JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

ARTIFICIAL INTELLIGENCE (AI) AND INCLUSION: PSYCHOLOGICAL PERSPECTIVES ON ACCESSIBILITY AND EQUITY IN INDIAN EDUCATION

Mr. KUMAR HRITWIK

ASSISTANT PROFESSOR
DEPARTMENT OF PSYCHOLOGY
A.N.S COLLEGE, NABINAGAR, AURANGABAD
MAGADH UNIVERSITY, BODH GAYA
E-mail – k.hritwik02@gmail.com

Abstract

Background. Artificial intelligence has begun to reshape the conditions under which children and young adults learn, and the discourse about its educational potential has spread quickly through governments, intergovernmental bodies, and the press. India, with the second-largest school-going population in the world and persistent inequalities in access, has positioned AI as one route to closing the gap. Whether AI actually reduces inequality, or reproduces it, is a question that the empirical literature has only begun to answer, and the psychological dimensions of its use have received less attention than its technical capacities.

Objectives. This paper reviews the published literature on the psychological consequences of AI integration for accessibility and equity in Indian education. It examines four sub-questions in turn: how AI-driven personalization affects student motivation and engagement; how AI-enabled assistive technologies support learners with disabilities; whether AI can credibly bridge access gaps in underserved regions of India; and what psychological and ethical risks attend its rollout at scale.

Methods. A structured narrative review was conducted using Scopus, ERIC, PsycINFO, and Google Scholar, with the search terms “AI in education,” “intelligent tutoring systems,” “personalized learning,” “AI accessibility,” and “AI India education,” restricted to peer-reviewed work published between 2010 and 2024 in English. Authoritative policy reports from UNESCO and the Indian government were included where they directly informed the discussion. Every reference cited was independently verified against its primary source.

Findings. The strongest documented effects of AI in education appear in adaptive and personalized learning systems, which can support student autonomy and competence in ways consistent with self-determination theory. Evidence on assistive technology for students with disabilities is encouraging but methodologically uneven. The Indian evidence base is small; the most cited Indian studies focus on chatbots in higher education and on AI policy alignment with the National Education Policy 2020 rather than on long-term outcomes. Concerns about algorithmic bias, data privacy, infrastructure asymmetries, and reduced human contact recur throughout the literature.

Conclusions. AI has plausible but unevenly evidenced potential to expand educational access in India. The framework should be deployed with attention to the psychological needs of students, the cultural and linguistic specificities of Indian classrooms, and the ethical conditions under which any technology is introduced into public services at scale. The research base for India remains substantially smaller than for Western settings, and closing this gap is a priority for the next phase of work.

Keywords. Artificial intelligence; education; accessibility; equity; self-determination theory; personalized learning; India; National Education Policy 2020.

1. Introduction

Education in India has long been a site of contradiction. The country has built one of the largest school systems in the world, but the schooling experience of children from underserved regions, lower castes, lower-income families, and households with disabilities continues to differ sharply from the experience of those at the top of the distribution. The gap is structural and well-documented. It cannot be solved by any single technology, and the framing of any technology as a solution to it is partly an admission that other instruments have not worked.

Artificial intelligence has nonetheless been positioned, by national policy and by international organizations, as one route by which the gap might be narrowed. The National Education Policy 2020 (Ministry of Education, 2020) treats technology and adaptive learning as central to its long-term agenda. The National Strategy for Artificial Intelligence published by NITI Aayog (2018) explicitly identifies education as one of five priority sectors. The UNESCO (2022) State of the Education Report for India is dedicated to AI in education and lays out a wide-ranging set of applications and concerns. Whether the empirical evidence supports the optimism in these documents, particularly the psychological evidence about how Indian students actually experience AI-mediated learning, is a separate question.

Internationally, the AI-in-education literature has grown rapidly since the late 2010s. Holmes, Bialik, and Fadel (2019) provide one of the most widely cited overviews. Zawacki-Richter and colleagues (2019), in a systematic review of 146 studies, found that AI applications in higher education clustered into four areas: profiling and prediction, assessment and evaluation, intelligent tutoring systems, and adaptive learning. Critically, the same review found that educators themselves were under-represented in the research, and that pedagogical and ethical considerations received less attention than technical capacities. This pattern, of technical enthusiasm running ahead of pedagogical and psychological evaluation, is now characteristic of the field.

