly reachable is also unevenly feared and loved. It judges students through detection and assessment systems they cannot see into; it carries, or fails to carry, their language and their knowledge; and it reprices the futures their degrees are based on. These sit beyond outcome domains of access in van Deursen and Helsper's (2015) sense. They are additional dimensions of the relationship itself, each with its own theoretical inheritance and each plausibly structured by the same positions mapped by research on digital access and use and by research on equity in higher education in India. Research on digital access and use, extended to AI, captures the material dimension of a relationship that has at least three more. Characterising all four, and what is known and unknown about each in India, is the work of the next section, which is this article's centrepiece.

## 6. Four dimensions of students' relationship with generative AI

A student's relationship with generative AI has multiple dimensions, of which this article develops four: *material*, *affective*, *epistemic* and *temporal*. The four are not offered as exhaustive, as a relationship this consequential surely has further dimensions, *relational* and *ethical* among the candidates, that other literature explores. The claim made here is more specific, proposing that these four are foundational for the stratification question, each draws on established theory, and each is plausibly structured by position. Every student occupies a location in all four at once, and the organising proposal for the framework this article posits is that social and institutional position shapes all four. What follows gives each dimension the same treatment, including definition and scope, theoretical inheritance, the extension claims this framework adds, the assembled evidence across international and Indian sources, and the dimension's specific unanswered question in India. The identical treatment should not be read as identical evidential weight. The material dimension rests on population data, technical benchmarks and a documented pandemic precedent, while the affective, epistemic and temporal dimensions rest more heavily on theory and international evidence, an asymmetry the unanswered questions recorded under each dimension make explicit. A fifth subsection assembles the cross-cutting apparatus, including the position axis, language's three functions, Indian knowledge's two functions and the integrative conjecture. Table 1 presents the apparatus whole.

Table 1. *The four dimensions of students' relationship with generative AI, and the cross-cutting apparatus*

| Dimension | What it concerns | Theoretical inheritance | What the framework this article introduces adds | Unknown in India |
|---|---|---|---|---|
| Material: access as a power relation | What AI a student can reach and how securely; tier of access; devices, data, sharing; skills; benefit; the worth of one's own effort | Three-level model of access, use and benefits (van Dijk; Hargittai; van Deursen & Helsper); AI extension (Carter et al.; Beckman et al.) | Access as power relation, not possession; precarity as distinct from absence; language as access condition; the worth of effort (Danaher & Nyholm) | Tier, security and precarity of student access; differential benefit; prompting competence by medium of instruction |
| Affective: feeling, distributed | The full spectrum of feeling in studying with AI, from obsolescence anxiety, detection anxiety and shame to relief, comfort, attachment and joy | Cultural theory of risk (Douglas & Wildavsky); social amplification (Kasperson et al.) | Affect as positionally distributed; the positive spectrum as data; institutions as amplification stations | Which affects concentrate at which positions |
| Epistemic: knowing, being known, being legible | Seeing into the systems that judge; being rendered into data; the legibility of one's language, name, knowledge | Opacity (Burrell); black box (Pasquale); datafication (Williamson et al.); decolonial AI (Mohamed et al.; Birhane; Nyaaba et al.); recognition (Honneth) | Two directions held as one condition; the testimonial edge (who gets believed); legibility and representation as epistemic | Contestability, the experience of being flagged, legibility of language and Indian content |
| Temporal: aspiration under acceleration, and the devalued loop of practice | Imagined futures under a force outside one's control; the destabilised promise of slow, effortful practice | Capacity to aspire (Appadurai; Ball et al.); deskilling and offloading (Fan et al.; Lee et al.) | Aspiration under acceleration; the devalued practice loop; openness to counter-frames regarding the meaning of effort | Imagined futures by position; the changed relationship to practice |
| Cross-cutting: position | The explanatory axis spanning institution type and tier, class, gender, first-generation status, medium of instruction and language background, rurality and social group, with caste in standard administrative categories | Measurement approaches to studying social exclusion in India (Thorat & Sabharwal, 2015) and to student diversity in higher education (Sabharwal & Malish, 2018); intersectionality (Crenshaw, 1991) | Axes treated as intersecting positions whose joint effects may exceed their parts, in place of independent covariates | Joint effects of position on all four dimensions |

| Dimension | What it concerns | Theoretical inheritance | What the framework this article introduces adds | Unknown in India |
|---|---|---|---|---|
| Cross-cutting: language | Three functions kept apart in measurement, including the access condition (models perform worse in Indian languages), legibility (whether the student's words and context are readable to the systems) and stratifier (medium of instruction as class-marked sorting) | Performance benchmarks (Singh et al., 2024); English-medium stratification (Borooah & Sabharwal, 2021) | Language as constitutive of the relationship in each dimension, with code-mixed practice as a possible bypass | Which function dominates for which students; whether code-mixing recovers capability |
| Cross-cutting: Indian knowledge | Two functions, including representation within the epistemic dimension (whether Indian knowledge traditions are present and given standing in what the AIs return) and analytical resource within the temporal dimension | Decolonial AI (Mohamed et al., 2020; Birhane, 2023); NEP 2020 Indian Knowledge Systems agenda | Knowledge representation as an equity question inside higher education rather than a purely cultural one | Whether students perceive their knowledge traditions in AI outputs, and with what standing |
| Cross-cutting: compounding | The integrative conjecture, held loosely, including precarious access, concentrated anxiety, low contestability and constrained futures converging on the same students, each step of the candidate mechanism individually testable | Detector bias as the demonstrated material-epistemic instance (Liang et al., 2023) | A five-step candidate mechanism from precarity through anxiety, suppressed use and thinned standing to devalued effort | Whether the dimensions compound at the same positions |

### 6.1 The material dimension: access as a power relation

The *material* dimension concerns resources defined more widely than access. It encompasses what AI a student can reach and how securely while distinguishing institutionally provided, privately paid and free-tier access as distinct material conditions. It treats the precarity and revocability of access as analytically distinct from its absence, covering device, data and device-sharing constraints under which use happens. It further includes skills of use, specifically prompting effectively in English as a stratified competence in a multilingual system. Finally, it measures actual derived benefit, asking what a student's own effort and developing competence are worth in a world where a machine produces polished output on demand.

Its theoretical inheritance is the three-level model of access, use and benefits, reflecting van Dijk's (2005, 2020) sequential access model, Hargittai's (2002) second-level skills divide, van Deursen and

Helsper's (2015) third-level outcomes, Scheerder et al.'s (2017) determinants review, Carter et al.'s (2020) early carriage of the structure into AI, and the emerging generative AI divide literature culminating, for higher education, in Beckman et al. (2025). The framework reflects this scholarship and then extends it with three claims concerning *access*, the *worth of effort* and *language*.

*Access* is treated as a power relation rather than a possession. What a free-tier student holds is not simply a smaller quantity of the same good held by an institutionally provisioned student. Instead, it represents a standing relation of dependence on a provider who owes the user nothing and can change terms without notice. As the COVID experience in India already demonstrated for educational technology, such free offerings often function as market entry rather than durable provision. Republican political theory names this condition domination, conceived as subjection to another's arbitrary capacity to interfere, whether or not interference occurs (Pettit, 1997). The framework posits that precarity and revocability must be measured as *properties of access* in their own right, invisible to any measure that records only what a student currently reaches.

