


POLICY PAPER

AI and K-12 Education in India and the Global South:

Opportunities, Risks, and Policy Directions



AUTHORS

Dr. Neethi S

Sriya Sridhar

AUGUST 2026

CENTRE FOR RESPONSIBLE AI
INDIAN INSTITUTE OF TECHNOLOGY MADRAS

Contents

1	Introduction	8
2	A Widening Global Divide	10
3	The Indian Context	12
3.1	Existing Policy Architecture and Initiatives	12
3.2	Current State	13
3.3	Pandemic disruption and persistent regional disparity	14
3.4	Ground-level evidence	14
3.5	AI in the Teacher Workflow	15
4	Where AI Can Help: Opportunities	17
4.1	Three AI approaches with direct evidence for foundational literacy and numeracy	18
4.2	Evidence from other Global South school systems	20
5	Risks and Open Questions	22
6	Toward a Policy Framework	25
6.1	For teachers and school administrators	26
6.2	An age-tiered access model for schools	27
6.3	For developers of AI systems used by children	31
6.4	For developers and researchers: aligning AI models to India's own vision of learning	32
6.5	For policymakers and regulators	34
7	Conclusion	37
—	References	38

Authors

Dr. Neethi S: Professor of Practice, Wadhvani School of Data Science & AI, Indian Institute of Technology, Madras

Sriya Sridhar: Senior Policy Analyst, Centre for Responsible AI, Indian Institute of Technology, Madras (CeRAI)

Acknowledgments

The authors express their deepest gratitude to Prof. Balaraman Ravindran, Head, Wadhvani School of Data Science and Artificial Intelligence (WSAI), and Centre for Responsible AI, (CeRAI) IIT Madras for valuable feedback and mentorship throughout the course of this work, and to Prof. Karthik Raman, Professor, WSAI, for valuable feedback and engagement with this paper. We further thank Omir Kumar, Policy Analyst, CeRAI, for assistance with copy editing and design.

This paper is published by the Centre for Responsible AI (CeRAI), Indian Institute of Technology Madras, Chennai – 600036, Tamil Nadu, India (www.cerai.iitm.ac.in). CeRAI is a multi-disciplinary, non-profit research centre positioned in the Global South, as one among the few global institutions that specialises in both technical and policy research to ensure and enable the responsible development and deployment of AI systems.

Errors in the paper, if any, are the sole responsibility of the authors. The copyright in this work vests with CeRAI.

Recommended citation: S, Neethi & Sridhar, S. (2026, August). AI and K-12 Education in India and the Global South: Opportunities, Risks, and Policy Directions. Centre for Responsible AI, Indian Institute of Technology Madras.

Executive Summary

Artificial Intelligence (AI) can support learning outcomes for K-12 learners, but only when it supports teaching and educators' judgment, rather than replacing it. In addition, it should be built around processes learning science already recognises as central to durable learning: diagnosing what a student knows, scaffolding just beyond their level, structuring retrieval practice, identifying specific misconceptions, and supporting progression toward mastery. Content personalisation alone does not meet this bar. The central question this paper asks is not whether India and the Global South should adopt AI in classrooms, but how to make it a genuine lever for equity, narrowing gaps in outcomes while protecting the autonomy of both students and teachers.

This paper examines the landscape of AI deployment in primary and secondary (K-12) education across India, informed by evidence from the broader Global South. Grounded in learning science and child-rights frameworks, our paper assesses how AI can act as an equalizer to narrow educational gaps—specifically in Foundational Literacy and Numeracy (FLN)—while establishing essential safeguards for student autonomy, data privacy, and pedagogical integrity.

The equity case is urgent because the divide is widening. Global adoption data shows a disparity in access to generative AI, with the Global North outpacing the Global South by a widening margin. This must be viewed within a broader digital divide, driven by differences in reliable electricity, connectivity, and digital skills, while acknowledging that increased AI adoption does not automatically translate to responsible AI adoption. Within India, around 5% of rural students had any access to online learning during the pandemic, and research on the Global South suggests generative AI tends to empower students who have already crossed a basic literacy threshold while leaving those who have not further behind. This is precisely the population India's own foundational learning gap is most concerned with. ASER 2024 shows nearly 45% of Class 5 students reading at a Class 2 level, but only 31% able to solve a basic division problem. This is a gap present in every state and rooted in the cumulative, hierarchical structure of mathematics itself. Seen in this context, generative AI tends to empower students who have already crossed a basic threshold of information literacy, while further disconnecting those who have not yet overcome their own digital divide.

The risks of AI in the K-12 context extend beyond privacy and safety, with the most consequential risks being hardest to spot. The paper situates India's approach within the international children's-rights consensus built around the UN Convention on the Rights of the Child, General Comment No. 25, and UNICEF's Policy Guidance on AI for Children. It organises practical risk clusters around communication, transparency, continuous monitoring, child-centred design, and teacher autonomy. We also identify a distinct category of pedagogical and cognitive risk that no safety audit catches: cognitive offloading, excessive scaffolding that prevents durable skill-building, answer dependence, and automation bias. These can only be detected through learning outcomes and behaviour over time, not through one-off technical checks. A more structural risk, is that general-purpose models trained on

English-language, Global North data are poorly equipped for Indian curricula, languages, and common misconceptions. We emphasise that AI in education has to be built around specific, well-diagnosed learning mechanisms rather than deployed as a generic tool.

Policy should be organised around a systems standard. A systems approach to incorporating AI in K-12 education involves integrating technology in ways that align with local realities, fostering institutional collaboration, and considering education systems, curriculum, teacher development, assessment, and governance as one connected system. This paper argues for adopting this system standard, and recommendations are addressed to four audiences. Interventions should be tested against the question of whether a given measure strengthen the system as a whole, starts from real classroom needs, and fits the Indian context, rather than being imported wholesale. Further, rather than viewing the teacher and student in isolation, a systems perspective calls for considering the impact of technological integration on the community, society and educational values at large.

Whether India and the rest of the Global South realise AI's potential in K-12 education will depend less on the sophistication of the technology, and more on these policy choices around infrastructure, teacher support, child rights, and equity of access. This paper's overarching recommendation therefore, is that AI should be introduced only where there is a clearly identified educational need, a sound pedagogical rationale, appropriate teacher mediation, evidence of benefit, and adequate safeguards for equity, child rights, and data protection. Solutions should be adopted not because the technology exists, but because they demonstrably serves the child in front of it.

The key arguments and recommendations we make in this paper are as follows:

1. Key Arguments & Opportunities

- **Scaffolding, Not Substitution:** AI must supplement and scaffold human teaching, not replace critical thinking or educators. AI is best deployed across four stages of the learning pathway: Diagnosis (evaluating student gaps), Selection (matching activities to level rather than grade), Guided Practice (adaptive feedback), and Reassessment/Progression.
- **Workflow based differentiation:** AI should streamline teacher administrative burdens, progress tracking, and lesson planning, while keeping high-stakes assessments firmly under human judgment.
- **Differentiated Evidence Profiles:** Opportunities to integrate AI in K-12 education need to be situated within the context of different types of AI systems. Each solve different problems, put the teacher in a different position, and rest on evidence bases of very different maturity:
 - Adaptive Learning Systems: Strong localized evidence for math remediation, providing gains equivalent to nearly a full grade level in months.
 - AI-Assisted Diagnostics: Strong field evidence for oral reading fluency and screening.

- Generative/Conversational Tutors: High potential (e.g., rapid gains shown in short pilots in Nigeria) but carries higher risks regarding accuracy, automation bias, and cognitive over-reliance.

2. Primary Risks & Governance Challenges

- **Cognitive & Pedagogical Risks:** Excessive scaffolding, cognitive offloading, answer dependence, and automation bias can impair durable learning and critical thinking.
- **Model & Data Risks:** General-purpose AI models trained predominantly on Global North, English-language datasets often fail to address regional language nuances, local contexts, or specific Indian curricula.
- **Child Rights & Safety:** Conversational AI models risk creating emotional dependency or exposing minors to non-age-appropriate content without proper guardrails. Solutions should prioritise childrens' data protection and privacy rights.

3. Policy Recommendations for Different Stakeholders

FOR TEACHERS & SCHOOL ADMINISTRATORS

- **Keep Humans-in-the-Loop:** Retain human judgment for grading, emotional support, and instructional selection across three critical loops: moment-to-moment teaching, planning/reflection, and tool selection.
- **Structured Capacity Building:** Adopt UNESCO's AI Competency Framework for Teachers to build capability across five dimensions (Human-Centred Mindset, Ethics, Foundations, Pedagogy, Professional Learning) rather than relying on one-off technical workshops. Teacher capacity building should thus be structured around progression through these five dimensions, with particular attention to the dimensions of ethics and pedagogy, rather than a tool-based notion of AI literacy. AI Samarth offers a working framework for the same.

FOR SCHOOL SYSTEMS: AN AGE-TIERED ACCESS MODEL

- **Primary (Classes 1–5):** Teacher-mediated access only. Focus on FLN diagnostics and adaptive support. Student exposure to AI concepts remains unplugged (computational thinking) per CBSE frameworks.
- **Middle School (Classes 6–8):** Supervised access. Introduction to approved no-code tools, gamified learning, and foundational AI literacy (bias, data privacy, verifying output).
- **Secondary (Classes 9–12):** Independent, regulated access. Use of AI for research, coding, and project work. Clear academic integrity rules classifying assignments as AI-Prohibited, Permitted-with-Disclosure, or AI-Integrated.

FOR DEVELOPERS & RESEARCHERS

- **Child-Centred & Strengths-Based Design:** Implement Child Rights Impact Assessments (CRIA), transparently disclose non-human AI identity, and build automated escalation paths to human adults when distress is detected.
- **Safety testing:** Developers must conduct extensive pre-deployment adversarial testing ("red-teaming") using child-safety matrices to catch hallucinations, implicit bias, and toxic content tailored to minor age groups.
- **Pedagogical Alignment:** Design inspectable, explainable models tailored to Indian languages (leveraging tools like Bhashini) and aligned with local learning progressions.
- **Local Context & Multilinguality:** Models must move beyond Global North benchmarks, integrating regional language datasets (e.g., leveraging Indian DPIs like Bhashini). Systems should accommodate code-switching, local idioms, and context-specific math/science curricula.
- **Architectural & Safety Delineation: Student-Facing vs. Teacher-Facing Systems**
 - Student-Facing Systems (High Risk): Must be strictly guarded against answer-giving and over-scaffolding by enforcing Socratic questioning, while incorporating real-time child-safety classification, distress escalation protocols, transparent non-human identity disclaimers, and rigorous age-appropriate guardrails.
 - Teacher-Facing Systems (Co-Pilot Role): Should focus on streamlining lesson planning and administrative tasks while maintaining full inspectability and teacher override capability, ensuring educators can inspect diagnostic reasoning and retain complete authority over generated content.