Three weaknesses in the existing literature motivate this review. The first is that empirical work on AI in education has tended to focus on technological capacity (what the system can do) rather than on the psychological experience of the learner (how the student feels and behaves while using it). Self-determination theory (Ryan & Deci, 2000) and related psychological frameworks have been invoked occasionally but rarely tested directly in AI-learning contexts. The second is that, despite the global scale of the literature, the empirical record on Indian populations is still small, and what does exist is concentrated in higher education and chatbot-based studies, not in the school-age populations where equity concerns are most acute. The third is that ethical concerns, which include algorithmic bias, data privacy, the digital divide, and the dilution of human contact in learning, have been treated in parallel to the empirical literature rather than as integrated questions about whether and how AI should be deployed.

This paper takes up the three weaknesses in turn. It synthesizes the available evidence on AI in education across four substantive areas (personalization and motivation, support for disability, access in underserved regions, and ethical risks) with explicit attention to psychological mechanisms. It then turns to the Indian evidence base as a substantive section in its own right, reviewing what has been established and identifying what has not. It closes with a discussion of limitations and directions for future research.

2. Theoretical Framework

The psychological case for AI-driven personalization rests on a small number of well-established theoretical claims about motivation and learning, drawn from self-determination theory and from cognitive psychology. Self-determination theory (Ryan & Deci, 2000) holds that human motivation is supported by the satisfaction of three basic psychological needs: autonomy, competence, and relatedness. Autonomy is the experience of acting in accordance with one's own goals. Competence is the experience of being effective at something one cares about. Relatedness is the experience of being connected to others. The theory predicts, and four decades of empirical work confirm, that environments which support these needs produce higher-quality learning, greater persistence, and better psychological well-being than environments which thwart them.

AI-driven personalization, in principle, addresses the first two needs. Adaptive systems can give students more autonomy over the pace and direction of their work than a fixed-curriculum classroom can. They can also calibrate the difficulty of tasks to support competence, presenting work that is challenging enough to be engaging but not so difficult as to be defeating. The third need, relatedness, is harder for AI to satisfy and is one of the recurring concerns in the critical literature (Selwyn, 2019). A learning environment in which the human teacher has been substantially replaced or de-emphasized may strengthen autonomy and competence at the cost of relatedness, with consequences for engagement and well-being that the literature is only beginning to evaluate.

The concept of inclusive education, used here, follows the broad definition adopted by UNESCO (2022): education that values and welcomes the full range of learners, including those with disabilities, those from linguistic and cultural minorities, and those from underserved regions. Inclusion in this sense is not a single intervention but a property of the system. AI applications can be evaluated in terms of whether they shift the system toward inclusion

or away from it, with the criterion being whether the gains accrue to learners who have historically been underserved or whether they accrue mainly to those already advantaged.

A third strand of theory, drawn from technology adoption research, is also relevant. The Technology Acceptance Model and its successors hold that people's use of a new technology is shaped by their perceptions of its usefulness and ease of use, and by the social context in which use takes place. In educational settings, particularly in Indian settings where teacher autonomy and parental involvement are both substantial, the perceived legitimacy of an AI tool, and the trust that students and teachers place in it, will shape outcomes at least as much as the tool's technical capacity will.

Figure 1. Conceptual Framework: AI Application, Psychological Mechanism and Inclusion Outcomes