The *worth of effort* belongs to the material dimension. Danaher and Nyholm (2021) argue that automation threatens an achievement gap, understood as a decline in the availability of the goods of achievement, the meaning that attaches to outcomes because one's own effortful agency produced them. Translated to students, the material dimension includes both what AIs one can reach and what one's own developing competence now feels like it is worth. This question has important stakes where families have invested heavily in competence converting into a future. This layer connects forward to the temporal dimension's devalued loop of practice.

*Language* is also an access condition. An AI that performs well only in English is materially less accessible to a Hindi-, Tamil-, or Telugu-medium student regardless of their device and data situation. Furthermore, prompting effectively in English is itself a competence distributed by the same schooling stratification that Borooah and Sabharwal (2021) documented for English-medium instruction generally. Private unaided schooling is largely English-medium, and the medium of schooling is associated with the range of subjects students go on to study in higher education (Varghese et al., 2019). English-medium schooling is associated with advantage across South Asia (Rahman, 2005), and in India the medium of schooling already differs by social group before any AI enters the picture.

The claim now has a quantified technical base, though its strongest current sources appear as preprints pending peer review.[3] Peer-reviewed studies confirm the underlying pattern. On the largest Indic generation benchmark, spanning cross-lingual summarisation, translation, and question answering, even the best-performing frontier AI model showed a significant performance gap relative to English in every one of the twenty-nine languages evaluated (Singh et al., 2024). Two features of this evidence are important beyond the headline gaps. The proprietary-to-open gap matters distributionally, because free tiers and open models are what less-resourced students disproportionately reach, so the language penalty and the tier penalty compound in the same students. The code-mixing finding points the other way. If mixed-language input recovers most of the lost performance, the code-mixing fluency common among Indian students may function as a vernacular skill that partially bypasses the penalty, an empirical question about practice that no study reviewed has asked.

[3] Benchmark evaluations report an average 8.3-point English-to-Hindi performance gap across tasks in frontier AI models (Chigrupaatii et al., 2025); best-model accuracy of only 38 to 45 per cent across nine Indic languages on a demanding multi-task benchmark, with open-source models far lower (Sankalp KJ et al., 2025); and native-language knowledge-recall accuracy gaps relative to English of up to approximately 0.50 across nine open-weight models, on a benchmark covering eighteen scheduled languages. For the eleven language pairs with code-mixed variants verified by native speakers, code-mixed inputs substantially reduced this gap, bringing performance within approximately 0.05 of English without model-level intervention (Mazumder et al., 2026).

The assembled *provision*-side evidence completes the picture. Institutional provision decisions, not only consumer purchasing, enact the divide. Which universities license which AIs for which members is a stratifying act, and in India the national technical education regulator has identified the sustainability of student access beyond expiring free pilot licences as a policy question (The Secretariat, 2025). Furthermore, the population-level infrastructure beneath any provision varies across the country (IAMAI & Kantar, 2026; Telecom Regulatory Authority of India, 2026). Data are cheap yet finite against study-scale use, and connection itself is contingent on electricity and network quality that the statistics do not measure. The COVID baseline demonstrated how quickly device and connectivity constraints translate into stratified educational experience (Choudhury & Kumar, 2024).

What is unknown in research in India is nearly everything above the first level. It appears that few if any Indian studies measure the tier, security or precarity of students' generative AI access. It also appears that few if any measure differential benefit, connect prompting competence to medium of instruction, or ask what students feel their effort is now worth. The material dimension's Indian evidence base consists of population infrastructure data, technical benchmarks and a COVID precedent, with the student layer not yet studied, as far as the searches show. A substantial share of that infrastructure picture necessarily rests on government and industry sources rather than peer-reviewed research, and this reflects the state of the field. Higher education research in India is a developing field with particular strengths, including the equity scholarship this article builds on, and the scale of the system leaves ample scope for further study. In these conditions, and for policy scholarship generally, the appropriate discipline is to draw on official statistics and policy instruments whose provenance and methods are documented, and to weigh them by triangulation. The figures relied on here agree in direction across a statutory regulator's administrative statistics, an industry survey with a published methodology and the academic COVID literature, which is the strongest confirmation currently available.

**6.2 The affective dimension: feeling, distributed**

The *affective* dimension concerns feeling, and its defining commitment is breadth, spanning the full spectrum of emotion in living and studying with generative AI. That spectrum includes fear of obsolescence and replacement, anxiety about detection and about being judged, shame at error and concern about potential dependence or cognitive offloading. It equally includes relief, comfort, gratitude, attachment and joy. For many students, generative AI is the first time that explanation, practice, feedback and translation have been available on demand, at any hour, in their own mixture of languages. The affect that meets a capacity which was literally unavailable only a few years earlier is not adequately captured by any anxiety scale. What the dimension claims is distributional, focusing on where each feeling lives rather than how much anxiety students feel on average, which is what existing instruments measure. The dimension examines which students, at which social and institutional positions, feel threatened, which feel relieved, which feel joy, which feel several of these at once and what each feeling then does to their study. The framework expects, for instance, detection anxiety to concentrate where institutional signals are most contradictory and standing to contest is lowest, and relief and joy to concentrate where the AI substitutes for support that position never supplied. Both expectations are open to refutation.

The foundational literature concerns the sociology of risk. Douglas and Wildavsky (1982) established that the selection of dangers is social. What a group fears tracks its position, institutions and commitments, not hazard magnitude alone. Kasperson et al. (1988) supplied the transmission mechanism. Risk signals amplify and attenuate as they pass through social stations, institutions prominent among them, so that the same hazard arrives at different social locations at very different intensities, a finding whose continuing relevance its authors have recently reviewed (Kasperson et al.,

2022). Together these anchor the dimension's two structural expectations, namely that affect should track position rather than exposure alone, and institutions that emit contradictory signals, promoting AI in one breath and policing it in the next, should function as amplification stations, concentrating anxiety among students most dependent on institutional cues about what is safe. The term expectation is deliberate. Consistent with the framework's sensitising status, these are directions along which to look, licensed by the inherited theory, rather than hypotheses the study is committed to confirming. The contradictory signals are documented fact. For example, fewer than a third of United Kingdom students said their institution encouraged AI use while nearly a third reported a ban (J. Freeman, 2025), and the Australian evidence showed the resulting risk discourse acting differentially, held off by confident users and absorbed by cautious and novice ones (Beckman et al., 2025).

The framework this article posits extends this scholarship in two ways. Affect is treated as positionally distributed, which converts the psychology of AI anxiety into a stratification question, where the object of study is a distribution across the position axis rather than a population mean. The framework also extends the affective spectrum to its positive end. A first-generation student with no family member to ask may experience a patient, non-judgemental, endlessly available assistant as relief and even attachment. Similarly, an English-insecure student may experience drafting help as anxiety reduced rather than integrity threatened. A fear-only instrument would code these students as untroubled or miss them entirely.