FOR POLICYMAKERS & REGULATORS

- **Strict Legal Enforcement:** Enforcing existing applicable laws such as the Digital Personal Data Protection (DPDP) Act 2023 and DPDP Rules 2025, mandating verifiable parental consent for children's data processing and prioritising confidentiality and children's privacy.
- **Evidence Standards & Needs-Based Procurement:** Mandate rigorous, independent impact evaluations before scaling edtech platforms nationally. Use collective, needs-first procurement through Digital Public Infrastructure (DPI) frameworks.
- **Interoperability & DPG Status:** Promote AI tools as Digital Public Goods (DPGs) with open licensing to avoid vendor lock-in and foster sustainable, sovereign educational technology.

1 Introduction

K-12 education (i.e., education from kindergarten until the end of high school), is being reshaped by the same technological forces that are reshaping the rest of society, in India and worldwide. Children today are growing up immersed in an increasingly AI-mediated world, and pedagogy needs to keep pace with the environment they will eventually work and live in. However, this transition will not be uniform: how AI enters classrooms, and the impacts of its design and deployment, is likely to look very different across the Global North and the Global South.

Artificial Intelligence has often been framed as a potential equalizer and as a way to narrow persistent gaps in access to quality K-12 education across the Global South. While there are significant opportunities to improve educational outcomes, this claim must also be interrogated critically, to ensure that the lifecycle of AI deployment in this sector is implemented responsibly. While AI can support learning outcomes, it should not be mistaken for a substitute for critical thinking. As a solution, it works best when it scaffolds existing teaching and learning processes rather than replacing them outright. The central policy question, then, is not whether to adopt AI in classrooms, but how to make it a genuine lever for equity. It should narrow gaps in K-12 outcomes while protecting the autonomy of both students and teachers.

Because the end users of many of these systems are children, the frame of responsibility applied to them has to be grounded in developmental and learning science principles. Any technology that children interact with directly needs to be designed and governed in a way that protects the cognitive, social, and emotional skills that are central to healthy development, rather than quietly displacing them. This involves distinguishing between AI that merely delivers personalised content and AI that supports evidence-based processes such as diagnosis of prior knowledge, scaffolding, retrieval and practice, formative feedback, misconception identification, metacognition and progression towards mastery.

AI can be developed to support processes that learning science has long identified as central to durable learning: diagnosing what a student already knows before instruction begins, scaffolding a task just beyond a student's current level, structuring retrieval practice, identifying a specific misconception, delivering formative feedback a student can act on, and prompting the kind of metacognitive reflection that helps a learner track their own progression toward mastery. The strongest evidence discussed later in this paper, on foundational literacy and numeracy specifically, comes from tools built around these deeper instructional functions rather than from content personalisation alone.

It is also worth noting at the outset, that AI's arrival in the classroom marks a qualitative shift. For most of the history of educational technology, machines were passive: they delivered content, but did not respond to it. Generative AI systems are different. They are interactive agents that can converse, adapt, and respond in ways that let them act as tutors, study companions, or even stand-ins for instructors. Some see in this the outline of a future where every student, teacher, school leader, and policymaker has

access to a personalized AI assistant, provided the underlying infrastructure is in place to support it: connectivity, devices, and teacher capacity, among other things [1]. Whether that infrastructure exists, and for whom, is precisely where the Global North/South divide becomes most consequential.

AI applications in K-12 education fall into a few distinct categories, each with a different evidence base and risk profile. Diagnostic and assessment AI identifies what a student already knows and where specific gaps or misconceptions lie, as in automatic speech recognition tools that screen for reading difficulty from a short recording. Adaptive or personalised learning systems adjust content, sequencing, or difficulty to an individual student's demonstrated level. AI-supported teacher tools assist with lesson planning, differentiated instruction, and administrative work, but leave instructional judgment with the teacher. Conversational or generative AI tutors hold an open-ended dialogue with a student, and carry a different set of risks around accuracy, emotional dependence, and age-appropriateness than the other categories. System-level analytics and decision-support tools operate above the level of the individual student entirely, tracking learning outcomes across districts and schools to guide policy and resource allocation. Distinguishing between these categories matters because a policy or safety concern that applies to one, such as the child-safety questions raised by conversational tutors, does not automatically apply to the others.

At the outset, we distinguish between two different senses in which "AI in K-12 education" is used. Most of what follows treats AI as a tool for teaching and learning: the diagnostic, adaptive, tutoring, and system-level applications described in the five categories above, addressing how AI can support instruction, assessment, and administration. A smaller part of the paper, principally Section 6.2, treats AI differently: as a subject of curriculum content, i.e, what children are formally taught about AI itself. These are different questions, with different stakeholders, evidence bases, and risks. The age-tiered access model developed in Section 6.2 is built primarily around the first sense, i.e, a student's access to AI as an instructional tool. The discussion of curricula such as CBSE's Computational Thinking and Artificial Intelligence framework, answers the second. This paper's central contribution is to the first question. Where it engages the second, this is mentioned explicitly.

In this report, therefore, we aim to understand and document the rapidly developing global landscape of deploying AI solutions in K-12 education, examine potential opportunities and risks, and accordingly suggest policy recommendations suitable for the Indian context. We hope that our recommendations further the conversation on developing guiding policy frameworks to support the improvement of educational outcomes for India's children, while ensuring that these systems are developed responsibly.

2 A Widening Global Divide

The gap in who benefits from AI more broadly, not just in K-12 education, is already visible in adoption data. Microsoft's AI Economy Institute, which tracks generative AI usage through aggregated and anonymized telemetry, found that by the second half of 2025, roughly 24.7% of the working-age population in the Global North had used a generative AI tool, compared with 14.1% in the Global South. This gap widened over the course of the year rather than closing, as adoption in wealthier regions grew nearly twice as fast [2]. By the first quarter of 2026, the gap had widened further still, reaching 12.1 percentage points. Microsoft attributes much of the divergence to differences in reliable electricity, connectivity, and digital skills, rather than to differences in interest or willingness to use the tools [2].

A parallel pattern shows up in institutional readiness. A 2025 UNESCO survey of nearly 400 university-level Chairs and networks found that about 70% of institutions in Europe and North America already had, or were actively developing, formal guidance on AI use, compared with 45% in Latin America and the Caribbean [3]. That particular figure describes higher education rather than K-12 systems specifically, but it is a reasonable proxy for the broader institutional gap in AI governance capacity between richer and poorer regions. K-12 systems in the Global South are likely to face this gap even more acutely, given tighter budgets and thinner administrative bandwidth.

In the Indian context, a National Sample Survey Office study conducted during the pandemic found that only around 5% of students in rural India had access to any online learning resources at all, at exactly the moment when access to digital learning determined whether a child's education continued or stalled [4]. The same rural-urban divide that shaped access to basic online learning during COVID-19 is likely, absent deliberate intervention, to shape access to AI-enabled tools as well. Beyond the infrastructural limitations posed by uneven internet connectivity, India is also faced with more basic infrastructural issues such as reliable supply of electricity especially in rural areas, and compatibility of low-end devices with e-learning materials (for example, the requirement to view a document on a smartphone), offline/low-bandwidth operation. Specific to AI, there are issues such as the lack of local-language speech and language models, supporting cloud/edge infrastructure, teacher support capacity in implementing AI solutions within the classroom, data infrastructure and interoperability with existing public education platforms.

Research specific to the Global South has argued that generative AI tends to empower students who have already crossed a basic threshold of information literacy, while further disconnecting those who have not yet overcome their own digital divide. The same tool that helps a confident reader ask better questions can leave a student who has not yet acquired basic reading fluency further behind, since the tool assumes a baseline of literacy. That is precisely the population this paper is most concerned with in the Indian context: children who have not yet cleared the FLN bar described in Section 3, for whom generic AI tools risk widening the gap between the already-capable and the still-struggling rather than narrowing it [5].

This provides a stronger pedagogical rationale for AI-enabled FLN diagnostics and level-based intervention: For children who have not acquired foundational literacy or numeracy, AI systems should first establish the learner's current competencies and misconceptions, and only then adapt instruction and practice.

Rather than being an argument against incorporating AI solutions in K-12 education in India, it is a reminder that closing the equity gap in who can access and benefit from these tools deserves policy attention alongside considerations of technical capabilities and developments. Without deliberate safeguards, digitalization and AI in education risk widening inequality. Therefore, it should be anchored explicitly in learners' right to education.

3 The Indian Context

3.1 Existing Policy Architecture and Initiatives

India's National Education Policy (NEP) 2020 named foundational literacy and numeracy (FLN) the cornerstone of all future learning [6]. Over the following five years, that intent was translated into a concrete implementation architecture. The NIPUN Bharat Mission, (National Initiative for Proficiency in Reading with Understanding and Numeracy), was launched by the Ministry of Education in 2021 as a flagship push for universal FLN by the 2026-27 academic year [7]. It functions through a five-tier structure spanning national, state, district, block, and school levels, intended to keep the mission accountable and responsive to local conditions. It combines teacher capacity-building, learner-friendly materials, periodic assessment, community engagement, and technology-enabled tracking of learning outcomes [7].

The National Curriculum Framework for the Foundational Stage (NCF-FS, 2022) is India's first integrated curriculum framework for children aged 3–8, developed by NCERT under NEP 2020's 5+3+3+4 structural reform. It covers the newly defined "Foundational Stage," spanning three years of pre-school and Grades 1–2, and is grounded in neuroscience research showing that over 85% of brain development occurs by age 6, alongside evidence that children's learning trajectories only begin to converge around age 8. The framework places play, defined broadly to include free, guided, and structured forms, at the centre of pedagogy, content organisation, and assessment, rather than treating it as separate from 'real' instruction. It draws simultaneously on global developmental theorists and Indian intellectual traditions, and derives a formal chain from Aims of Education to Curricular Goals to Competencies to Learning Outcomes. The National Curriculum Framework (NCF) 2023 followed as a further structural reform, aligning curriculum design with NEP's broader goals [8]. Through these initiatives, FLN outcomes have started to trend upwards, a trend confirmed by both government monitoring and independent surveys such as the Annual Status of Education Report (ASER) [9].