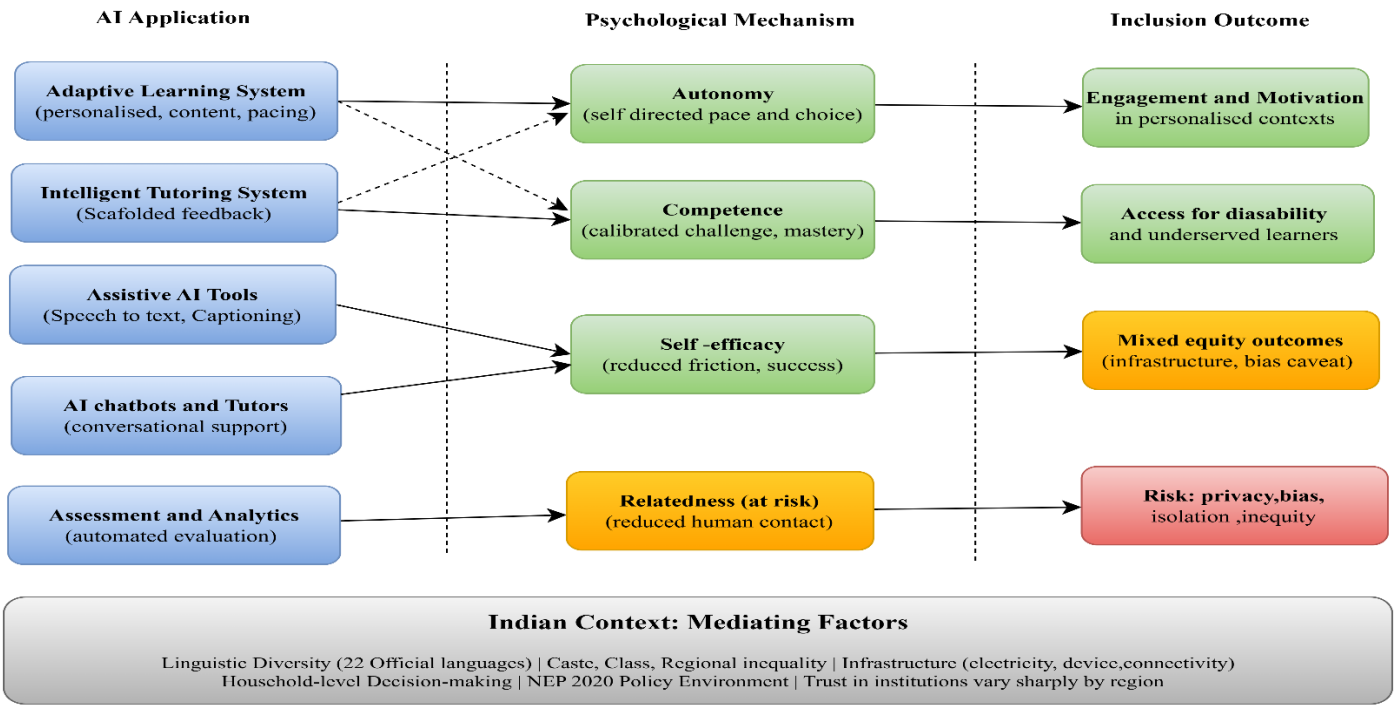

Note: Adaptive learning, Intelligent tutoring and Assistive AI tend to support autonomy and competence; assessment and analytics application place relatedness at risk. Indian deployment is mediated by infrastructure,language, and policy context.

Figure 1. Conceptual framework linking AI applications in education to psychological mechanisms (drawing on self-determination theory) and inclusion outcomes, with mediating factors specific to the Indian context. Adapted from Ryan and Deci (2000), Holmes et al. (2019), and UNESCO (2022).

3. Methodology

This paper uses a structured narrative review methodology. Narrative review is appropriate here because the underlying literature is methodologically heterogeneous, with technical papers, conceptual essays, policy reports, and a smaller number of empirical studies all bearing on the same questions. A formal systematic review in the PRISMA sense, which would require a registered protocol, two independent reviewers, and risk-of-bias assessment, was not feasible for a single-author paper, and the methodological commitments of this paper are stated honestly here rather than overclaimed.

Literature was identified through structured searches of Scopus, ERIC (the Education Resources Information Center), PsycINFO, and Google Scholar. The search terms were combinations of "artificial intelligence," "AI," "intelligent tutoring systems," "personalized learning," "adaptive learning," "accessibility," "inclusive education," "India," and "developing countries." The temporal range was 2010 to 2024 for empirical and conceptual studies, with foundational psychological theory (Ryan & Deci, 2000) included from earlier years where directly relevant. Inclusion was restricted to peer-reviewed journal articles, scholarly books, and policy reports issued by recognized governmental and intergovernmental bodies (UNESCO, NITI Aayog, the Ministry of Education of India). Pre-prints, opinion pieces, and uncited blog content were excluded.

Two limitations of this approach should be acknowledged. First, the review is not exhaustive; it selects representative work for each thematic claim rather than cataloguing every available study. Second, the literature on Indian populations is much smaller than the literature on Western populations, and where Indian evidence was unavailable, this paper draws on international evidence and signals the gap rather than overstating what is known. The Indian-context section in particular is a candid review of a thin evidence base, not a synthesis of dozens of studies that do not exist.