The assembled evidence shows a measurement field that is real but partial. Validated instruments exist. These include a general AI anxiety scale (Wang & Wang, 2022), a brief AI attitude scale (Grassini, 2023), and, closer to this framework's needs, two recent Global South validations. The Ethical AI Dilemma Anxiety Scale, validated on two samples of Egyptian university students totalling 1,530 respondents, resolves into three factors (academic integrity anxiety, professional future anxiety and societal impact anxiety) (Helali et al., 2026). An academic-research AI anxiety scale validated in Saudi Arabia resolves into four factors including fear of committing plagiarism and fear of losing research skills (Alshaibani et al., 2025). These factor structures are independently useful to this framework, because they show distinct anxieties loading separately. Most notably, integrity anxiety sits on the epistemic dimension's detection edge, professional future anxiety is the temporal dimension's future half rendered psychometric, and skill-loss fear adjoins the devalued practice loop. Experimental evidence on the positive end is beginning to accumulate. Bai (2026) compared university students working on creative problem-solving tasks with generative AI, with an internet search engine or without technology. The generative AI condition lowered anxiety and shame while raising both the correctness and the novelty of solutions, with metacognitive regulation moderating the effect on novelty. That evidence, however, registers the positive end only as the lessening of anxiety and shame. No existing instrument appears to measure relief, comfort, gratitude, attachment or joy in their own right, and none is India-validated. Both absences define instrument-development work rather than settled measurement.

Early signals that *affect* is positionally patterned already exist at covariate level, and they point the way while falling short of the question. Women in the United Kingdom survey were more deterred than men by accusation risk and by fear of false output (J. Freeman, 2025), and higher-income United States students held more positive attitudes than lower-income students (Baek et al., 2024). A further signal concerns the institutional shaping of affect. Among 410 students across eight departments of a Turkish university, Üstün et al. (2026) found employment anxiety to be the most pronounced of three anxiety dimensions. The association between AI use and anxiety became statistically insignificant once departmental differences were controlled, a result the authors interpret as showing that emotional responses to AI reflect the institutional contexts in which technological change is encountered rather than individual engagement with the technology. The United Kingdom and

United States findings concern single demographic variables in general-purpose surveys, and the Turkish study concerns disciplinary context within one institution. None analyses how affect distributes across a structured position space, and none appears to exist for India. What is unknown in India is therefore the distribution itself, as no identified study anywhere, and none in India, maps whether AI-related fear, or relief, concentrates at particular social positions among higher education students. The *affective* dimension's distributional question is timely.

### 6.3 The epistemic dimension: knowing, being known and being legible

The *epistemic* dimension concerns knowing and being known, and it extends in two directions. Outward, it asks whether students can see into, question and contest the algorithmic systems that increasingly judge them, including AI-text detection, automated proctoring and automated assessment. Inward, it examines whether and how those systems see students, who are rendered into data, scores and flags with little standing to contest the rendering. Underlying both directions are questions of legibility and representation, regarding whether a student's language, name and context are legible to the AIs at all, and whether Indian knowledge, the very content many Indian students study, is present, accurate and given standing in what the AIs return.

The scholarly literature is layered. Burrell (2016) distinguishes three forms of machine learning opacity, as corporate secrecy, as technical illiteracy and as the intrinsic opacity of high-dimensional models, establishing that the last form does not dissolve with disclosure. Pasquale (2015) supplies the political economy of the black box and the principle with the sharpest educational consequence, that judgements that cannot be explained cannot be contested. Turned inward, the dimension inherits the critique of datafication in higher education, whereby students become objects of continuous measurement whose data doubles carry consequences their originals cannot audit (Williamson et al., 2020), the older sociology of reactive measurement, in which institutional metrics remake what they claim to describe (Espeland & Sauder, 2007), and the documented record of automated administrative systems bearing hardest on the disadvantaged. Eubanks (2018) showed digitised eligibility and risk systems continuing the moralistic surveillance and punishment of the poor, and Couldry and Mejias (2019) frame the wholesale rendering of human life into extractable data as a colonial continuity. That algorithmic judgement errs along social lines is a matter of record (Angwin et al., 2016). Selbst et al. (2019) caution that fairness cannot be secured by abstracting a system from the social context it operates in, which is precisely the abstraction an institutional detector performs on a multilingual student body. In terms of decolonial AI, Mohamed et al. (2020) frame algorithmic systems as sites where colonial patterns of knowledge extraction and imposition recur. Similarly, Birhane (2023) argues that algorithmic colonisation "occurs not through brute physical force, but rather through invisible and nuanced mechanisms, for example, control of digital ecosystems and infrastructure... [and] assumes that human personality, behaviour, and action is raw material" (p. 247). Nyaaba et al. (2026) reported generative models defaulting to Western epistemologies, marginalising non-dominant languages and knowledge, in education specifically. Finally, and integratively, the *epistemic* dimension is read through Honneth (1995), for whom recognition is a moral infrastructure, and its denial, misrecognition, is an injury with institutional sites. The claim that places this study within research on equity in higher education in India is that the algorithmic layer is a new institutional site of recognition and misrecognition, added to the campus sites that scholarship has already documented, including those in Indian academic life traced by Sabharwal et al. (2020). A detector that wrongly flags a student's own words, or a model to which a student's knowledge tradition is invisible, is misrecognition in Honneth's sense, delivered by infrastructure.

The framework introduced in this article extends this literature in three ways. The two directions are held together, as opacity of the systems to the student and exposure of the student to the systems constitute one condition, and their meeting point, the moment a student meets an AI-mediated

judgement they cannot audit, is the dimension's political centre. The sharpest equity edge of that centre is testimonial. Where an AI detector is used, the question becomes who is believed, and whose word counts against a machine's. The answer plausibly tracks the same positions that determine whose word counts in institutional life generally. Legibility and representation belong to the epistemic dimension as well as the material one. A model that cannot handle a student's language or renders their knowledge tradition thinly does not merely serve them less. It fails to know them and their world, a limitation in an education system whose national policy (Ministry of Human Resource Development, 2020) has made Indian Knowledge Systems an explicit agenda. Whether the dominant AI models carry Indian knowledge at all is a live policy question with no apparent empirical answer.

The assembled evidence is dominated by one mechanically demonstrative finding. Liang et al. (2023) tested seven widely used GPT detectors against non-native English writers' essays and found over 61 per cent falsely flagged as machine-generated, some unanimously, against near-perfect accuracy on native-speaker essays. The underlying mechanism is perplexity, whereby constrained linguistic expression reads as machine-like. The sting lies in the second half, where prompting ChatGPT to enrich the essays' word choice cut false positives below 12 per cent, so the escape route from false accusation itself requires AI access and skill. The students most exposed to false accusation are those with the least native-like English and the least secure access, who also have the least means of defence. Here, the material and epistemic dimensions compound in a single mechanism, a point Section 6.5 returns to. Detector unreliability is independently established (Weber-Wulff et al., 2023). The Indian institutional regime gives these findings their local force. Global guidance exists but is soft, as UNESCO's (2023) first international framework for generative AI in education urges human-centred use, capacity building and validated systems, and binds no one. National policy on AI in higher education is set within broader frameworks. The India AI Governance Guidelines set out principles that apply across sectors (Press Information Bureau, 2026b), and NEP 2020 encourages the introduction of AI as a subject of learning across stages of education (Ministry of Human Resource Development, 2020). The UGC's plagiarism regulations (UGC, 2018), as applied in practice to unacknowledged AI-generated content, form the current framework. Practice varies across institutions. Many universities across India now run research submissions through AI content detection, a form of scrutiny that did not apply to earlier cohorts, and practitioner guidance written for research students documents both the spread of this practice and differences in how it is applied (The Editorial, 2026). Detection AIs are licensed products, so the checks a student is subject to, from mandatory screening to none, vary with institutional type and resources. On representation, the first large-scale India-centred multimodal benchmark indicated performance degrading for culturally specific Indian content and low-resource languages (Faraz et al., 2025, preprint), initial technical support for the thin-canon concern.