DIKSHA, the Ministry of Education's national digital content platform, has been live since 2017 and is now used across 36 states and union territories. Since 2023, the National e-Governance Division has been working to integrate Personalized Adaptive Learning into it, so that a calculation mistake in a Class 9 geometry problem, for instance, can be flagged and looped back to a simpler explanatory video automatically [10]. Separately, Vidya Samiksha Kendras (AI-powered dashboards that track learning data across districts and schools) were pioneered in Gujarat and have since been adopted by several other states as a governance tool for identifying where interventions are most needed [10].

Other initiatives include the UDISE+ (Unified District Information System for Education Plus is an educational management information system) [11], NDEAR, a Digital Public Infrastructure (DPI) blueprint for India's educational ecosystem, and the expansion of multilingual model capabilities through Bhashini, the National Language Translation Mission launched in 2022 [12]. Andhra Pradesh has gone further still, partnering with private ed-tech providers such as Embibe, ConveGenius, and Mindspark

specifically to train teachers in using AI-generated analytics for remedial learning [10].

The Central Board of Secondary Education's Computational Thinking and Artificial Intelligence curriculum, developed under the Ministry of Education's guidance and rolling out for Classes 3 to 8 from the 2026-27 academic year, is the first system-wide curriculum to formally sequence computational thinking and AI content across the primary and middle stages nationally [35]. Separately, AI Samarth, launched in September 2025 by the Central Square Foundation in partnership with IIT Madras's Wadhvani School of Data Science and AI, is an AI literacy curriculum framework aimed at middle and high school students, teachers, and parents in government and affordable private schools, targeting over 5 million learners [34]. Both are discussed in more detail in Section 6, where they inform this paper's specific policy recommendations. Bodhan AI, the Centre of Excellence in AI for Education is an IIT Madras-incubated initiative supported by the Ministry of Education, building the foundational AI layer for India's education system, enabling equitable learning through personalized practice, continuous feedback, and access to high-quality learning resources in multilingual environments.

The above detailed landscape of initiatives in the space of educational technology in India shows that the approach is moving towards emerging as a form of DPI, rather than as an amalgamation of individual platforms.

3.2 Current State

Despite this progress, numeracy continues to trail literacy nationally. ASER 2024 found that while close to 45% of Class 5 students could read at least at a Class 2 level, only around 31% could correctly solve a basic division problem, and no state in the country reports higher numeracy outcomes than literacy outcomes [9].

Language learning tolerates partial understanding: a child with gaps in vocabulary can still follow much of a story. Mathematics does not work the same way. It is inherently cumulative and hierarchical, so a missed foundational concept, such as place value, can make everything built on top of it, including addition, decimals, and fractions, effectively inaccessible later on. Left unaddressed, small early gaps compound into large ones.

Curriculum-paced, syllabus-driven teaching tends to make this worse, because it advances the class regardless of whether individual students have actually grasped what came before. Pratham's Teaching at the Right Level (TaRL) approach, developed and refined over two decades in partnership with researchers at the Abdul Latif Jameel Poverty Action Lab (J-PAL), offers a well-evidenced alternative: group and teach children according to their demonstrated learning level rather than their grade, using low-cost, activity-based instruction, and re-assess regularly so groups can shift as students progress [13].

A related and important finding, from a large 2025 study led by MIT and Harvard researchers with support from J-PAL, is that the disconnect between classroom mathematics and real-world mathematics

runs in both directions. Surveying over 1,400 children working in markets in Kolkata and Delhi, the researchers found that these children could perform complex mental arithmetic fluently in the context of a real transaction, but struggled with equivalent problems posed in the abstract, symbolic form typically used in schools. Conversely, children with no market experience, who were only taught arithmetic in the classroom, did better on abstract problems but performed far worse when the same arithmetic was embedded in a realistic, concrete scenario [14]. Neither group's skills transferred cleanly to the other setting, which is a strong signal that Indian mathematics pedagogy needs to more deliberately connect abstract procedures to the situations in which students will actually use them.

Taken together, this evidence points toward a specific pedagogical agenda: extend the child-friendly, activity-based methods already used for early FLN into higher-level content; keep instruction aligned to a child's actual learning level rather than a fixed grade-based syllabus; and deliberately embed every day, real-life problem-solving into numeracy instruction so that classroom learning and lived experience reinforce rather than bypass each other. This makes it especially important to align AI deployment in this sector with learning progressions in FLN. AI-enabled diagnosis is educationally useful when assessment results can be mapped to meaningful competencies, prerequisite relationships and next-step learning opportunities. This would also provide a conceptual foundation for adaptive practice and personalised intervention.

3.3 Pandemic disruption and persistent regional disparity

The COVID-19 pandemic widened foundational learning deficits considerably, and did so unevenly. Countries with wide digital divides, India among them, were hit especially hard, since remote learning was simply not a realistic option for many students. There has been some recovery in recent years. For example, ASER 2024 recorded a significant rebound in both reading and arithmetic compared with the pandemic-era low points of 2022, in some cases exceeding pre-pandemic levels for the first time since ASER began tracking in 2005 [9]. But the gap between the best- and worst-performing states widened rather than narrowed between 2022 and 2024, even as the national average improved [9].

FLN deficits are largest among poorer, rural, and otherwise marginalized children, and inequities in foundational learning correlate strongly with gender, geography, and socioeconomic status. Children who fall behind in FLN are considerably more likely to eventually drop out of school and less likely to access decent work later in life, which turns an early learning gap into a lifelong one and makes closing it a meaningful contribution. State-level variation illustrates the same dynamic: states such as Kerala and Himachal Pradesh continue to lead on literacy and skill acquisition, while states such as Jharkhand and Bihar continue to lag, with children from low-income and marginalized households facing compounding disadvantages in access to quality schooling, digital resources, and experienced teachers [15].

3.4 Ground-level evidence

Much of the discussion so far has drawn on national policy documents and aggregate survey data. A 2025 qualitative study offers ground-level insight. Researchers at BITS Pilani conducted 23 semi-structured interviews with student volunteer teachers working in rural schools in Rajasthan and Delhi, most through the National Service Scheme and the Nirmaan Organization, to understand how people actually teaching in under-resourced classrooms think about large language models [16]. These volunteers are not full-time teachers, but their dual exposure to urban AI tools and rural classroom realities makes them a useful bridge population for understanding how AI-related aspirations meet ground conditions.

Their reported barriers further support structural gaps. Fifteen of the 23 volunteers named unreliable internet, limited devices, and underfunded schools as the primary obstacle to any AI-based learning tool, a finding that lines up closely with the rural access data in Section 2. Thirteen cited a near-total lack of AI training among government-school teachers. Language emerged as a recurring frustration: most available AI tools operate in English or Hindi, while rural students are frequently more comfortable in regional languages, and NCERT textbook terminology in English compounds the difficulty, reinforcing for treating multilingual infrastructure as a necessity [12,16].

12 of the 23 reported that students were performing well below grade level, in one case describing first-graders functioning at a preschool level, and noted that AI-driven personalization is only as useful as a student's foundational skills allow it to be. An AI tutor that assumes baseline reading or numeracy proficiency is of limited use to a child who has not yet acquired it, which is precisely the concern that motivated TaRL's grade-independent, level-based approach to instruction in the first place [13,16].

Volunteers also described social friction. Nine reported that persuading parents of AI's value was an uphill task, since many rural families see it as a distraction from, rather than a support for, traditional learning, and several volunteers described government-school teachers treating tools like ChatGPT with outright suspicion, as a shortcut for cheating.

Despite these frictions, and despite worries that heavy AI use could blunt students' independent thinking and creativity, the volunteers converged on a clear, consistent conclusion: AI should be a supplementary, teacher-supported layer within existing instruction, never a replacement for it, echoing the human-in-the-loop principle that runs through this paper's later policy recommendations [1,16].

Ground-level accounts such as these confirm that the opportunity for AI deployment in this sector is real, but need to be supplemented with the right infrastructure, teacher training, and address language gaps.

3.5 AI in the Teacher Workflow

It is also important to assess which parts of a teacher's actual workflow AI can plausibly support, since the answer varies considerably by task (see the below diagram for a visual representation). Diagnosis, identifying what a student already knows and where a specific misconception lies, is where the evidence is strongest [17,18]. Grouping students by demonstrated level rather than grade, the core operational task

behind TaRL, is a natural extension of this diagnostic function, and is precisely the kind of continuous, individual-level tracking that annual surveys like ASER cannot provide but a classroom tool plausibly could.

Lesson planning and practice generation, producing worksheets, examples, and practice questions pitched at a given level, are comparatively low-risk uses well suited to AI support. Feedback generation is more contested: useful for routine, low-stakes correction, but it needs clear limits around higher-stakes assessment, where a teacher's judgment should remain central. Progress monitoring over time is where a tool could meaningfully reduce a teacher's manual record-keeping burden, provided it feeds back into decisions the teacher still makes, echoing the human-in-the-loop principle discussed later in this paper [1]. Administrative reporting, the task furthest from direct instruction, is already being addressed at scale through Vidya Samiksha Kendra's district-level dashboards, suggesting this is a workflow layer better handled at the system level than left to individual teachers [10]. Different parts of a teacher's workflow will need to be assessed differently in terms of incorporating AI. For example, while diagnosis and grouping already have real evidence behind them, while others, like feedback on high-stakes work, still need a teacher's judgment in the loop.

FIG. 01

AI readiness across the teacher workflow

01	Diagnosis	Strong evidence	STRONG EVIDENCE
02	Grouping	Level-based groups	STRONG EVIDENCE
03	Lesson plans	Low risk, ready now	READY NOW, LOW RISK
04	Practice items	Low risk, ready now	READY NOW, LOW RISK
05	Feedback	Needs human judgment	NEEDS HUMAN JUDGMENT
06	Progress monitoring	Human-in-the-loop	READY NOW, LOW RISK
07	Admin reporting	System-level (VSK)	SYSTEM-LEVEL

Tasks flow roughly diagnosis to reporting, but readiness for AI support varies far more than the sequence suggests.

4 Where AI Can Help: Opportunities

Given this landscape, with real but uneven progress on FLN, a specific and well-diagnosed numeracy problem, and large regional and socioeconomic disparities, there is a real opening for AI to support Indian K-12 education, provided it is aimed at these specific gaps.

It is useful to visualise integrating AI in the K-12 ecosystem as stages in a single instructional pathway: diagnosing what a student currently knows, selecting a learning activity appropriate to that diagnosis, guiding the student through practice with feedback, and then reassessing to decide whether the student has progressed enough to move on. AI does not replace any of these four stages, but it can support each of them differently, with a different role for the teacher and a different level of evidence behind it.