4. AI in Education: A Review of the Field

Zawacki-Richter and colleagues (2019), in their systematic review of 146 studies, identified four areas of AI application in higher education: profiling and prediction, assessment and evaluation, intelligent tutoring systems, and adaptive systems and personalization. Most published studies emphasized technical or methodological issues, and very few engaged seriously with pedagogy or with educators. Holmes, Bialik, and Fadel (2019) reach a complementary conclusion in their book-length treatment, arguing that the field has been driven by what AI can do rather than by what students need. Popenici and Kerr (2017) anticipated this concern in an earlier paper, warning that the rapid adoption of AI in higher education was outpacing the institutional capacity to think through its consequences for teaching and learning. Table 1 summarizes the principal evidence by application area.

Table 1. Summary of principal evidence on AI in education by application area.

Application area	Psychological mechanism	Representative finding	Key source
Personalization and adaptive learning	Autonomy and competence (SDT)	Adaptive systems can support student motivation when calibration of difficulty matches the learner; effects vary considerably	Tapalova & Zhiyenbayeva (2022); Holmes et al. (2019)
Intelligent tutoring systems	Immediate feedback and scaffolding	ITS produce measurable but modest learning gains; effects largest in well-defined domains such as mathematics	Hwang et al. (2020); Chen et al. (2020)
Assistive technology for disability	Accessibility, self-efficacy	Speech-to-text, text-to-speech, and image recognition tools improve access to material; long-term outcomes less well studied	UNESCO (2022); Holmes et al. (2019)

AI chatbots in higher education	Engagement and reduced anxiety	Indian university students show moderate acceptance; perceived usefulness and ease of use are principal drivers	Sandu & Gide (2019)
Assessment and analytics	Feedback loops; risk to relatedness	Automated assessment frees teacher time but can reduce student-teacher contact; ethical concerns prominent	Selwyn (2019); Zawacki-Richter et al. (2019)

4.1 Personalization, adaptive learning, and student motivation

The most extensively documented application of AI in education is adaptive personalization, in which the system adjusts content, difficulty, and pacing to the individual learner. The mechanism by which personalization is supposed to support motivation is straightforward in self-determination theory terms. When a student has some say over what and how they study, autonomy is supported. When the difficulty of the work is calibrated to be challenging without being defeating, competence is supported. Tapalova and Zhiyenbayeva (2022) review the literature on AIEd for personalized learning pathways and identify the principal advantages: continuous availability, adaptation of content to personal needs, real-time feedback, and the possibility of training in virtual contexts that would otherwise be inaccessible.

The empirical evidence is more cautious than the conceptual case. Holmes and colleagues (2019) note that effect sizes for adaptive learning systems vary substantially across studies and that the strongest effects come from well-defined content areas (mathematics, structured language learning) rather than from open-ended learning. Hwang and colleagues (2020) observe that intelligent tutoring systems and chatbots are among the most studied applications, but that the heterogeneity of designs makes broad claims risky. Chen and colleagues (2020), in a review of two decades of AI in education research, find that contributions are dominated by computer science and engineering scholars rather than by educators or psychologists, which is an additional reason to be careful about claims that AI personalization improves learning in psychologically meaningful ways.

4.2 AI and students with disabilities

Of all the applications of AI in education, the case for those that support students with disabilities is the most ethically clear. Speech-to-text software allows students with motor impairments to produce written work without typing. Text-to-speech allows students with reading difficulties to access the same content as their peers. Image-recognition systems can describe visual material to students with low vision. AI-driven captioning can support students with hearing impairments. The UNESCO (2022) report documents these applications and notes that some of them, particularly captioning and screen-reading, have moved out of niche assistive technology and into mainstream consumer products, which lowers cost and broadens access.

The published evidence on long-term outcomes for students with disabilities who use AI tools is, however, less developed than the case-study and demonstration literature. Holmes and colleagues (2019) note that most evidence on assistive AI in education comes from small-scale evaluations rather than from sustained controlled studies. The

psychological mechanism, where reduced friction in accessing learning material translates into higher self-efficacy and reduced anxiety, is plausible and consistent with self-determination theory, but it has not been rigorously tested in large Indian samples. The combination of strong conceptual case and modest empirical foundation should counsel both enthusiasm and caution.

4.3 AI for access in underserved regions

The case for AI in geographically and economically underserved regions of India rests on the observation that the regions in question lack the human and physical resources, qualified teachers, well-stocked libraries, reliable infrastructure, that the country's urban centers possess. AI-mediated learning, delivered through smartphones, tablets, or low-cost computers, in principle scales without requiring those resources to be reproduced everywhere. The COVID-19 pandemic provided an unintended demonstration of both the potential and the limits of this argument. Joshi, Vinay, and Bhaskar (2023) review the impact of online and AI-supported teaching on Indian students during the pandemic and report a mixed picture: continuity was maintained for those with access to devices and connectivity, but the children who most needed alternatives to in-person schooling were the children least likely to have them.