What is unknown in India is the experiential layer entirely. No study was identified that measures Indian students' perceived understanding of the AI systems that judge them, their perceived capacity to contest AI-mediated judgements, their experience of being seen and flagged, whether believed or disbelieved, or the perceived legibility of their language, names and knowledge to the AIs they use daily.

### 6.4 The temporal dimension: aspiration under acceleration, and the devalued practice loop

The *temporal* dimension concerns time, in two halves. The future half concerns aspiration under acceleration, regarding how students imagine degrees, careers and selves under a technological acceleration they do not control, which continuously reprices what their credentials will be worth. The present half concerns what this article calls the *devalued practice loop*. Education's implicit temporal promise is that repeated, effortful practice transforms the person, and that promise is destabilised when a machine already produces the output the practice was supposed to yield, leaving

students to ask why they should endure the slow loop at all. Its object is provisionality, where futures are experienced as contingent on a force outside one's control.

The future half's theoretical anchor is Appadurai's capacity to aspire (2004, 2013). Aspiration is a cultural and navigational capacity rather than an individual psychological trait, formed in social life, consisting substantially in maps, the learned connections between immediate choices and distant possibilities, and distributed unequally. Aspiration's skewed distribution is a constitutive feature of poverty rather than an incidental one. The anchor earns its place three times over. It converts aspiration under acceleration from metaphor into theory. According to this framing, if aspiring is navigating by maps of pathways, an acceleration that scrambles the maps, determining which credentials will hold value and which occupations will persist, should injure precisely those whose maps were thinnest, anticipating the stratification the framework expects rather than merely permitting it. Aspiration relates to the terms of recognition, aligning the temporal dimension to the epistemic dimension. Furthermore, it is Indian-grounded scholarship, developed through Appadurai's Mumbai work. The higher education application is established. Ball et al. (2002) demonstrated advantaged youth moving steadily toward university while disadvantaged youth navigated the same passage with doubt, deliberation and thin information, a finding that the account of aspiration under acceleration extends rather than replaces.

The present half reflects the deskilling and cognitive offloading literatures as they reach generative AI. Experimental evidence associates generative AI use with what Fan et al. (2024) name 'metacognitive laziness'. In their randomised comparison, learners working with ChatGPT improved their essay scores but showed no corresponding gain in knowledge or transfer, having offloaded the regulation of their own learning to the AI. Among 319 knowledge workers reporting 936 real instances of generative AI use, greater confidence in the AI predicted less self-reported critical-thinking effort, while greater confidence in one's own ability predicted more (Lee et al., 2025). This study observed trust in the machine and trust in one's own formation pulling in opposite directions. Our reading of this literature is temporal rather than merely cognitive. What is at stake is students' relationship to practice, struggle and slow formation, the plateau phases of learning whose value is invisible at the time and vindicated only afterwards, and whose perceived necessity the machine's instant output undermines. Qualitative evidence from a comparably stratified system already registers this. South African academics interviewed about generative AI observed that the "seeming 'instant perfection' of AI generated work" deters students from working through the slow work of forming their own identities (Schoon et al., 2026, p. 248), illustrating the devalued practice loop described from the teaching perspective. The same concern has now been raised from within India. Writing in the Economic and Political Weekly, Mathews (2026) draws on the cognitive debt findings of Kosmyna et al. (2025) to argue that routine reliance on generative AI in student writing shifts authorship and agency toward the machine while leaving judgement with the writer. The resulting homogenisation of voice, on this account, carries particular costs in a multilingual setting, notwithstanding the easier access to English that AI-assisted writing gives many learners.

The framework introduced in this article extends this literature in three ways. The future and present halves belong together as one temporal dimension. The future that acceleration renders provisional and the present whose practice loop is devalued are the same temporality seen from opposite perspectives, and instruments should carry both. The dimension is proposed as varying by position. A stable imagined future may be least secure, and the mobility calculation most exposed, where the equity scholarship locates disadvantage. Early labour-market evidence, discussed below, suggests the destination end of the mobility question may already vary by caste, so the framework asks whether the student-side experience of the future is graded the same way. Furthermore, the dimension is open to counter-frames students may themselves bring. The available theory of aspiration presumes a

returns-to-credential subject, whereas Indian intellectual traditions offer various lenses where the meaning of effort does not depend on its returns.

The evidence for the temporal dimension is the thinnest of the four, which is itself a finding. Professional future anxiety exists as a validated psychometric factor (Helali et al., 2026), and fear of losing skills as another (Alshaibani et al., 2025), giving both aspects of the dimension measurable analogues. The Indian equity scholarship documents occupational mobility and its obstacles for students from disadvantaged groups (Sabharwal & Malish, 2018; Tierney et al., 2019), and one recent analysis supplies the destination backdrop. Mapping three independent occupational AI-exposure indices onto India's redesigned Periodic Labour Force Survey, Mishra's (2026, preprint) study of roughly 83,000 employed graduates reported that Scheduled Caste and Scheduled Tribe graduates had lower measured exposure to generative AI than non SC/ST/OBC category graduates within the same district, by between 0.24 and 0.37 standard deviations, and Other Backward Classes graduates by about 0.14 standard deviations. The study attributed the difference to occupational distribution, noting that about a quarter of Scheduled Caste and a third of Scheduled Tribe graduates worked in farm or elementary occupations, and those in white-collar work were less represented in managerial, software and finance roles. Because exposure commanded a wage premium of up to 20 per cent in the study's estimates, the pattern bears on the occupational destinations toward which students' mobility aspirations are directed. The Scheduled Caste and Other Backward Classes gaps were widest among labour-market entrants aged 15 to 30 and narrowed across older cohorts, so they are largest among the graduates now entering employment, while the Scheduled Tribe gap was similar across cohorts (Mishra, 2026, preprint). The destination end of the mobility question therefore appears, on the evidence of the preprint, to vary by caste in the terrain this technology reprices, and no study documents whether the formation and imagined futures of students still inside the system vary in the same way. What is unknown in India is the entire student side, how students at different positions imagine their futures alongside AI, how their relationship to practice and slow learning has changed, and what frames they use to relate effort to outcome. It does not appear that any study asks these important questions.

### 6.5 The cross-cutting apparatus: position, language, Indian knowledge and compounding

Four *elements* cut across the dimensions and complete the framework.

*Position* provides the explanatory axis. The dimensions describe the relationship. Position explains its variation, and their interaction is the forthcoming study. Position comprises institution type and tier, class, gender, first-generation status, medium of instruction and language background, rurality and social group. Social group is studied with the concepts, indicators and measures developed in research on social exclusion (Thorat & Sabharwal, 2015) and in research on the experiences of students from socially disadvantaged groups in higher education (Sabharwal & Malish, 2018). These axes are treated as intersecting positions whose joint effects may exceed their parts (Crenshaw, 1991).