Diagnosis of learning needs. The starting point of the pathway is establishing what a student already knows, where a specific gap or misconception lies, and what they bring to their learning beyond a simple right-or-wrong score. This is where AI-enabled assessment can, in principle, give teachers richer, faster information across a wider range of skills than periodic standardized testing allows, provided the goal remains informing a human decision rather than replacing it [1]. Shifting from the deficit-only model that most educational technology has historically used, toward one that also identifies a student's existing strengths, matters especially for students with disabilities and English-language learners, who are often the worst served when diagnosis stops at what a student got wrong [1]. None of this removes the underlying risk of algorithmic bias, which is inherent to how these systems are built and trained, and which requires deliberate, ongoing correction [1].

Selection of an appropriate learning activity. Once a diagnosis exists, the next step is matching a student to the right activity at the right level, rather than advancing every student through the same fixed, grade-locked sequence regardless of whether they are ready for it. This is the operational core of Pratham's Teaching at the Right Level approach, and AI-enabled tools can support the same logic at an individual level: grouping students by demonstrated level, adapting materials for neurodiverse learners, and translating and localizing content into different Indian languages so that the activity a student receives actually fits their context [1,13]. Bodhan AI's model of keeping teachers involved in selecting and evaluating the tools used in their own classrooms is directly relevant here: activity selection works best when a teacher, not just an algorithm, has a say in what gets used [19].

Guided practice and feedback. This is the stage most commonly referred to in the context of AI tutoring, and it is also the stage with the widest range of evidence quality. Mindspark's adaptive system for mathematics, which detects specific misconceptions and adjusts question difficulty accordingly, produced gains equivalent to almost a full grade level in 4.5 months, with the largest relative gains for the weakest students [18]. A GPT-4 tutoring program evaluated in Nigeria produced gains equivalent to 1.5 to 2 years of schooling in six weeks [20], and Kenya's EIDU program, pairing personalized digital learning with structured pedagogy, produced roughly 0.8 additional years of learning [21]. These results should be

read against a more modest global baseline: a meta-analysis of intelligent tutoring systems in K-12 mathematics found next to no overall effect across 26 studies [22], a reminder that guided practice works when it is built around a specific, well-diagnosed learning mechanism and implemented consistently, not because personalization is inherently effective. Throughout this stage, AI assistants can also lighten a teacher's load directly, helping plan lessons, generate practice materials in multiple languages, and surface the information needed to work closely with individual students, while keeping the teacher informed, connected, and in control of substantive instructional decisions [1,19].

Reassessment and progression. The pathway closes, and restarts, with reassessment, i.e., deciding whether a student has learned enough of the current activity to progress to the next one, or needs to return to an earlier stage of diagnosis. This is where AI's role shifts from the individual classroom to the system level. India's Vidya Samiksha Kendras already track learning outcomes across districts and schools using this kind of data, giving administrators a live view of where progression is and is not happening, at a scale no annual survey could match [10]. At the classroom level, the same reassessment logic is what allows a system to loop a student back into diagnosis rather than simply advancing them by age or grade. This is the central lesson of the numeracy gap discussed earlier: without this loop, gaps in learning that form early widen over time [14].

FIG. 02

The AI-supported learning pathway

STAGE 1	STAGE 2	STAGE 3	STAGE 4
Diagnosis of learning needs	Selection of learning activity	Guided practice and feedback	Reassessment and progression
What does this student already know, and where is the specific gap?	Match the activity and level to the diagnosis, not to the grade.	Adaptive tutoring, practice, and teacher-supported correction.	Has the student learned enough to move on, or loop back to diagnosis?

← Cycle repeats as the student progresses

Each stage puts AI in a different role and rests on a different level of evidence. See Section 4.1 for the evidence behind Stages 1 to 3 specifically.

4.1 Three AI approaches with direct evidence for foundational literacy and numeracy

The opportunities described above need to be situated within the context of different types of AI systems. Adaptive learning systems, AI-assisted diagnosis, and generative or conversational tutors solve different problems, put the teacher in a different position, and rest on evidence bases of very different maturity.

Adaptive learning systems. These systems solve a pacing and misconception problem: a fixed, grade-locked syllabus advances regardless of whether a student has actually understood the previous concept,

while an adaptive system instead adjusts question difficulty and content to what a given student has demonstrated they know. Mindspark, India's flagship example, uses this approach for mathematics, detecting specific misconceptions and adjusting accordingly. The teacher's role here is one of facilitation and monitoring rather than direct instruction during system use: teachers typically supervise a dedicated lab session and use the system's reports to inform their own teaching, but the moment-to-moment pedagogical decisions are made by the software. This is also the category with the most rigorous evidence in India specifically: a landmark randomized evaluation found gains equivalent to almost a full grade level in mathematics in 4.5 months, and a separate three-year trial in Rajasthan government schools replicated smaller but still significant gains [18]. However, that evidence should not be generalized to adaptive learning as a category. A meta-analysis of intelligent tutoring systems in K-12 mathematics more broadly found next to no average effect across 26 studies [22], which suggests that positive results are not inherent to adaptive software, and may reflect a specific, well-implemented system.

AI-assisted diagnosis. This is a narrower and, in some ways, more foundational category: rather than teaching directly, these tools identify what a student already knows and where a specific gap lies, leaving the response to that diagnosis in the teacher's hands. RightToRead's reading assessment component and automatic speech recognition tools that screen for oral reading difficulty both belong here. The teacher's role is central by design: RightToRead is delivered by regular government teachers working through the existing textbook, and a diagnostic signal on its own does nothing without a teacher, or a TaRL-style regrouping process, acting on it. Evidence maturity varies sharply within this category. RightToRead has the strongest field evidence in this paper, a randomized trial across 5,000 government schools and over 1 million students with a two-year follow-up showing gains held over time [17]. Speech-recognition-based screening for reading difficulty is considerably newer, with promising accuracy in research settings but little evidence yet at the scale or duration of RightToRead or Mindspark [23].

Generative and conversational tutors. This category addresses a different problem again: open-ended, dialogic support across a broader range of questions than a narrow diagnostic or adaptive tool can handle, closer to what a human tutor does in an unstructured conversation. A GPT-4-based virtual tutoring program evaluated in Nigeria is the clearest evidence available, producing learning gains equivalent to 1.5 to 2 years of typical schooling in six weeks [20]. But this evidence should be read more cautiously than the other two categories for several reasons. The Nigeria trial paired the tool with active teacher guidance rather than deploying it autonomously. The trial ran for six weeks on secondary students in English-medium classrooms, and has not yet been replicated at the scale, duration, or primary-FLN focus of RightToRead or Mindspark. Because conversational tutors interact far more openly with a child than an adaptive math system or a reading-fluency screen, they carry the pedagogical and cognitive risks discussed in Section 5, i.e., hallucinated or incorrect explanations, reinforcement of a student's existing misconception, and automation bias, in a way that narrower, more constrained tools largely do not.

Read together, these three approaches point to a useful principle for Indian policymakers and developers:

the type of AI that matters is not a generic chatbot layered on top of a curriculum, but a technology built around a specific, well-diagnosed learning mechanism, whether that is multi-sensory letter-sound acquisition, adaptive misconception-correction in mathematics, or automated fluency screening, each matched to a specific point in a child's FLN development.

4.2 Evidence from other Global South school systems

India is not alone in testing these ideas, and the strongest evidence so far comes from a small set of rigorously evaluated programs across Africa, Latin America, and Southeast Asia. In Sub-Saharan Africa, a World Bank randomized controlled trial in Edo State, Nigeria, paired GPT-4-powered virtual tutoring with teacher guidance for first-year secondary students over six weeks and found learning gains equivalent to roughly 1.5 to 2 years of typical schooling [20]. In Kenya, a peer-reviewed randomized trial across 291 government pre-primary schools and nearly 2,000 learners found that EIDU's combination of digital personalized learning and structured pedagogy led to learning gains equivalent to roughly 0.8 extra years of learning [21]. In Rajasthan, state authorities have used AI-powered tools to score paper-based worksheets for 4.5 million students, freeing up teacher time for instruction, and World Bank-supported pilots are underway in Cote d'Ivoire, the Gambia, Mali, and Ghana [24]. A November 2025 gathering of over 100 education and technology leaders in Nairobi, convened by the World Bank and EdTech Hub, converged on similar priorities for the region: teacher support, personalization, and assessment [24].

Southeast Asia illustrates a version of the same divide described in Section 2. Vietnam, Indonesia, and Singapore show real diversity of AI-in-education initiatives, while EdTech Hub's regional mapping found no comparable initiatives at all in Lao PDR, Myanmar, Brunei, or Cambodia [26]. Indonesia's CoLearn platform, which offers affordable AI-supported tutoring to substitute for the country's expensive private tutoring culture, is a useful illustration of AI addressing a specific, well-diagnosed access gap, in much the same spirit as the numeracy-focused interventions described in Section 3 [26].

A recent open-access systems analysis from Harvard's Global Education Innovation Initiative draws a similar lesson across all of these cases: AI's educational value in the Global South depends on treating curriculum, teacher development, assessment, and system governance as one connected system, rather than adopting AI tools piecemeal and hoping for systemic change to follow. The same analysis explicitly warns that educational technology has disappointed before, and that marginal, tool-level improvements should not be mistaken for the deeper systemic transformation India's own foundational learning gaps actually require [27].

Together, this cross-regional evidence supports two conclusions relevant to India:

First, evidence demonstrates that AI-enabled interventions can produce meaningful learning gains in Global South contexts, but effects appear highly contingent on implementation conditions, pedagogical design and teacher involvement.

Second, the biggest gap across nearly every region studied is not technical capability but evaluation rigor, which argues for treating rigorous, independent evaluation as a precondition for scaling AI in Indian classrooms [13,21,25].

5 Risks and Open Questions

The same features that make generative and conversational AI powerful in education, such as its responsiveness, its capacity to sound empathetic, its ability to hold a sustained back-and-forth with a child, are also where the risks concentrate. These risks apply with particular force to children. Therefore, when considering AI in K-12 education, it must be kept in mind that AI can automate processes in service of goals that humans set, but it should not be setting those goals itself. The vision of good teaching and learning will come from educators' understanding of their students' actual strengths and needs, with AI treated as an instrument for realizing that vision rather than an independent source of educational priorities [1].