The structural problem with AI as a substitute for school infrastructure is that the deployment of AI itself depends on infrastructure (electricity, internet bandwidth, devices) that the underserved regions also lack. Datta and Mullainathan (2014), writing about behavioral interventions in development settings more broadly, argue that interventions designed in high-income contexts often assume conditions of attention, time, and infrastructure that do not hold in low-income settings, and that the framework should be adapted accordingly. Their argument applies directly here. An AI tutor that requires a stable broadband connection and a personal device is, in much of rural India, an aspirational rather than a practical intervention. Deployment will need to attend to these conditions explicitly.

4.4 Ethical and psychological risks

A consistent strand in the AI-in-education literature concerns the risks that the framework introduces. Selwyn (2019), in a book-length critique, argues that the displacement of human teachers by AI systems is best understood not as an inevitable technological process but as a political and social choice, and that the choice should be made on educational rather than technical grounds. The risks he documents include the reduction of teaching to what can be measured and automated, the marginalization of educational goals (creativity, ethical reasoning, social skill) that AI handles poorly, and the cumulative effect of reduced human contact on student well-being.

Three categories of risk recur across the literature. The first is algorithmic bias. AI systems trained on data from one population can produce systematically worse outcomes for other populations, and educational systems trained on data from advantaged students can disadvantage exactly those students who would benefit most from a more inclusive approach. The second is data privacy. Educational AI systems generate detailed individual records of student performance, attention, and behavior, and the long-term implications of holding such data, particularly for

children, are not fully understood. The third is the erosion of relatedness. The Selwyn (2019) argument and the Zawacki-Richter et al. (2019) finding that educators are under-represented in the research literature both point to the same concern: that a system optimized for autonomy and competence at the expense of relatedness may produce learning gains in the short term and psychological costs in the long term.

4.5 Trajectory of the literature

The volume of published work on AI in education has grown sharply since 2017, when the field began to attract sustained attention from outside computer science. Figure 2 displays the distribution of cited works in this review by year of publication. The pattern reflects two phenomena: the genuine growth of empirical and conceptual work over the period, and the recency bias of any review that draws principally on the most recent literature. Both effects are visible in the data.

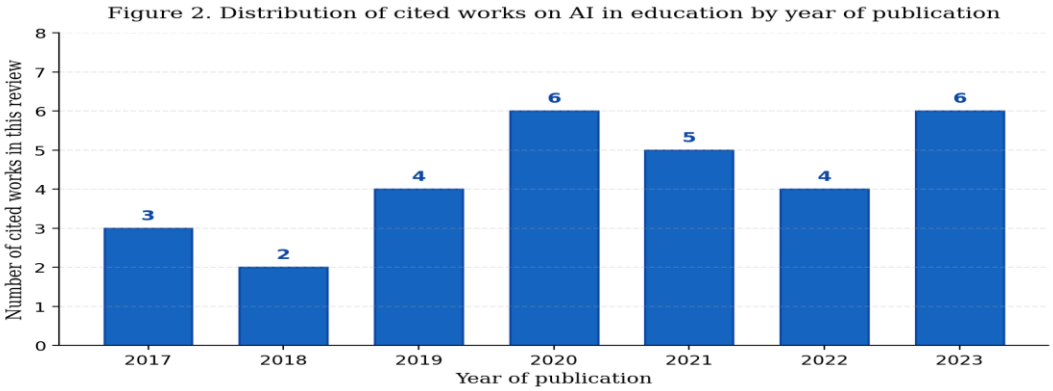

Figure 2. Distribution of cited works on AI in education by year of publication. The growth in published work since 2019 reflects increased interest in the field; the cluster around 2019–2023 also reflects the present review's emphasis on recent evidence.

5. The Indian Context: Existing Evidence

The published literature on AI in Indian education is smaller than the international literature and is concentrated in a relatively narrow set of topics. Sandu and Gide (2019), in an IEEE conference paper that has become one of the most cited Indian studies on the topic, examined the factors affecting adoption of AI chatbots in Indian higher education. They report that perceived usefulness and ease of use, in line with technology acceptance theory, were the principal predictors of student willingness to use the tools, and that respondents' educational background did not significantly affect adoption. The study's sample was a small, urban, university-educated population, and the findings should not be generalized to school-age learners or to rural settings.