*Language* performs three functions, and these must be kept apart in measurement despite compounding in life. It is an access condition within the material dimension, as AIs perform worse in Indian languages, so language background conditions what capability a student actually reaches. It is a legibility question within the epistemic dimension, regarding whether the student's own words, name and context are readable to the AIs and their adjacent detectors, with Liang et al. (2023) demonstrating the stakes of illegibility. Further, it is a stratifier within the position axis: medium of instruction is a class-marked sorting mechanism in Indian education (Borooah & Sabharwal, 2021). Instruments must ask about each job separately, including language mixing and not merely language choice, since code-mixed practice may largely bypass the performance penalty (Mazumder et al., 2026, preprint). Analysis is then free to trace how the three jobs compound. The mixing finding points

to something larger the instruments should stay open to, namely a repertoire of improvisational workarounds, familiar in India as *jugaad*, the art of making constrained resources do the work of absent ones, through which students may be actively converting shared devices, free tiers and mixed-language prompting into capability that research on digital access and use would not predict from their measured access.

*Indian knowledge* performs at least two functions. Within the *epistemic* dimension it is a representation question, examining whether the knowledge traditions many students study are present and given standing in what the AIs return, a question with national policy weight (Ministry of Human Resource Development, 2020) and initial technical evidence of thin coverage (Faraz et al., 2025, preprint). Within the *temporal* dimension it is an analytical resource, potentially available to students as sense-making equipment.

The preliminary integrative conjecture, held loosely, is that the four dimensions compound into one condition. The framework this article posits conjectures that precarious access, concentrated anxiety, low contestability and constrained imagined futures converge on the same students, particularly those from socially and economically disadvantaged groups (SEDGs). It further conjectures that their convergence is mutually reinforcing. A candidate mechanism can be sketched. Precarious access raises the stakes of every use, and raised stakes feed anxiety. Anxiety absorbed as institutional risk discourse suppresses use among the students least equipped to contest an accusation. Suppressed use slows the growth of skill and of the standing to question what the systems return. Thinned skill and narrowed standing then shorten the futures that appear reachable, and futures that appear constrained devalue the effortful loop through which skill would have grown. Each step is individually testable, which is what converts the conjecture from an image into a research design. One instance of the conjectured mechanism is already demonstrated. Liang et al. (2023) found that the students most likely to be falsely accused were those whose English was least native-like, and the demonstrated escape route, AI-assisted polishing, requires exactly the access and skill that systems have not yet extended to the same students. That is material-epistemic compounding shown mechanically, in a single published experiment. Whether the full four-way convergence occurs in Indian students' lives is an important integrative question.

## 7. Discussion and conclusion

Three things follow from this review. The *field* comes first. The literature on students and generative AI is large, fast and productive, and it is misshaped for the most consequential question before it. Its strongest studies keep discovering the edges of the paradigm, with income and institution surfacing as predictors in covariate designs, economic context patterning perceptions where demography barely does, a national survey watching its own measured divide widen year on year. The field's problem is direction rather than quality. The four-dimensional framework developed in this article is offered as the reorientation, centred on one relationship, of which four dimensions are developed here, each located within the scholarly literature, each proposed as structured by position and each carrying a specific, currently unanswered question. Indian empirical work on students and generative AI exists in quantity, multi-site perception studies have appeared, and cross-national surveys include Indian students. The searches identified no research that treats students' relationship with generative AI as multi-dimensional, beyond adoption or attitude. They identified no studies that analyse that relationship in terms of social stratification, with social position as the unit of analysis, and none engages socially and economically disadvantaged groups (SEDGs). None builds on Indian research on who enrols in higher education, which institutions and fields students enter, and what students experience once they are there. The gap exists in the international literature as well. India is a suitable setting in which to address it. Its higher education system is the second largest in the world, its student body is socially diverse, and it has a substantial body of research on equity in higher

education to draw on. One system elsewhere is asking the right question: South African research poses generative AI explicitly as an equity problem for a highly unequal, multilingual system with a colonial-language instruction legacy, asking whether AI can enable or will further obstruct epistemological access (Schoon et al., 2026), with epistemological access (Morrow, 2009) the South African example of the *epistemic* dimension. However, there too, the project interviews academics, while the student layer is missing.

The second concerns the *setting.* Higher education in India brings together several conditions that bear on how generative AI reaches students. These are a student body in which many students belong to the socially and economically disadvantaged groups, following expansion of the system and the reach of affirmative action; a mobility calculation resting on credentials whose destination labour market appears to vary by caste in its exposure to AI; a language ecology in which the dominant AIs perform less well in most students' languages while the state builds supporting infrastructure; a detection regime that varies with institutional capacity; and a divide whose levels are live simultaneously and not in sequence. These conditions bear on every dimension of the framework, and institutions are the main site where they meet students.

*Institutions* come third, and here the bounded account of other actors earns its place. On the assembled evidence, institutions shape every dimension of the student relationship. They provision or withhold AIs, and the sustainability of free provision is an open policy problem at the national level. They emit the signals, frequently contradictory, that amplify or dampen anxiety differentially across their student bodies. They operate the detection and assessment systems that constitute the epistemic dimension's sharpest edge, and their assessment and skills responses will substantially decide what devaluation does to students' relationship with practice. The practical implication is symmetrical with the analytical one: if the four dimensions compound, institutional action on any one of them is partial, and equity-serious institutional strategy must be built across all four.

This review set out to examine the fields bearing on students' relationship with generative AI in higher education in India, and to do the first conceptual work on a framework adequate to that relationship. Its conclusion is deliberately an opening. The framework now stands in the literature with its inheritances named, its extension claims stated, its evidence assembled and its unknowns specified. It is sensitising, and it is developed to be tested, revised or broken by Indian students' own accounts. That test is the purpose of the empirical study this review foregrounds, Students, Generative AI and Equity in Higher Education, a mixed-methods study of Indian students, with a fieldwork design drawn from research on equity in higher education in India. Students are the primary informants across all four dimensions, with institutional leaders and academics informing the institutional-shaping question. The study's ambition is continuous with the research it extends, to ask how the equity aims of one of the world's largest higher education systems are shaped by its newest technology, and to ask it with the students whose experience the answer concerns.

**References**


Alshaibani, M. H., Al-Rahmi, W. M., & Tawafak, R. M. (2025). Psychometric validation of the artificial intelligence anxiety scale: A confirmatory factor analysis for academic research. *Contemporary Educational Technology, 17*(4), ep595. https://doi.org/10.30935/cedtech/17309

Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016, May 23). Machine bias. *ProPublica*.

Appadurai, A. (2004). The capacity to aspire: Culture and the terms of recognition. In V. Rao & M. Walton (Eds.), *Culture and public action* (pp. 59-84). Stanford University Press.

Appadurai, A. (2013). *The future as cultural fact: Essays on the global condition*. Verso.

Baek, C., Tate, T., & Warschauer, M. (2024). "ChatGPT seems too good to be true": College students' use and perceptions of generative AI. *Computers and Education: Artificial Intelligence*, 7, Article 100294. https://doi.org/10.1016/j.caeai.2024.100294

Bai, Y.-Q. (2026). An experimental study on the impact of generative AI on university students' emotions and performance in creative problem-solving tasks. *Learning and Instruction*, 102, Article 102316. https://doi.org/10.1016/j.learninstruc.2026.102316

Baidoo-Anu, D., Asamoah, D., Amoako, I., & Mahama, I. (2024). Exploring student perspectives on generative artificial intelligence in higher education learning. *Discover Education*, 3(1), Article 98. https://doi.org/10.1007/s44217-024-00173-z

Ball, S. J., Davies, J., David, M., & Reay, D. (2002). 'Classification' and 'judgement': Social class and the 'cognitive structures' of choice of higher education. *British Journal of Sociology of Education, 23*(1), 51-72. https://doi.org/10.1080/01425690120102854

Bhashini. (2026). *National Language Translation Mission*. Digital India Bhashini Division, Government of India.