These risks should be situated within the broader international framework. The UN Convention on the Rights of the Child, ratified by nearly every country in the world, sets out four general principles that anchor how any technology aimed at children should be judged: non-discrimination, the best interests of the child as a primary consideration, the right to life, survival, and development, and the right to be heard. In 2021, the UN Committee on the Rights of the Child adopted General Comment No. 25, its first formal clarification that all of these rights apply fully in digital environments, developed with direct input from more than 700 children across 27 countries, and including provisions addressed specifically to the right to education [28]. UNICEF's Policy Guidance on AI for Children, now in its third edition, translates these principles into concrete requirements for anyone building AI systems that children use directly: supporting children's development and wellbeing, ensuring inclusion, prioritizing fairness and non-discrimination, protecting data and privacy, ensuring safety, and providing transparency, explainability, and accountability, among others [29].

A 2024 study by the Vidhi Centre for Legal Policy and the Centre for Responsible AI (CeRAI) at IIT Madras, proposed a participatory approach to the design, development, deployment and adoption of AI. This involves identifying stakeholders, mapping their concerns across the various stages of the AI lifecycle, and organising their inputs based on the axes of power, harm, legitimacy and urgency [30]. Based on this framing, a participatory approach can be applied to AI deployment for K-12 education involving children, caregivers, and domain experts, organized these risks into several practical clusters, as follows:

- **Communication and understanding:** whether an AI system can appropriately interpret children's non-literal, playful, or emotionally loaded communication, and whether there are clear protocols preventing responses that could harm a child's wellbeing.
- **Transparency and authenticity:** whether it is clear to a child that they are interacting with a machine and not a person, and whether the system actively avoids letting children form an inflated sense of its empathy or understanding. This includes whether it reminds children, and escalates to a human teacher or caregiver, when a child shows signs of distress.

- **Continuous monitoring and improvement:** whether the system's real-world interactions with children are regularly audited for inaccurate or inappropriate responses, and whether there are mechanisms to track its effects on children's wellbeing as it evolves.
- **Child-centred design:** whether children's own perspectives and vulnerabilities are built into the design process from the start, whether child-development experts are actually involved, and whether the system itself can help build children's own understanding of what AI is, and is not, capable of.
- **Supporting teacher autonomy:** Ensuring that children's use of AI in an educational context is continually monitored and supported and mediated by teachers. To enable this, it is important to enable teacher capacity and document their concerns during the process of design and development.

These are not abstract concerns. UNESCO's 2025 guidance warns that AI and digital tools, deployed without deliberate safeguards, risk eroding rights and deepening inequality rather than reducing it [31]. A related risk, particularly in the Indian context, is that AI tools deployed unevenly across well-resourced and under-resourced schools could reproduce, or even sharpen, the same regional and socioeconomic disparities already visible in ASER data, unless deployment is explicitly designed to reach the students furthest behind, rather than the students whose schools can already afford good technology.

Another category of risk are pedagogical and cognitive risks. This is the effect AI systems can have on how students actually think and learn. These can manifest as follows:

- **Cognitive offloading:** letting a tool do work a student would otherwise have practised doing themselves.
- **Excessive scaffolding:** when a system supports a student so heavily that the underlying skill is never independently developed. By removing the productive struggle involved in reasoning, retrieval, and problem-solving, AI assistance may appear effective in the short term while producing little durable learning or transfer.
- **Answer dependence:** where a student increasingly relies on an AI tool to complete tasks rather than developing the underlying knowledge or skill independently. This is a related but distinct failure mode that may remain hidden when evaluation focuses primarily on engagement or task completion instead of evidence of learning and independent performance.
- **Automation bias:** the tendency to trust an AI-generated answer or recommendation simply because it came from a system, can undermine exactly the critical thinking this paper has argued AI should support. AI systems can reinforce a student's misconceptions, particularly when a model is not actually diagnosing the underlying error but pattern-matching to a plausible-sounding response.

What makes this category of risk different from the others discussed in this section is that it cannot be caught by testing a system for data breaches, inappropriate content, or bias in isolation: it only shows up in how students actually perform and behave over time, which means it has to be evaluated through

learning outcomes and learner behaviour directly, not inferred from a system passing technical safety checks.

Model-related risks arise because most generative AI systems are trained on general-purpose data and optimised for plausible, fluent output rather than pedagogical correctness, which means a model can state a wrong explanation as confidently as a right one, and is not designed to recognise a student's specific misconception. Data-related risks compound this: models trained predominantly on English-language, Global North content are poorly equipped to reflect Indian curricula, regional languages, or locally common misconceptions, which is the same underlying gap the discussion in this paper addresses for language access, but applies equally to the pedagogical content itself. Both risks point to the same conclusion already reached elsewhere in this paper: a general-purpose AI model is not a substitute for one built and evaluated against a specific, well-diagnosed learning mechanism, and the absence of pedagogical or data grounding is itself a risk that deserves the same scrutiny as privacy or child-safety failures.

6 Toward a Policy Framework

A useful way to organize policy recommendations is by who is being asked to act, since teachers, developers, and regulators each control different levers.

Adopting a Systems Approach

Before turning to those specific audiences, we suggest an overarching approach or framing to guide policy direction. A systems analysis of AI and education in the Global South argues that the countries getting the most out of AI are not the ones adopting the most tools, but the ones treating curriculum, teacher development, assessment, and governance as one connected system rather than a set of separate technology decisions [27].

Applied to India, this means asking how a given tool changes the whole system it enters: what it does to teacher workload and authority, what it does to how FLN progress gets measured, and what it does to the gap between well-resourced and under-resourced schools. The same analysis insists that solutions be human-centered, built around what teachers and students actually need rather than around what a given AI model happens to do well, and locally adapted, meaning designed for India's specific languages, classroom conditions, and existing FLN architecture [27]. Every recommendation that follows, for teachers, developers, and policymakers alike, should be read against that standard: Whether it strengthens the system as a whole, whether it starts from real classroom needs, and whether it is genuinely adapted to the Indian context.

Toward a Policy Framework: recommendations by audience

Guiding standard: systemic, human-centered, and locally-adapted design (does it strengthen the whole system, start from real classroom needs, and fit the Indian context?)

Teachers & Schools	Developers & Researchers	Policymakers
Stay in all three loops moment-to-moment decisions, planning & reflection, and tool selection/evaluation	Test communication & transparency safeguards before deployment, not just after harm	Enforce DPDPA 2023 & applicable laws verifiable parental consent, plus education-specific guardrails
Delegate routine tasks only lesson planning, feedback, admin; review AI output before class use	Keep models inspectable explainable, and overridable so teachers can reject recommendations	Use collective procurement commission tools for identified FLN gaps via DIKSHA / VSK, not ad hoc buying
Protect student wellbeing shield from harmful/age-inappropriate content; teach online safety directly	Align models to India's own vision of learning strengths-based, socially collaborative, open-ended	Require evidence standards TaRL-level rigor before scaling any AI-enabled edtech nationally
Apply the age-tiered model Classes 1–5, 6–8, 9–12 each get different access and aims	Design for children's rights best interests as primary consideration, evolving capacities, Child Rights Impact Assessment before scale-up	Invest in shared infrastructure Bhashini for Indian languages; AI literacy as a system-wide competency

Shared priority: close the adoption and infrastructure gap first

6.1 For teachers and school administrators

For administrators contemplating incorporating AI for K-12 education within their institutions, it is crucial to understand where AI can genuinely augment student learning and where it cannot, and use it in ways that build students' critical thinking and their capacity to engage authentically with the world around them, rather than substituting for either. This includes being deliberate about which decisions remain human, such as grading judgment calls, handling a distressed student, or deciding what a class actually needs, and which can reasonably be supported by AI tools.

A 2023 framework from the US Department of Education's Office of Educational Technology offers guidance, albeit in the context of a different educational system. It describes three distinct loops in which teachers need to stay actively involved [1]. The first is the moment-to-moment loop of everyday teaching decisions, such as which student needs attention right now, or how to respond to an unexpected answer. The second is the planning and reflection loop: preparing lessons, reviewing how a class went, and adjusting for next time. The third, is the loop in which teachers actually help design, select, and evaluate the AI tools their schools adopt, rather than having those tools handed down to them [1,19].

We recommend a guiding, national policy which permits teachers to use institutionally approved AI tools

for lesson planning, content creation, assessment design, differentiated instruction, feedback generation, and routine administrative tasks. However, teachers and institutions must retain the responsibility for accuracy, fairness, and educational quality squarely with the teacher: AI-generated material should be reviewed and validated before it ever reaches a classroom, not published straight from the tool. The same framework also supports student privacy and confidentiality whenever they use these tools.

Going beyond advocating for AI literacy as a basic competency, teacher capacity-building should also be anchored to a specific standard. UNESCO's AI Competency Framework for Teachers, published in 2024, is the most detailed reference point available for teacher training and capacity building. It defines fifteen competencies across five dimensions, and utilises a human-centred mindset that keeps AI subordinate to educators' own judgment and values. This aligns with our recommendations on keeping humans-in-the-loop, addressing privacy, data protection, and algorithmic bias; AI foundations and applications, the technical literacy needed to understand what a given tool can and cannot do; AI pedagogy, integrating AI specifically into lesson planning, instruction, and assessment; and AI for professional learning, using AI to support a teacher's own ongoing development [32]. Each dimension is further structured across three progression levels, Acquire, Deepen, and Create, so that a teacher with no prior AI exposure and a teacher already integrating AI into their practice are given different, appropriately pitched learning objectives [32].

This distinction between generic AI literacy and teacher capacity is important; by UNESCO's account, only seven countries had developed any AI framework or programme for teachers as of 2022 [32]. Teacher capacity building should thus be structured around progression through these five dimensions, with particular attention to the dimensions of ethics and pedagogy, rather than a tool-based notion of AI literacy.

6.2 An age-tiered access model for schools

In this regard, guidance can be drawn from AI Samarth (introduced in Section 3.1) [34]. Rather than technical AI skill-building, it defines AI literacy as the knowledge and practices needed to understand AI, recognise its impact on daily life, and use it meaningfully, safely, and responsibly, deliberately prioritising awareness and safe use over programming or system-design skills at this stage.

The curriculum is structured around four common pillars for both students and teachers: Understanding AI (what AI is, its history, and how it differs from human intelligence), Essential AI Concepts (data literacy, machine learning, NLP, computer vision), Social Impact and Ethics of AI (bias, privacy, accountability, transparency, environmental impact), and Practical Applications of AI (prompting, using AI for learning and creative work, and critically verifying AI-generated output). The teacher framework additionally covers using AI for lesson planning, administrative tasks, and professional development, while reinforcing that teachers remain the guiding authority in the classroom. The framework draws on UNESCO's AI Competency Frameworks for Students and Teachers, Day of AI, Experience AI, and

Digital Promise's AI Literacy Framework, and will be disseminated as an open-source, multilingual (English plus four Indian languages initially) content stack through online self-learning modules and teacher-led cascaded training.