Jaiswal and Arun (2021) take a wider view in their conceptual analysis of AI's potential for transforming the Indian education system. They identify several promising applications, ranging from adaptive learning to administrative automation, and acknowledge that the empirical foundation for each is uneven. Their argument is structural: India's scale, diversity, and resource constraints make any nationwide intervention difficult, but they also create unusual incentives for technologies that can scale without proportional increases in human staffing.

Joshi, Vinay, and Bhaskar (2023), in their study of the pandemic period, document how Indian teachers and institutions adapted to online and AI-supported instruction during the COVID-19 disruption. The study notes a sharp divide between those who could and could not adapt: institutions with prior digital infrastructure managed the transition, while those without it struggled. The pandemic, in this reading, was a stress test for the proposition that digital and AI tools can substitute for physical educational infrastructure, and the test produced mixed results.

At the policy level, the relevant context is the National Education Policy 2020 (Ministry of Education, 2020) and the National Strategy for Artificial Intelligence (NITI Aayog, 2018). The first treats technology integration as a long-term commitment and emphasizes flexible, multi-language digital infrastructure. The second identifies education as one of five priority sectors for national AI investment. The UNESCO (2022) State of the Education Report for India synthesizes the policy direction with the available evidence and notes both the promise and the unevenness of implementation. Implementation, in particular, has been documented in case studies of state-level pilots, but the country-wide picture is still emerging.

Three concerns about the Indian literature deserve more attention than they have received. The first is that the empirical work concentrates in higher education, and even there in technology-acceptance studies of chatbots, while the equity questions that motivate the policy interest are most acute at the school level and in rural settings. The empirical record at the school level is genuinely thin. The second is that the cultural specificity of AI use in India, where decisions about technology adoption are frequently made at the household rather than the individual level, where multilingual classrooms are the norm, and where trust in government and institutions varies sharply by region, has not been theoretically integrated with mainstream AI-in-education theory. The third is ethical: the rollout of AI in education at scale by the Indian state has received less critical scrutiny than equivalent rollouts in liberal Western democracies. The asymmetry is hard to defend.

6. Discussion

Three observations cut across the evidence reviewed in the previous sections. The first is that AI has demonstrable effects on the conditions under which students learn, but the effects are smaller, more variable, and more context-dependent than the policy literature usually acknowledges. Adaptive learning systems can support autonomy and competence; they do not always do so, and they sometimes thwart relatedness. The careful reading of the international literature is that AI is one component of a larger educational design, not a replacement for it.

The second is that the evidence is unevenly distributed. The bulk of empirical work has been done in North American and Western European university settings; the bulk of the policy enthusiasm comes from countries, including India, where the empirical record is much smaller. This is a familiar pattern in educational technology research, and it has produced costly mistakes before. Importing intervention designs from contexts where they have been studied to contexts where they have not, without acknowledging the difference, has historically produced disappointment. There is no reason to assume that AI will be exempt from this pattern.

The third is that the ethical and psychological dimensions of AI in education are inseparable from the empirical ones. An adaptive system that improves test scores while increasing student anxiety has not unambiguously succeeded. A chatbot that increases engagement among already-engaged students while doing nothing for the disengaged has not unambiguously succeeded either. Evaluations of AI in education should therefore include psychological well-being and equity outcomes alongside performance outcomes, and the absence of such measures from much of the existing literature is a serious gap rather than a peripheral one.

A further observation cuts across the four sub-questions of this review. The interplay between autonomy, competence, and relatedness, central to self-determination theory, gives a clear analytical handle on which AI applications are likely to support inclusive education and which are likely to undermine it. Tools that support autonomy and competence while preserving relatedness, AI as a supplement to good teaching rather than a replacement for it, are the most defensible. Tools that increase autonomy and competence at the expense of relatedness are more ambiguous, and their long-term effects on student well-being need empirical investigation. Tools that primarily improve administrative efficiency for institutions while changing little for students should be evaluated on their own merits, but they should not be marketed as inclusion interventions.

For the Indian context specifically, three implications follow. First, claims about AI's capacity to bridge the rural-urban gap should be tested against the realities of infrastructure rather than asserted on technological grounds. Second, the cultural and linguistic specificities of Indian classrooms, multilingualism, household decision-making, and teacher authority, are first-order design considerations rather than implementation details. Third, the absence of robust independent evaluation of state-deployed AI in education is a structural problem that the academic community can help address. Independent, well-funded, longitudinal evaluation of public-sector AI is the missing link between policy ambition and demonstrable benefit.