Bhatia, K. V., Arora, P., & Gupta, S. (2025). Edtech platforms from below: A family ethnography of marginalized communities and their digital learning post-pandemic. *Learning, Media and Technology, 50*(4), 482–496. https://doi.org/10.1080/17439884.2024.2328693

Beckman, K., Apps, T., Howard, S. K., Rogerson, C., Rogerson, A., & Tondeur, J. (2025). The GenAI divide among university students: A call for action. *The Internet and Higher Education, 67*, 101036.

Birhane, A. (2023). Algorithmic colonization of Africa. In S. Cave & K. Dihal (Eds.), *Imagining AI*. Oxford University Press.

Blumer, H. (1954). What is wrong with social theory? *American Sociological Review, 19*(1), 3-10.

Borooah, V., & Sabharwal, N. S. (2021). English as a medium of instruction in Indian education: Inequality of access to educational opportunities. In M. Arsel, A. Dasgupta, & S. Storm (Eds.), *Reclaiming development studies* (pp. 197-226). Anthem Press. https://doi.org/10.2307/j.ctv22d4tj9.16

Borzanović, A., Simić, A., Papić-Blagojević, N., Klašnja-Milićević, A., Sanin, C., Levula, A., Islam, M. R., & Ivanović, M. (2026). Student perceptions of generative AI tools in higher education: A cross-regional study of use, ethics, and educational utility. *Education and Information Technologies, 31*(12), 4327-4367. https://doi.org/10.1007/s10639-026-13964-8

Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society, 3*(1). https://doi.org/10.1177/2053951715622512

Capraro, V., Lentsch, A., Acemoglu, D., Akgun, S., Akhmedova, A., et al. (2024). The impact of generative artificial intelligence on socioeconomic inequalities and policy making. *PNAS Nexus*, 3(6), Article pgae191. https://doi.org/10.1093/pnasnexus/pgae191

Carter, L., Liu, D., & Cantrell, C. (2020). Exploring the intersection of the digital divide and artificial intelligence: A hermeneutic literature review. *AIS Transactions on Human-Computer Interaction, 12*(4), 253-275. https://doi.org/10.17705/1thci.00138

Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20(1), Article 43. https://doi.org/10.1186/s41239-023-00411-8

Chigrupaatii, R., Kanishka, P. S. T., Routhu, L. C., Patel, M., Reddy, S. S., Gupta, D., Srikar, D., Kuchimanchi, K. T., Misra, R., & Tripathi, R. (2025). *HinTel-AlignBench: A framework and benchmark for Hindi-Telugu with English-aligned samples* [preprint]. arXiv 2511.15183. https://arxiv.org/abs/2511.15183

Choudhury, P. K., & Kumar, A. (2024). COVID-19 and digital learning in higher education in India: Unpacking the inequity in students' experiences. In S. Bhushan (Ed.), *The evolving landscape of higher education in India*. Springer. https://doi.org/10.1007/978-981-97-9270-2_9

Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press.

Crenshaw, K. (1991). Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review, 43*(6), 1241-1299.

Crompton, H., Burke, D., Nickel, C., Bozkurt, A., Miao, F., Sharples, M., Greene, J. A., Parsons, D., Gill-Simmen, L., Edmett, A., Pegrum, M., de Waard, I., Bonk, C. J., Garcia, M. B., Curry, J. H., Lindsey, L., Yang, M., Marshall, S., Bali, M., … Yu, S. (2026). Governing generative AI in higher education: A global Delphi study on policy and practice. *International Journal of Educational Technology in Higher Education, 23*(1), Article 21. https://doi.org/10.1186/s41239-026-00602-z

Danaher, J., & Nyholm, S. (2021). Automation, work and the achievement gap. *AI and Ethics, 1*, 227-237. https://doi.org/10.1007/s43681-020-00028-x

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, *13*(3), 319–340. https://doi.org/10.2307/249008

Digital Education Council. (2024). *Digital Education Council global AI student survey 2024*. https://www.digitaleducationcouncil.com/resource-library-items/digital-education-council-global-ai-student-survey-2024

Digital Education Council. (2025). *Digital Education Council global AI faculty survey 2025.* https://www.digitaleducationcouncil.com/post/digital-education-council-global-ai-faculty-survey

Digital Education Council. (2026). *AI in higher education global survey 2026*. https://www.digitaleducationcouncil.com/resource-library-items/ai-in-higher-education-global-survey-2026

Douglas, M., & Wildavsky, A. (1982). *Risk and culture: An essay on the selection of technological and environmental dangers*. University of California Press.

Espeland, W. N., & Sauder, M. (2007). Rankings and reactivity: How public measures recreate social worlds. *American Journal of Sociology, 113*(1), 1-40.

Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor* (1st ed.). St. Martin's Press.

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2024). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2), 489-530. https://doi.org/10.1111/bjet.13544

Faraz, A., Akash, Khan, S., Kolla, R., Patidar, A., Goswami, S., Ravi, A., Khatri, C., & Agarwal, S. (2025). *IndicVisionBench: Benchmarking cultural and multilingual understanding in VLMs* [preprint]. arXiv 2511.04727. https://arxiv.org/abs/2511.04727

Freeman, J. (2025). *Student generative AI survey 2025*. Higher Education Policy Institute. https://www.hepi.ac.uk/2025/02/26/student-generative-ai-survey-2025/

Gilbertson, A., Dey, J., Deuchar, A., & Grills, N. (2021, August 1). India's COVID-19 divide in digital learning. *Pursuit*, University of Melbourne. https://pursuit.unimelb.edu.au/articles/india-s-covid-19-divide-in-digital-learning

Gmyrek, P., Winkler, H., & Garganta, S. (2024). *Buffer or bottleneck? Employment exposure to generative AI and the digital divide in Latin America* (ILO Working Paper 121). ILO and The World Bank. https://doi.org/10.54394/TFZY7681

Grassini, S. (2023). Development and validation of the AI attitude scale (AIAS-4): A brief measure of general attitude toward artificial intelligence. *Frontiers in Psychology*, 14, Article 1191628. https://doi.org/10.3389/fpsyg.2023.1191628

Hargittai, E. (2002). Second-level digital divide: Differences in people's online skills. *First Monday, 7*(4).

Helali, M. M., Mostafa, M. M., Aldawsari, H. K., Ibrahim, A. R., Bukhamseen, A. M., & Nemt-allah, M. A. (2026). Development and psychometric validation of the ethical AI dilemma anxiety scale among university students. *Frontiers in Psychology, 17*, 1810679. https://doi.org/10.3389/fpsyg.2026.1810679

Honneth, A. (1995). *The struggle for recognition: The moral grammar of social conflicts*. Polity Press.

International Labour Organization, & United Nations Office of the Secretary-General's Envoy on Technology. (2024). *Mind the AI divide: Shaping a global perspective on the future of work*.