One practical question is at what age, and under what supervision, a child should get any access to AI at all. This should also be assessed based on developmental stages, that is, what a student should actually be able to do with AI at a given stage, not just how closely supervised their access is. UNESCO's AI Competency Framework for Students provides guidance. It defines twelve competencies across four dimensions, human-centred mindset, ethics of AI, AI techniques and applications, and AI system design, spanning three progression levels: Understand, Apply, and Create [33].

When considering the question of age based access, we must consider two aspects: the first being how much a student can do with AI unsupervised, and the second being what kind of relationship a student has with AI based on their competencies, from user to critical evaluator to creator. An age-tiered access model, as described below, should therefore be viewed as single developmental arc: supervision loosens as competency deepens, moving from a teacher-mediated tool in primary school, to a supervised object of study in middle school, to an independently used and eventually co-designed system by the end of secondary school.

The AI Samarth framework offers a real-world test of this age-tiered structure. Within its target population, (middle and high school students), the framework maps onto the Middle and Secondary tiers of the model above, and prioritises foundational literacy readiness for primary school students. AI Samarth uses the same four pillars for both middle and high school students, Understanding AI, Essential AI Concepts, Social Impact and Ethics of AI, and Practical Applications of AI, which map closely onto the four dimensions of UNESCO's student framework discussed above [33,34]. Its design process prioritises awareness and safe use over technical skill-building at this stage, deliberately holding back from UNESCO's Create level for the general student population and reserving system-building exposure for a narrow "Builders of AI" careers strand aimed at students pursuing STEM subjects [34]. Apply is the realistic target for most students in Classes 9 to 12, and Create is appropriate only for students in more advanced, specialised computer science or project-based coursework, not as a universal expectation.

FIG. 04

An age-tiered model: access permissions and AI competency progression

	ACCESS MODEL	AI COMPETENCY FOCUS (UNESCO)	PEDAGOGICAL OBJECTIVE
Primary Classes 1–5	Teacher-mediated only. No independent, unrestricted student access to generative AI.	Below Level 1 (Understand). Building the FLN readiness and simple awareness that AI is a tool a teacher uses, not a decision-maker itself.	Strengthen FLN through teacher-directed diagnostics, adaptive learning, pronunciation support.
Middle Classes 6–8	Supervised student access to institutionally approved AI tools (gamified learning, coding, formative assessment).	Level 1, Understand. Human-centred mindset and ethics of AI: recognizing bias, privacy, and misinformation; first hands-on AI techniques.	Build AI literacy alongside subject learning; practice verifying AI output.
Secondary & Sr. Secondary Classes 9–12	More independent access to approved tools for research, coding, and project work, subject to school policy.	Level 2, Apply, moving into Level 3, Create for advanced coursework: using AI critically in real work, and beginning to build or evaluate AI system design.	Critical evaluation and academic integrity; assessments classified as AI-prohibited, permitted-with-disclosure, or integrated.
Understand → Apply → Create (UNESCO AI CFS) INCREASING ACCESS, AUTONOMY & AI COMPETENCY →			

For primary school (Classes 1 to 5), the model restricts students to teacher-mediated use only: no independent, unrestricted access to generative AI, but teacher-directed use for FLN diagnostics, adaptive learning, pronunciation support, and personalized intervention is explicitly permitted. The pedagogical goal at this stage is to strengthen foundational literacy, numeracy, and problem-solving without letting the technology substitute for the learning experiences that build those skills in the first place, while giving children an early, age-appropriate sense that AI is a tool that assists learning rather than a decision-maker in its own right. On the competency axis, this stage sits below UNESCO's framework: the priority is building the foundational literacy, numeracy, and language skills a child will need before any of the twelve competencies become meaningful, alongside simple, age-appropriate exposure to the idea that AI is a tool a teacher uses and not a decision-maker. India's National Education Policy structures schooling into four stages under its 5+3+3+4 model: Foundational (ages 3 to 8, spanning pre-school through Grade 2, covered by the NCF-FS discussed in Section 3.1), Preparatory (Grades 3 to 5), Middle (Grades 6 to 8), and Secondary (Grades 9 to 12). CBSE's own Computational Thinking and Artificial Intelligence curriculum uses this Preparatory Stage terminology for Classes 3 to 5, distinct from NCF-FS's Foundational Stage, even though this paper's own Primary tier (Classes 1 to 5) spans both bands. Because the CT and AI curriculum only begins at Class 3, it speaks to the upper part of the Primary tier

and is silent on Classes 1 and 2, which remain governed by NCF-FS and the FLN-focused recommendations in Section 3.1 alone.

For Classes 3 to 5, the CBSE's curriculum deliberately excludes AI entirely. It embeds computational thinking, decomposition, pattern recognition, abstraction, and algorithmic reasoning into existing Mathematics and "The World Around Us" lessons. This rationale is drawn from Asunda et al's argument that the abstraction, decomposition, and algorithmic reasoning underlying computational thinking supply the cognitive scaffolding which AI understanding depends on. Thus, computational thinking is a cognitive prerequisite AI cannot be meaningfully taught without [35,44]. This deliberately unplugged, screen-free design is not in tension with the Primary tier's own access model above, once the distinction we have introduced earlier in Section 1 is kept in view: the age-tiered model concerns AI as a tool teachers may use to deliver instruction, while CBSE's curriculum concerns AI as content children are formally taught. A school can consistently use a teacher-directed adaptive learning tool with a Class 4 student for FLN diagnostics while that same student's formal CT and AI curriculum remains entirely unplugged, since the first is an instructional means and the second is a curricular subject.

For middle school (Classes 6 to 8), the model allows supervised student access to institutionally approved tools for gamified learning, simulations, coding and computational thinking, language learning, and formative assessment. Alongside using AI to learn, students start learning about AI itself, including ideas like algorithmic recommendation, bias, privacy, misinformation, and deepfakes, and practicing the habit of checking AI-generated output. This is where UNESCO's Understand level becomes directly applicable: students begin recognizing what AI is and is not, why an AI response might be biased or wrong, and how their own data is collected and used, corresponding to the framework's human-centred mindset and ethics of AI dimensions specifically [33]. Structured coding and simulation activities also begin building the AI techniques and applications dimension, giving students hands-on exposure to how a simple AI system actually works.

For Classes 6 to 8, the CBSE curriculum introduces AI for the first time, alongside advanced computational thinking and interdisciplinary projects, with explicit learning outcomes covering supervised, unsupervised, and reinforcement learning at a conceptual level and basic data handling. Bias, fairness, and digital footprints appear among the Class 6 curricular goals, extending to informed consent and data-privacy practice by Class 7, and deepening into hands-on no-code AI tools and a full "AI project lifecycle" by Class 8 [35]. Those no-code tools are limited to platforms such as Teachable Machine and Machine Learning for Kids, which let students build and test simple classifiers without writing code; generative AI does not appear anywhere in the Classes 3 to 8 curriculum. This is a narrower scope than the Middle tier described in this paper, which contemplates supervised access to a wider range of institutionally approved tools, potentially including generative and conversational systems. The omission is deliberate: CBSE sequences CT before AI and no-code tools before generative ones, on the view that a student should be able to reason about how a system arrives at an output before being handed a system whose reasoning is opaque. This is one of the more distinctive features of the Indian approach, and

schools following CBSE's own sequencing would in practice restrict Classes 6 to 8 more tightly than the access model above allows on its own.

For secondary and senior secondary school (Classes 9 to 12), students get more independent access to institutionally approved AI tools for research, coding, project work, and career exploration, alongside explicit academic integrity expectations: AI-generated work cannot be presented as entirely the student's own, and disclosure is required whenever a teacher asks for it. The framework classifies individual assessments as AI-prohibited, AI-permitted-with-disclosure, or AI-integrated, which gives schools a more precise tool than an outright ban for managing AI's growing presence in coursework. This sits alongside CBSE's existing Artificial Intelligence elective (Subject Code 417), introduced in January 2019 for Classes 9 to 12 and now taken by more than 750,000 students annually, which predates the 2026-27 Computational Thinking and AI curriculum discussed below and already gives schools a dedicated, assessed pathway into AI at this tier [43]. On the competency axis, this stage corresponds to UNESCO's Apply level across all four dimensions, using AI tools critically and responsibly in real academic work. For students engaged in more advanced computer science or project-based coursework, this is the beginning of the Create level: building, modifying, or critically evaluating the design of a simple AI system [33].

This age-tiered structure is a complement to the child-safety framework in Section 6.3, not a substitute for it. The same framework also flags a related responsibility: protecting students from harmful, misleading, or age-inappropriate AI-generated content, and actively teaching them about misinformation, deepfakes, and online safety. CBSE's design is also a concrete, national-scale instance of the systemic approach this paper has otherwise argued for: embedding CT and AI into existing subjects, sequencing content to match cognitive readiness [27,35]. However, the effectiveness of this approach will lie in its implementation and teacher capacity, as we have emphasised elsewhere in this paper.

6.3 For developers of AI systems used by children

Our recommendations in this regard build directly on the practical framework developed through the participatory AI framework, and on the wider body of expert literature and policy analysis it draws together [30], along with the initial experience from the pilot phase of the Bodhan-AI programme [19]. In practice, this means involving child-development expertise in design from the outset, building in transparency about the system's non-human nature, establishing clear escalation paths to a human adult when a child shows signs of distress, and committing to ongoing, independent audits of how the system actually behaves with children in the field, not just how it performs in a controlled evaluation.

More specifically, developers need to test for a set of design questions before deployment. On the communication side, this means checking whether the system can pick up on a child's playful or non-literal way of talking, whether it recognizes emotional cues and responds to them appropriately, and whether clear protocols exist to stop it from producing anything that could harm a child's wellbeing.

On transparency, the system should make its non-human nature obvious to a child and based on the age-tiered model, should be mediated through a teacher or parent. On an ongoing basis, developers need real mechanisms for monitoring how the system actually performs with children over time: regular audits that catch inaccurate or inappropriate responses, and a standing process for assessing the system's effect on children's wellbeing as it evolves and as new risks emerge. Underlying this, the design process itself needs to start from children's perspectives and vulnerabilities, with genuine involvement from child-development experts and educators, and with an explicit goal of helping children build their own understanding of what AI can and cannot do.