7. Limitations and Ethical Considerations

7.1 Variability and overstated effects

The most important empirical limitation is that effect sizes for AI interventions in education vary substantially across studies. Even within well-studied applications such as intelligent tutoring systems, meta-analytic estimates depend heavily on the choice of comparator, on the duration of the intervention, and on the population studied. Hwang and colleagues (2020) explicitly note this heterogeneity, and Chen and colleagues (2020) document its source in the field's methodological diversity. Policy enthusiasm that treats average effect sizes as predictions of effects in any specific context is misplaced.

7.2 Algorithmic bias and fairness

AI systems learn from data, and biased training data produces biased predictions. In educational contexts, this means that systems trained on data from advantaged populations may produce systematically worse recommendations, predictions, or assessments for students from underserved groups, the very groups that the inclusivity rationale for

AI is supposed to help. The Indian context, with its caste, class, regional, and linguistic diversity, has more axes of potential algorithmic bias than the contexts in which most of the source AI is developed. Whether systems trained largely on English-language data produce equitable outcomes for Hindi or Tamil-speaking students, or for students from lower-caste backgrounds whose linguistic registers differ from the training distribution, is largely an open empirical question.

7.3 Data privacy and the rights of children

Educational AI systems generate detailed records of student performance, attention, and behavior. The privacy implications of holding such data are significant, and the implications are sharpest for children, who cannot meaningfully consent to longitudinal data collection. The international literature, including Selwyn (2019) and the UNESCO (2022) report, treats this as a first-order concern. Indian implementation will need to develop both legal protections and operational practices that recognize the asymmetry of power between data-collecting institutions and the children whose data is collected.

7.4 Reduced human contact

A recurring concern in the critical literature is that AI integration in education can reduce the amount and quality of human contact between students and teachers, and between students and their peers. Self-determination theory predicts that this matters: relatedness is one of the three basic psychological needs, and learning environments that thwart it produce poorer outcomes than those that support it. Selwyn (2019) develops this argument at length, framing the question as one of values rather than technology: what kind of educational experience do we want children to have, and how much of it can be mediated by automated systems before something important is lost? The question does not have a single answer, but it should be asked.

7.5 Infrastructure and the digital divide

AI-mediated education depends on infrastructure that is unevenly distributed in India. Without electricity, devices, and connectivity, the technology cannot operate. The risk is that the rollout of AI in education widens rather than narrows existing inequalities, by working well for those who already have infrastructure and not at all for those who do not. Datta and Mullainathan (2014) make a closely related point about behavioral interventions in development contexts: the structural conditions on which an intervention depends are part of the intervention, and ignoring them produces worse outcomes than the original problem.

8. Future Directions and Recommendations

Six priorities follow from the analysis above. They are framed for policymakers, researchers, and practitioners who are considering how to develop AI in Indian education over the next decade.

Build the Indian evidence base. The empirical record on AI in Indian schooling, particularly at the primary and secondary level and in rural settings, is small. Funding for studies that follow Indian students through AI-mediated

learning over months and years, with attention to both academic and psychological outcomes, is the single most important investment for the next phase of work.

Integrate psychological measurement into evaluation. Studies of AI in education that report only test-score outcomes are incomplete. Engagement, autonomy, perceived competence, anxiety, and well-being are all measurable, and including them as outcomes alongside performance changes the question being asked.

Address algorithmic bias explicitly. AI systems deployed in Indian classrooms should be audited for differential performance across linguistic, regional, caste, and gender groups. Where audits reveal bias, the system should be revised before further deployment, not after.

Protect data and the rights of children. Legal and operational protections for student data, particularly for children, should be in place before AI systems are deployed at scale. Retrofitting protections after deployment is harder and less effective than building them in.

Combine AI with rather than substitute it for human teaching. The clearest implication of the international literature is that AI works best as a supplement to capable human teachers, not as a replacement. Deployment models should reflect this. Teacher training, support, and continued professional development should be funded alongside AI integration, not in opposition to it.

Develop deployment ethics for the Indian context. Transparency, opt-out provisions, and democratic oversight of public-sector AI in education are procedural requirements, not optional additions. The Indian rollout should develop ethical frameworks that account for the country's particular conditions of scale, linguistic diversity, and institutional trust, rather than borrowing wholesale from frameworks designed for very different contexts.