Internet and Mobile Association of India, & Kantar. (2025). *Internet in India report 2024*. Internet and Mobile Association of India. https://www.iamai.in/sites/default/files/research/Kantar_%20IAMAI%20report_2024_.pdf

Internet and Mobile Association of India, & Kantar. (2026). *Internet in India report 2025*. Internet and Mobile Association of India. https://www.scribd.com/document/1064177762/Kantar-IAMAI-Report-2025-V6-Print-and-Web-Version-1

Jandrić, P., Knox, J., Besley, T., Ryberg, T., Suoranta, J., & Hayes, S. (2018). Postdigital science and education. *Educational Philosophy and Theory, 50*(10), 893-899. https://doi.org/10.1080/00131857.2018.1454000

Jin, Y., Yan, L., Echeverria, V., Gašević, D., & Martinez-Maldonado, R. (2025). Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. *Computers and Education: Artificial Intelligence*, 8, Article 100348. https://doi.org/10.1016/j.caeai.2024.100348

Kasperson, R. E., Renn, O., Slovic, P., Brown, H. S., Emel, J., Goble, R., Kasperson, J. X., & Ratick, S. (1988). The social amplification of risk: A conceptual framework. *Risk Analysis, 8*(2), 177-187.

Kasperson, R. E., Webler, T., Ram, B., & Sutton, J. (2022). The social amplification of risk framework: New perspectives. *Risk Analysis*, 42(7), 1367-1380. https://doi.org/10.1111/risa.13926

Kitsara, I. (2022). Artificial intelligence and the digital divide: From an innovation perspective. In Bounfour, A. (eds) Platforms and Artificial Intelligence. Progress in IS. Springer. https://doi.org/10.1007/978-3-030-90192-9_12

Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task [preprint]. arXiv 2506.08872. https://doi.org/10.48550/arXiv.2506.08872

LaFrance, J. (2026). Governing generative artificial intelligence: Institutional policies and guidelines at America's flagship universities. *Educational Policy*. https://doi.org/10.1177/08959048261450995

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1-22). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *Patterns, 4*(7), 100779. https://doi.org/10.1016/j.patter.2023.100779

Lodge, J. M., & Loble, L. (2026). *Artificial intelligence, cognitive offloading and implications for education*. University of Technology Sydney.

Lythreatis, S., Singh, S. K., & El-Kassar, A.-N. (2022). The digital divide: A review and future research agenda. *Technological Forecasting and Social Change*, 175, Article 121359. https://doi.org/10.1016/j.techfore.2021.121359

Maleki, A. (2026). Rethinking ethical academic integrity ecosystems in higher education in the age of artificial intelligence. *Discover Artificial Intelligence, 6*(1), Article 145. https://doi.org/10.1007/s44163-026-00887-z

Mateen, A., Kumar, V., Singh, A. K., Yadav, B., Mahto, M., Hassan, A., & Nasir, N. (2025). Impact of generative AI in medical education in India: A systematic review. *Frontiers in Artificial Intelligence, 8*, 1704785. https://doi.org/10.3389/frai.2025.1704785

Mathews, S. M. (2026). Writing in the age of generative AI: Cognitive debt, human agency, and homogenisation of voice. Economic and Political Weekly, 61(44). https://www.epw.in/journal/2026/44/perspectives/writing-age-generative-ai.html

Maxwell, D., Oyarzun, B., Kim, S., & Bong, J. Y. (2025). Generative AI in higher education: Demographic differences in student perceived readiness, benefits, and challenges. *TechTrends, 69*(6), 1248-1259. https://doi.org/10.1007/s11528-025-01109-6

Mazumder, D., Pathak, D., Kodali, P., Joshi, A., Agarwal, A., & Patro, J. (2026). *Evaluating cross-lingual knowledge consistency in code-mixed vis-a-vis Indian languages using IndicKLAR* [preprint]. arXiv 2605.29637. https://arxiv.org/abs/2605.29637

Menon, D., & Shilpa, K. (2023). "Chatting with ChatGPT": Analyzing the factors influencing users' intention to use the Open AI's ChatGPT using the UTAUT model. *Heliyon*, 9(11), Article e20962. https://doi.org/10.1016/j.heliyon.2023.e20962

Ministry of Education. (2026). *All India Survey on Higher Education (AISHE) 2023-24*. Government of India. https://aishe.gov.in/

Ministry of Education of the People's Republic of China. (2025, June 11). *2024 nian quanguo jiaoyu shiye fazhan tongji gongbao* [Statistical bulletin on the development of national education in 2024]. http://www.moe.gov.cn/jyb_sjzl/sjzl_fztjgb/202506/t20250611_1193760.html

Ministry of Electronics and Information Technology. (2024). *IndiaAI Mission*. Government of India.

Ministry of Human Resource Development. (2020). *National Education Policy 2020*. Government of India.

Mishra, K. (2026). *The privilege of exposure: Caste and generative AI in India's graduate labour market* [preprint]. arXiv 2606.13314. https://arxiv.org/abs/2606.13314

Morrow, W. (2009). *Bounds of democracy: Epistemological access in higher education*. HSRC Press.

Mohamed, S., Png, M.-T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology, 33*(4), 659-684. https://doi.org/10.1007/s13347-020-00405-8

Niazi, S. (2025, October 16). AI is already transforming Indian higher education. *University World News.* https://www.universityworldnews.com/post.php?story=20251016140031657

Nyaaba, M., Wright, A. L., & Choi, G. L. (2026). Generative AI and digital neocolonialism in global education: Towards an equitable framework. *Journal of Digital Educational Technology, 6*(1), ep2608. https://doi.org/10.30935/jdet/17862

Oyelere, S. S., & Aruleba, K. (2025). A comparative study of student perceptions on generative AI in programming education across Sub-Saharan Africa. *Computers and Education Open*, 8, Article 100245. https://doi.org/10.1016/j.caeo.2025.100245

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

Pettit, P. (1997). *Republicanism: A theory of freedom and government*. Oxford University Press.

Press Information Bureau. (2026a, March 3). *AI in education: Building India's talent pipeline for global leadership*. Government of India.

Press Information Bureau. (2026b, February 15). *India AI governance guidelines*. Government of India. https://www.pib.gov.in/PressReleasePage.aspx?PRID=2228315&reg=48&lang=2

Ragnedda, M. (2017). *The third digital divide: A Weberian approach to digital inequalities*. Routledge.

Rahman, T. (2005). Passports to privilege: The English-medium schools in Pakistan. *Peace and Democracy in South Asia, 1*(1), 24-44.

Raman, R., Mandal, S., Das, P., Kaur, T., Sanjanasri, J. P., & Nedungadi, P. (2024). Exploring university students' adoption of ChatGPT using the diffusion of innovation theory and sentiment analysis with gender dimension. *Human Behavior and Emerging Technologies*, 2024(1), Article 3085910. https://doi.org/10.1155/2024/3085910

Ravšelj, D., Keržič, D., Tomaževič, N., Umek, L., Brezovar, N., et al. (2025). Higher education students' perceptions of ChatGPT: A global study of early reactions. *PLOS ONE*, 20(2), Article e0315011. https://doi.org/10.1371/journal.pone.0315011

Robinson, L., Cotten, S. R., Ono, H., Quan-Haase, A., Mesch, G., Chen, W., Schulz, J., Hale, T. M., & Stern, M. J. (2015). Digital inequalities and why they matter. *Information, Communication & Society, 18*(5), 569-582.