Drawing from international child rights frameworks, the best interests of the child must be a primary consideration in every decision about how an AI system aimed at children is designed, deployed, or governed [28]. A more granular version of the same principle translates this into eleven concrete design principles tied to specific articles of the Convention, including explicit requirements that products be age-appropriate by design and that a Child Rights Impact Assessment be conducted before deployment, which is a framework that developers can draw from [36,41].

The concept of children's evolving capacities, the recognition that a child's appropriate degree of autonomy and access to information grows as they mature, is the underlying justification for the age-tiered access model set out in Section 6.2: a system appropriate for a Class 11 student researching an independent project is simply not the same system that should be placed in front of a Class 2 student learning to read, and developers should treat that as a design constraint [28,36].

Alongside safety validation, developers should also validate AI systems against educational outcomes. Developers should be able to demonstrate that the system improves or supports a defined educational outcome, not merely that it is safe, accurate or engaging.

6.4 For developers and researchers: aligning AI models to India's own vision of learning

Beyond child safety specifically, systems should be judged on how closely they fit India's actual learning goals. Three shifts should be built into any Indian deployment. First, from deficit-based to strengths-based model: a system should identify what a student already brings to their learning, not only what they are getting wrong. Second, from narrowly cognitive to genuinely social: models should support group work, peer learning, and the kind of reasoning and justification skills that matter for numeracy specifically, rather than treating learning as a purely individual, right-or-wrong process. Third, from fixed, closed tasks to open-ended and creative ones: assisting students with connecting to classroom content [1,14].

Further, the acceptable autonomy, evidence requirements, explainability requirements and safety controls should differ between an AI system that assists a teacher and one that interacts directly with a child. For example, in a teacher facing system, teachers should be able to understand the basis for a learner diagnosis, recommended the next activity, feedback or intervention, and should be able to override or

defer an AI recommendation. In a student facing system, drawing from the risks identified earlier in this paper, attention should be paid to designing systems which are developmentally appropriate, based on evidence requirements stemming from educational outcomes, and allowing parents and/or teachers to review and override system outputs.

Continued research and development, in India and elsewhere, needs to focus specifically on how a given AI model performs for different kinds of learners, and teachers, in different classroom conditions [1]. There is also significant opportunity to leverage AI for teachers to strengthen core teaching practices such as lesson planning, differentiated instruction, content creation, formative assessment, feedback, multilingual teaching, and professional productivity, while providing them the autonomy to interrogate and override any AI-generated outputs.

Recommendations for policymakers and regulators

01 Regulatory compliance

DPDP Act 2023 & DPDP Rules 2025: verifiable parental consent before processing a child's data; penalties up to ₹200 crore. Applies to every classroom AI tool, government or private vendor.

02 Coherent, integrated governance

Align AI governance with existing FLN frameworks, NIPUN Bharat and NCF 2023.

03 Transparency and bias review

Schools and platforms disclose what data is shared with AI systems; periodic bias audits for any AI used in placement, assessment, or other educational decisions.

04 Multilingual equity and AI literacy

Build on Bhashini's translation and speech models across all 22 scheduled languages; treat AI literacy as a basic digital competency for students, teachers, and administrators alike.

05 Evidence standards and guardrails

Require TaRL-level evidence of effectiveness before scaling any AI-enabled edtech; set education-specific guardrails, including a guaranteed human fallback for high-stakes decisions.

06 Collective, needs-first procurement

Identify priority problems with teachers, students, and parents first, then commission tools to fit, as Bodhan AI's pilot-assess-monitor cycle does.

07 Digital public good status

Pursue open licensing and DPG standards for public AI-in-education platforms, building sovereign, locally-adapted, and interoperable infrastructure over time.

Shared foundation across all seven priorities: close the connectivity, device, and digital-skills gap first (Section 2)

6.5 For policymakers and regulators

Enforcing existing applicable laws: Policymakers and regulators must address AI governance in K-12 education as part of the broader regulatory landscape it intersects with, rather than as an isolated technology policy question. In India, this can, for example, be addressed through the Digital Personal Data Protection Act, 2023 and Digital Personal Data Protection Rules, 2025, which require verifiable parental or guardian consent before any platform processes a child's personal data, and set penalties of

up to 200 crore rupees specifically for violations involving children's data [37].

Any AI-enabled tool used in an Indian classroom, whether run by government or a private vendor, is directly covered by this requirement, alongside existing consumer protection law. More broadly, policymakers need a coherent governance framework for AI in K-12 education that is informed by the FLN-focused work already under way through NIPUN Bharat and the NCF 2023 [6,7,8]. Schools and platform providers should be transparent with students and families about exactly what information may or may not be shared with an AI system. AI systems used for any educational or administrative decision, including placement and assessment, should be periodically reviewed for unfair bias.

Multilingual capabilities as a priority: Multilingual AI capability should be actively encouraged so that AI-enabled tools support learning in Indian languages, not only English, while still safeguarding academic quality and originality. India already has infrastructure to build this on: Bhashini, the National Language Translation Mission launched in 2022, has built AI translation and speech models across all 22 scheduled languages and is already being used to translate learning material into regional and tribal languages [12]. This connects directly to the equity gap described in Section 2, since language is one of the more concrete, addressable barriers standing between rural and urban students' access to these tools. AI literacy itself, for students, teachers, and administrators alike, should be treated as a basic digital competency across the system.

Evidence based development and policy: Just as India already expects rigor from interventions like TaRL before scaling them nationally, AI-enabled edtech should be held to comparable evidence-of-effectiveness expectations before wide adoption [1,13,38].

Needs based development and a public goods framing: A more structural priority is to ensure that AI is developed for K-12 outcomes based on needs, instead of wholesale adoption of products on the market. Collective procurement can address this problem: states, districts, or school networks first identify a priority problem through consultation with teachers, students, and parents, and only then commission software built to that specification, rather than buying a generic product off the shelf [39].

Taking the example of Bodhan AI, this form of procurement based on piloting products, assessing need and continually monitoring through training exercises and observing school-level impact can ensure that AI solutions are tailored as a foundational infrastructural layer, towards building sovereign, context-aware AI tools for learning and teaching, ensuring interoperability at scale. Procurement should specify educational outcomes and evidence requirements, not merely technical specifications.

AI for K-12 education in India can usefully be understood through the lens of digital public goods (DPGs): open technologies that meet a recognized standard for open licensing [40]. Framed this way, digital public good status becomes an actionable target for India's AI-in-education platforms, and a further argument for the collective procurement approach discussed above, moving towards publicly built tools on locally adapted foundations.

Interoperability: Based on the above recommendation of treating educational AI as a digital public good, interoperability should be considered a core design principle for AI-enabled education. Where appropriate, AI components should be capable of integrating with existing public education infrastructure, assessment systems, language resources, and relevant data standards, thereby reducing fragmentation, avoiding unnecessary proprietary dependencies, and enabling reuse and sustainable scaling.

Across all these groups, closing the adoption and infrastructure gap documented earlier has to be treated as a first-order policy goal, requiring deliberate investment in connectivity, devices, and digital skills for under-resourced schools and states.

7 Conclusion

If used deliberately, responsibly, and designed in alignment with learning science and developmental principles under a child-rights framework, AI poses several opportunities to improve K-12 educational outcomes and bridge the educational divide in India. These systems must be aimed at well-diagnosed problems like the numeracy gap, designed with children's developmental needs built in from the start, and deployed in a way that actively narrows existing regional and socioeconomic disparities.

Whether India and other countries in the Global South realize this potential will depend less on the sophistication of the AI itself, and more on the policy choices made around infrastructure, teacher support, child safety, and equity of access. We thus hope that our recommendations in this paper provide a path forward towards developing evidence based policy which can support the growth of AI development in the education sector, while centering the growth and development of India's children. Our overarching conclusion and recommendation is that AI should be introduced where there is a clearly identified educational need, a sound pedagogical rationale for its use, appropriate teacher mediation, evidence of educational benefit, and adequate safeguards for equity, child rights and data protection.

References

- [1] U.S. Department of Education, Office of Educational Technology. (2023). Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations. <<https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf>>.
- [2] Microsoft AI Economy Institute. (2026). AI Diffusion Report 2025 H2: A Widening Digital Divide; and Q1 2026 update. <<https://www.microsoft.com/en-us/corporate-responsibility/topics/ai-economy-institute/reports/global-ai-adoption-2026-q1/>>.
- [3] UNESCO. (2025). AI and the Future of Education: Disruptions, Dilemmas and Directions. Released at Digital Learning Week 2025, Paris. <<https://www.unesco.org/en/articles/ai-and-future-education-disruptions-dilemmas-and-directions>>.
- [4] National Sample Survey Office (NSSO), Government of India. (2021). Survey findings on rural student access to online learning during the COVID-19 pandemic, as reported in Business Standard, 'EdTech woes: Why India's online & offline education platforms are crashing' (2025). <https://www.business-standard.com/india-news/india-education-sector-edtech-crisis-coaching-centers-downfall-125012400833_1.html>.
- [5] Artopolous, A. (2024). AI and Unequal Knowledge in the Global South. In Faul, M.V. (Ed.), AI and Digital Inequities, Policy Insights #04. NORRAG. <<https://www.norrageducation.org/ai-and-unequal-knowledge-in-the-global-south/>>.
- [6] Ministry of Education, Government of India. (2020). National Education Policy 2020.
- [7] Department of School Education and Literacy, Ministry of Education, Government of India. (2021). NIPUN Bharat Mission Guidelines. <https://ncert.nic.in/pdf/NIPUN_BharatStakeholders_RolesResponsibilities.pdf>.
- [8] National Council of Educational Research and Training (NCERT). (2023). National Curriculum Framework for School Education 2023. <https://ncert.nic.in/pdf/focus-group/NCF-SE_2023EN.pdf>.
- [9] ASER Centre. (2024). Annual Status of Education Report (ASER) 2024, National Findings.
- [10] National Council of Educational Research and Training (NCERT) and Ministry of Education, Government of India. DIKSHA (Digital Infrastructure for Knowledge Sharing) platform, live since 2017; National e-Governance Division (NeGD) integration of Personalized Adaptive Learning into DIKSHA (2023); Vidya Samiksha Kendra state dashboards; and Andhra Pradesh's partnerships with Embibe, ConveGenius, and Mindspark.
- [11] Ministry of Education, Government of India. UDISE+ (Unified District Information System for Education Plus) 2024-25 National Report, released August 2025, including national pupil-teacher ratio and teacher vacancy data.
- [12] Digital India Bhashini Division (DIBD), Ministry of Electronics and Information Technology, Government of India. Bhashini: National Language Translation Mission, launched July 2022.
- [13] Pratham and the Abdul Latif Jameel Poverty Action Lab (J-PAL). Teaching at the Right Level (TaRL): evidence base and impact evaluations, 2005-present. <<https://www.povertyactionlab.org/case-study/teaching-right-level-improve-learning>>.
- [14] Banerjee, A. V., Bhattacharjee, S., Chattopadhyay, R., Duflo, E., Ganimian, A. J., Rajah, K., & Spelke, E. S. (2025). Children's arithmetic skills do not transfer between applied and academic mathematics. *Nature*, 639, 673-681. doi: 10.1038/s41586-024-08502-w.
- [15] ASER Centre. (2023, 2024). Annual Status of Education Report (ASER) 2022 and 2023, National Findings (used for historical/regional trend comparisons). See also, Arpan Tulsyan and Ishika Ranjan, "Tracking Regional Disparities in Learning Outcomes," ORF Issue Brief No. 852, Observer Research Foundation, December 2025. <<https://www.orfonline.org/research/tracking-regional-disparities-in-learning-outcomes>>.
- [16] Goyal, H., Garg, G., Mordia, P., Ramachandran, V., Kumar, D., & Challa, J. S. (2025). Thematic insights into the impact of large language models on K-12 education in rural India from student volunteers' perspectives. *Scientific Reports*, 15, Article 45681. <https://doi.org/10.1038/s41598-025-18047-1>.