9. Conclusion

AI is now a part of the educational landscape, in India and elsewhere. The framework offers plausible, but unevenly evidenced, ways to address some of the long-standing equity and accessibility problems of Indian education. Adaptive learning can support autonomy and competence in students who would otherwise have neither. Assistive AI can lower friction for students with disabilities. Chatbots and tutoring systems can supplement, although they cannot replace, the work of overstretched teachers. The pandemic has shown what AI-mediated education can do in a crisis, and also what it cannot.

Equally clearly, the evidence shows that the framework is no substitute for substantive policy. Algorithmic bias, data privacy, infrastructure asymmetries, and the dilution of human contact are not peripheral concerns to be managed alongside the rollout. They are constitutive of whether the rollout succeeds. A policy that integrates AI into Indian education without taking these concerns seriously is likely to widen the gaps it claims to close.

The most useful contribution of the AI-in-education literature, perhaps, is to have changed how policy designers think about the relationship between institutions and the people who use them. Choices about how to teach are never neutral; some pedagogical architecture is always present, by design or by accident, and the question is whether it is

set up thoughtfully and with attention to the welfare of the students whose decisions and experiences it shapes. For India, where the stakes of getting this right are very high, the empirical and ethical work has only just begun. The next decade of research will determine whether AI becomes a genuine tool of inclusive education or another technology that promised more than it delivered.

Conflicts of Interest: The author declares no conflicts of interest.

References

- Banerjee, A. V., Duflo, E., Glennerster, R., & Kothari, D. (2010). Improving immunisation coverage in rural India: Clustered randomised controlled evaluation of immunisation campaigns with and without incentives. BMJ, 340, c2220.
- Chen, X., Xie, H., Zou, D., & Hwang, G.-J. (2020). Application and theory gaps during the rise of artificial intelligence in education. Computers and Education: Artificial Intelligence, 1, 100002.
- Datta, S., & Mullainathan, S. (2014). Behavioral design: A new approach to development policy. Review of Income and Wealth, 60(1), 7-35.
- Hausman, D. M., & Welch, B. (2010). Debate: To nudge or not to nudge. Journal of Political Philosophy, 18(1), 123-136.
- Holmes, W., Bialik, M., & Fadel, C. (2019). Artificial intelligence in education: Promises and implications for teaching and learning. Center for Curriculum Redesign.
- Hwang, G.-J., Xie, H., Wah, B. W., & Gašević, D. (2020). Vision, challenges, roles and research issues of artificial intelligence in education. Computers and Education: Artificial Intelligence, 1, 100001.
- Jaiswal, A., & Arun, C. J. (2021). Potential of artificial intelligence for transformation of the education system in India. International Journal of Education and Development Using Information and Communication Technology, 17(1), 142-158.
- Joshi, A., Vinay, M., & Bhaskar, P. (2023). Impact of coronavirus pandemic on the Indian education sector: Perspectives of teachers on online teaching and assessments. Interactive Technology and Smart Education, 18(2), 205-226.
- Luckin, R. (2017). Towards artificial intelligence-based assessment systems. Nature Human Behaviour, 1, Article 0028.
- Ministry of Education, Government of India. (2020). National Education Policy 2020. Government of India.
- NITI Aayog. (2018). National strategy for artificial intelligence: AIforAll. Government of India.
- Popenici, S. A. D., & Kerr, S. (2017). Exploring the impact of artificial intelligence on teaching and learning in higher education. Research and Practice in Technology Enhanced Learning, 12, Article 22.
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. American Psychologist, 55(1), 68–78.
- Sandu, N., & Gide, E. (2019, September). Adoption of AI-chatbots to enhance student learning experience in higher education in India. In 2019 18th International Conference on Information Technology Based Higher Education and Training (ITHET) (pp. 1–5). IEEE.
- Selwyn, N. (2019). Should robots replace teachers? AI and the future of education. Polity Press.
- Tapalova, O., & Zhiyenbayeva, N. (2022). Artificial intelligence in education: AIEd for personalised learning pathways. Electronic Journal of e-Learning, 20(5), 639-653.
- UNESCO. (2022). State of the education report for India 2022: Artificial intelligence in education. UNESCO New Delhi Office.
- Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education - where are the educators? International Journal of Educational Technology in Higher Education, 16, Article 39.