Sabharwal, N. S. (2021). Nature of access to higher education in India: Emerging pattern of social and spatial inequalities in educational opportunities. In M. S. Jaglan & Rajeshwari (Eds.), *Reflections on*

*21st century human habitats in India* (pp. 345-369). Springer. https://doi.org/10.1007/978-981-16-3100-9_14

Sabharwal, N. S. (2024). Understanding students' attitudes towards affirmative action policy in higher education in India. *Social Inclusion, 12*. https://doi.org/10.17645/si.7601

Sabharwal, N. S., Henderson, E. F., & Joseph, R. S. (2020). Hidden social exclusion in Indian academia: Gender, caste and conference participation. *Gender and Education, 32*(1), 27-42.

Sabharwal, N. S., & Malish, C. M. (2017). Student diversity and challenges of inclusion in higher education in India. *International Higher Education*, (91), 25-27. https://doi.org/10.6017/ihe.2017.91.10136

Sabharwal, N. S., & Malish, C. M. (2018). *Student diversity and social inclusion: An empirical analysis of higher education institutions in India* (CPRHE Research Paper 10). Centre for Policy Research in Higher Education, NIEPA.

Sankalp KJ, Kumar, A., Balaji, L., Kotecha, N., Jain, V., Chadha, A., & Bhaduri, S. (2025). *IndicMMLU-Pro: Benchmarking Indic large language models on multi-task language understanding* [preprint]. arXiv 2501.15747. https://arxiv.org/abs/2501.15747

Schoon, A., Smit, A., Uzuegbunam, C., & Bosch, T. (2026). Adoption and perceptions of generative AI among South African academics. *Teaching in Higher Education, 31*(2), 248-265. https://doi.org/10.1080/13562517.2025.2571911

Scheerder, A., van Deursen, A., & van Dijk, J. (2017). Determinants of Internet skills, uses and outcomes: A systematic review of the second- and third-level digital divide. *Telematics and Informatics, 34*(8), 1607-1624.

Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59-68). ACM.

Singh, H., Gupta, N., Bharadwaj, S., Tewari, D., & Talukdar, P. (2024). IndicGenBench: A multilingual benchmark to evaluate generation capabilities of LLMs on Indic languages. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 11047-11073). Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.acl-long.595

Stephenson, R., & Armstrong, C. (2026). *Student generative AI survey 2026 (HEPI Report 199).* Higher Education Policy Institute. https://www.hepi.ac.uk/wp-content/uploads/2026/03/HEPI-Report-199-Gen-AI-Survey-2026.pdf

Strzelecki, A. (2024). Students' acceptance of ChatGPT in higher education: An extended Unified Theory of Acceptance and Use of Technology. *Innovative Higher Education, 49*(2), 223–245. https://doi.org/10.1007/s10755-023-09686-1

Telecom Regulatory Authority of India. (2026). *The Indian telecom services performance indicators, January-March 2026*. TRAI.

The Editorial. (2026). *UGC AI plagiarism guidelines 2026: Complete guide for Indian PhD students*.

The Secretariat. (2025, September 25). *Forging the future: AI as a balancing act, AICTE as a catalyst for change*.

Thakur, D. (2026, June 22). AI divide threatens India's higher education goals. *Policy Circle*. https://www.policycircle.org/opinion/ai-divide-india-higher-education/

Thorat, S., & Newman, K. S. (Eds.). (2010). *Blocked by caste: Economic discrimination in modern India*. Oxford University Press.

Thorat, S., & Sabharwal, N. S. (2015). Caste and social exclusion: Concept, indicators, and measurement. In *India's children* (pp. 374-392). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199455287.003.0015

Tierney, W. G., & Sabharwal, N. S. (2018). Reimagining Indian higher education: A social ecology of higher-education institutions. *Teachers College Record*, 120(5), 1-32. https://doi.org/10.1177/016146811812000504

Tierney, W. G., Sabharwal, N. S., & Malish, C. M. (2019). Inequitable structures: Class and caste in Indian higher education. *Qualitative Inquiry*, 25(5), 471-481. https://doi.org/10.1177/1077800418817836

UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO.

University Grants Commission. (2018). *UGC (Promotion of Academic Integrity and Prevention of Plagiarism in Higher Educational Institutions) Regulations, 2018*.

University Grants Commission. (2023). *Guidelines for training/orientation of faculty on Indian Knowledge System (IKS)*. UGC.

Üstün, A., Akçelik, E., & Marabaoğlu, S. (2026). Artificial intelligence and technological anxiety among university students: A comparative analysis. *Journal of Mehmet Akif Ersoy University Economics and Administrative Sciences Faculty, 13*(3), 1121–1144. https://doi.org/10.30798/makuiibf.1897002

van Deursen, A. J. A. M., & Helsper, E. J. (2015). The third-level digital divide: Who benefits most from being online? In *Communication and information technologies annual* (Studies in Media and Communications, Vol. 10, pp. 29-52). Emerald. https://doi.org/10.1108/S2050-2060201510

van Deursen, A. J., & van Dijk, J. A. (2013). The digital divide shifts to differences in usage. *New Media & Society*, 16(3), 507-526. https://doi.org/10.1177/1461444813487959

van Dijk, J. A. G. M. (2005). *The deepening divide: Inequality in the information society*. Sage.

van Dijk, J. A. G. M. (2020). *The digital divide*. Polity Press.

Varghese, N. V. (2022). Directions of change in higher education in India: From massification to universalization. In S. Chattopadhyay, S. Marginson, & N. V. Varghese (Eds.), *Changing higher education in India*. Bloomsbury Academic. https://doi.org/10.5040/9781350192409.ch-001

Varghese, N. V., & Sabharwal, N. S. (2022). *The future of higher education in India: From massification to universalization*. National Institute of Educational Planning and Administration.

Varghese, N. V., & Sabharwal, N. S. (Eds.). (2024). *India higher education report 2022: Women in higher education*. Routledge.

Varghese, N. V., Sabharwal, N. S., & Malish, C. M. (Eds.). (2018). *India higher education report 2016: Equity*. Sage.

Varghese, N. V., Sabharwal, N. S., & Malish, C. M. (2019). *Equity and inclusion in higher education in India* (CPRHE Research Paper 12). Centre for Policy Research in Higher Education, NIEPA.

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly, 27*(3), 425–478. https://doi.org/10.2307/30036540

Wang, C., Boerman, S. C., Kroon, A. C., Möller, J., & de Vreese, C. H. (2025). The artificial intelligence divide: Who is the most vulnerable? *New Media & Society*, 27(7), 3867-3889. https://doi.org/10.1177/14614448241232345

Wang, Y. Y., & Wang, Y. S. (2022). Development and validation of an artificial intelligence anxiety scale: An initial application in predicting motivated learning behavior. *Interactive Learning Environments, 30*(4), 619-634. https://doi.org/10.1080/10494820.2019.1674887

Watermeyer, R., Lanclos, D., Phipps, L., Shapiro, H., Guizzo, D., & Knight, C. (2025). Academics' weak(ening) resistance to generative AI: The cause and cost of prestige? *Postdigital Science and Education, 7(*4), 1171–1191. https://doi.org/10.1007/s42438-024-00524-x

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., & Waddington, L. (2023). Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19, Article 26. https://doi.org/10.1007/s40979-023-00146-z

Williamson, B., Bayne, S., & Shay, S. (2020). The datafication of teaching in higher education: Critical issues and perspectives. *Teaching in Higher Education, 25*(4), 351-365. https://doi.org/10.1080/13562517.2020.1748811