- [17] Srinivasan, V., & Murthy, H. (2021). Improving reading and comprehension in K-12: Evidence from a large-scale AI technology intervention in India. *Computers and Education: Artificial Intelligence*, 2, 100019. <https://doi.org/10.1016/j.caeai.2021.100019>.
- [18] Muralidharan, K., Singh, A., & Ganimian, A. J. (2019). Disrupting Education? Experimental Evidence on Technology-Aided Instruction in India. *American Economic Review*, 109(4), 1426-1460. DOI: 10.1257/aer.20171112. Also: Muralidharan, K., & Singh, A. (2021). Evaluation of a technology-aided instruction platform (Mindspark) in government schools in Rajasthan, AEA RCT Registry. February 02. <https://doi.org/10.1257/rct.2546-1.3>.
- [19] IIT Madras and Ministry of Education, Government of India. (2026). Bodhan AI: a sovereign, AI-powered national education ecosystem (Bharat Bodhan AI Conclave, February 2026). <<https://bodhan.ai/>>.
- [20] De Simone, M., Tiberti, F., Barron Rodriguez, M., Manolio, F., Mosuro, W., & Dikoru, E. J. (2025). From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria. *World Bank Policy Research Working Paper* 11125. <https://doi.org/10.1596/1813-9450-11125>. Also: 'The ethics of AI or techno-solutionism? UNESCO's policy guidance on AI in education,' on the environmental footprint of generative AI. <https://doi.org/10.1080/01425692.2025.2502808>.
- [21] Digital personalised learning to improve literacy and numeracy outcomes: a randomised controlled trial in Kenyan pre-primary classrooms. (2025/2026). *Oxford Review of Education / EdTech Hub*, in partnership with EIDU and Women Educational Researchers of Kenya (WERK). <https://doi.org/10.1080/02671522.2025.2605645>.
- [22] Steenbergen-Hu, S., & Cooper, H. (2013). A meta-analysis of the effectiveness of intelligent tutoring systems on K-12 students' mathematical learning. *Journal of Educational Psychology*, 105(4), 970-987. DOI: 10.1037/a0032447.
- [23] Elmaadaway, M.A.N., El-Naggar, M.E., & Abouhashesh, M.R.I. (2025), Improving Primary School Students' Oral Reading Fluency Through Voice Chatbot-Based AI, *Journal of Computer Assisted Learning*. <<https://doi.org/10.1111/jcal.70019>>.
- [24] World Bank and EdTech Hub. AI in Education Global Community of Interest; World Bank blog (2026), 'The future is Africa: Shaping AI-enabled EdTech for skilling the next generation.' <<https://blogs.worldbank.org/en/education/the-future-is-africa--shaping-ai-enabled-edtech-for-skilling-the->>.
- [25] Inter-American Development Bank and Ceibal. (2025-2026). AI and Education: Building the Future Through Digital Transformation <<http://dx.doi.org/10.18235/0013500>>. and the "IDB, How AI Is Transforming Education in Latin America and the Caribbean: Lessons from 193 Solutions. Jan 8 2026. <https://www.iadb.org/en/blog/education/how-ai-transforming-education-latin-america-and-caribbean-lessons-193-solutions>>.
- [26] EdTech Hub. AI in Southeast Asia: Marginalised Learners, topic brief; <<https://docs.edtechhub.org/lib/ZAIZ22IV/download/MTDNTJ2U>>.
- [27] Reimers, F., Azim, Z., Palomo, M.R., & Thony, C. (2026). *Artificial Intelligence and Education in the Global South: A Systems Perspective*. Springer, Cham. Open access. DOI:10.1007/978-3-032-11449-5.
- [28] UN Committee on the Rights of the Child. (2021). General Comment No. 25 on Children's Rights in Relation to the Digital Environment. Includes the Convention's four general principles (non-discrimination, best interests of the child, right to life, survival, and development, and right to be heard), provisions on the right to education, the concept of children's evolving capacities, and the requirement for Child Rights Impact Assessments.
- [29] UNICEF, Office of Global Insight and Policy (Innocenti). Policy Guidance on AI for Children, version 2.0 (2021) and version 3.0 (2025). <<https://www.unicef.org/innocenti/reports/policy-guidance-ai-children>>.
- [30] Vidhi Centre for Legal Policy and Centre for Responsible AI (CeRAI), IIT Madras. (2024). Participatory Approaches in AI Development and Governance: A Principled Approach. <<https://doi.org/10.48550/arXiv.2407.13100>>.
- [31] UNESCO. (2025). AI and Education: Protecting the Rights of Learners. <<https://www.unesco.org/en/articles/ai-and-education-protecting-rights-learners>>.
- [32] UNESCO. (2024). AI Competency Framework for Teachers. Miao, F., & Cukurova, M. Paris: UNESCO. <<https://>>

www.unesco.org/en/articles/ai-competency-framework-teachers>.

- [33] UNESCO. (2024). AI Competency Framework for Students. Miao, F., & Shiohira, K. Paris: UNESCO. <<https://www.unesco.org/en/articles/ai-competency-framework-students>>.
- [34] Central Square Foundation. (2025, September). AI Samarth: AI Literacy Curriculum Framework for Students and Teachers. Developed in partnership with the Wadhvani School of Data Science and AI, IIT Madras. <<https://www.centralsquarefoundation.org/reports/ai-samarth-ai-literacy-curriculum-framework-for-students>>.
- [35] Central Board of Secondary Education. (2026). Computational Thinking and Artificial Intelligence Curriculum, Classes 3-8, 2026-27. <https://cbseacademic.nic.in/web_material/CurriculumMain27/CTAI_Pri_2026-27.pdf>.
- [36] Livingstone, S., & Pothong, K. (2023). Child Rights by Design: Guidance for Innovators of Digital Products and Services Used by Children. Digital Futures Commission, 5Rights Foundation. <<https://researchonline.lse.ac.uk/id/eprint/119724>>.
- [37] Ministry of Electronics and Information Technology (MeitY), Government of India. (2025). Digital Personal Data Protection Rules, 2025, notified November 13, 2025, under the Digital Personal Data Protection Act, 2023.
- [38] Fesler, L., Martinez, J., Agnew, C., Loeb, S. "The Evidence Base on AI in K-12: A 2026 Review," AI Hub for Education of the SCALE Initiative, Stanford University, 2026. <<https://scale.stanford.edu/sites/default/files/The%20Evidence%20Base%20on%20AI%20in%20K-12%20Report.pdf>>.
- [39] Macgilchrist, F. (2024). AI in Education: How Developers and Legislators Make Futures. In Faul, M.V. (Ed.), AI and Digital Inequities, Policy Insights #04. NORRAG. <https://www.norrageducation.org/wp-content/uploads/2024/02/2024_ToC-A5-flyer_PolicyInsights_AI_v5.pdf>.
- [40] Stinckwich, S., Wong, N., Nyamawe, A. S., Liu, J., Latif, F., & Stuart, J. (2026). AI Systems as Digital Public Goods: Evidence and Recommendations from a Multi-Stakeholder Assessment. Research report, United Nations University Institute in Macau, commissioned by the Asian Development Bank, in partnership with the UN Office for Digital and Emerging Technologies. <<https://unu.edu/publication/ai-systems-digital-public-goods>>.
- [41] Livingstone, S., & Shekhawat, G. (2024). AI and Children's Rights. NORRAG Blog, 29 May 2024, republished from Policy Insights #04: AI and Digital Inequities. The Digital Futures for Children Centre, London School of Economics and Political Science, with 5Rights Foundation. <<https://www.norrageducation.org/ai-and-childrens-rights/>>.
- [42] Livingstone, S., Atabey, A., & Pothong, K. (2021). Addressing the problems and realising the benefits of processing children's education data: Report on an expert roundtable. Digital Futures Commission, 5Rights Foundation. <https://researchonline.lse.ac.uk/id/eprint/119739/1/Roundtable_report_25112_final.pdf>.
- [43] Central Board of Secondary Education. (2019). Circular introducing Artificial Intelligence (Subject Code 417) as an elective for Classes 9-12, effective April 2019 academic session. See also Press Trust of India, 'Over 7.9 lakh class 9, 10 CBSE students opted for AI subject during 2024-25 session: Government tells LS,' December 2024.
- [44] Asunda, P., Faezipour, M., Tolemy, T., & Engel, R. (2023). Embracing Computational Thinking as an Impetus for Artificial Intelligence in Integrated STEM Disciplines through Engineering and Technology Education. *Journal of Technology Education*, 34(2). <https://doi.org/10.21061/jte.v34i2.a.3>



The Centre for Responsible AI (CeRAI) at the Indian Institute of Technology, Madras, is a multi-disciplinary, non-profit research centre positioned in the Global South, as one among the few global institutions that specialises in both technical and policy research to ensure and enable the responsible development and deployment of AI systems.

ADDRESS

Indian Institute of Technology Madras
Chennai – 600036, Tamil Nadu, India

FIND US

www.cerai.iitm.ac.in
[LinkedIn /cerai-iitm](#)
[X @cerai_iitm](#)

