

सर्वजन शिवाय | सर्वजन सुखाय
WELFARE FOR ALL | HAPPINESS OF ALL



INDIA AI IMPACT SUMMIT 2026

COMPENDIUM

Real-World Impact of AI in Education

kStep



CENTRAL SQUARE
FOUNDATION

Disclaimer:

The case studies included in this compendium have been evaluated based on information submitted by the respective authors and participating organizations. Responsibility for the accuracy of data, metrics, and representations rests solely with the submitting authors. The evaluation committee and partner institutions shall not be held liable for any discrepancies, omissions, or subsequent changes in the information provided.

This Casebook benefited from AI-assisted editing for clarity and readability. All substantive content, analysis, and conclusions were reviewed and approved by the case study authors and the technical partners

Partners:

IndiaAI Mission, Ministry of Electronics and Information Technology (MeitY), Government of India; Central Square Foundation (CSF), and EkStep Foundation.

We thank the Ministry of Skill Development and Entrepreneurship (MSDE), Government of India for facilitating this compendium.



Table of Contents

I. Foreword	IV-VI
II. Technical note	VII
III. Glossary	X
1. AI NESTHAM	01
2. Amira Learning	03
3. Answerr.ai	05
4. Aural AI	07
5. BharatGen Yojaka	09
6. Bookbot	11
7. Chimple	13
8. CrazyGoldfish Technologies	15
9. Darsel	17
10. Dhrishti	19
11. Directorate of Higher Education, Government of Goa	21



12. Duolingo English Test	23
13. Eedi Labs	25
14. EIDU	27
15. Government of Rajasthan	29
16. SATHEE	31
17. Indra AI	33
18. Khanmigo	35
19. Kibo	37
20. Lenovo Leap NextGen Scholar Program	39
21. Luma Learn	41
22. NavGurukul Foundation for Social Works	43
23. NetPractice	45
24. Paperpal	47
25. Pratham	49
26. QwiXGenie	51



27. Rabbitt AI	53
28. Rocket Learning	55
29. Sakhee	57
30. SocraticAI	59
31. Technovation	61
32. Top Parent	63
33. Torchit Electronics	65
34. TOY8	67
35. Transform Schools, People For Action	69
36. Udhyaam Saathi (AI Mentor)	71

Foreword

As global leaders, thinkers, and innovators convene at the India AI Impact Summit 2026, we look forward to showcase and deliberate on the transformative impact of AI across governance, innovation, and sustainable development. Through this Summit, India aims to deliver a coherent set of outcomes, translating technical discussions into actionable global and regional pathways for AI adoption with development impact.

With a focus on three key sutras—people, progress and planet—the Summit will highlight the real-world impact of AI innovation to bolster human development, catalyse socio-economic transformation, and accelerate the achievement of sustainable development goals.

At the IndiaAI Mission, India continues to advocate for the democratisation of AI—not just in terms of enabling access to AI infrastructure, but also in terms of skills and equal opportunity. This is reflected in India’s world-leading talent ecosystem, achieved as a result of India’s population-scale skilling efforts. We are also ensuring that AI literacy is accessible to everyone, regardless of background or experience. Key to this vision is leveraging AI for advancing educational outcomes, personalising pedagogy, and bridging the digital divide.

The Casebook on Real-World Impact of Artificial Intelligence in Education, developed in partnership with Central Square Foundation and EkStep Foundation, represents a key milestone in these efforts. This casebook is intended to serve as a reference for the global community—specifically policymakers, innovators, and researchers who are navigating the complexities of AI adoption. By documenting high-impact interventions that address critical challenges, it offers a practical blueprint for replication. Our objective is to empower stakeholders to adapt these proven successes to their unique local contexts, thereby accelerating the scaling of solutions across the Global South and catalysing the holistic digital transformation of education systems.

We acknowledge with appreciation the leadership of the Ministry of Skill Development and Entrepreneurship, Government of India in advancing AI deployment in the education sector, and the partnership of the Central Square Foundation and EkStep Foundation, and recognise the contributions of the innovator and research communities in developing AI-based solutions for the public good. Together, these collective efforts demonstrate that artificial intelligence can serve as a powerful instrument for equitable development—transforming lives while upholding the highest standards of safety, transparency, and human dignity.



Shri S. Krishnan

Secretary

Ministry of Electronics and Information Technology
Government of India

Foreword

India stands at a defining inflection point where Artificial Intelligence is transitioning from a frontier technology to a fundamental enabler of social transformation. Guided by the Hon'ble Prime Minister's vision, the Ministry of Skill Development and Entrepreneurship (MSDE) is committed to the principle of 'AI for All'—ensuring that the benefits of this revolution are inclusive, scalable, and accessible to every citizen.

Our strategy is anchored in a three-pronged framework: 'Skilling in AI' to build deep technical competencies; 'Skilling with AI' to leverage adaptive platforms for personalized and multilingual learning; and 'Skilling for AI' to prepare our broader workforce for an AI-augmented economy.

The "Casebook on Real-World Impact of AI in Education" serves as a vital testament to this vision. By documenting 36 scalable solutions from India and across the globe, this compendium moves beyond theoretical potential to demonstrate tangible impact. From enhancing foundational literacy and supporting teachers to empowering the differently-abled through assistive technologies, these case studies prove that technology can indeed bridge the access gap.

I congratulate the contributors for compiling these innovations. They offer a roadmap for how we can deploy AI responsibly to empower our youth and build a Viksit Bharat.



Ms. Debashree Mukherjee

Secretary

MSDE

Foreword

India's vision for the AI Impact Summit 2026 is to advance the global AI dialogue from principles to tangible impact, demonstrating how AI can deliver measurable outcomes for People, Planet, and Progress. As the first Global AI Summit in this series to be held in the Global South, India aims to advance an inclusive and resilient AI future that serves humanity, promotes equitable growth, fosters social development and protects the planet.

The foundational step to achieving this vision is to build a skilled and AI-ready citizenry. Through the IndiaAI Mission, India has prioritised skilling at scale, training students, public officials and professionals in AI and data capabilities, achieving one of the highest AI skill penetration rates in the world. We are committed to harnessing AI for the workforce of tomorrow by innovating learning and teaching practices.

The Casebook on Real-World Impact of Artificial Intelligence in Education, developed in partnership with the Central Square Foundation and EkStep Foundation, highlights the transformative potential of AI in improving education access, equity, quality and learning outcomes. This Casebook is intended to serve as a valuable reference for policymakers, innovators and researchers seeking to adopt and scale proven AI education solutions.

This joint initiative invited contributions from researchers, innovators and institutions worldwide, and received an enthusiastic response of 271 submissions across geographies and contexts. Submissions underwent a multi-stage screening and technical evaluation process, followed by expert committee evaluations with representatives from the Ministry of Skill Development and Entrepreneurship, Government of India, to secure the inclusion of the most impactful, evidence-based and replicable solutions.

The resulting Casebook presents a curated set of 36 AI education applications from across the globe, ranging from AI-enabled foundational literacy tools to innovations in teacher capacity building and inclusive education for neurodiverse learners.

We extend sincere appreciation to the Ministry of Skill Development and Entrepreneurship (MSDE) for their leadership in advancing AI deployment in the education sector, and to Central Square Foundation and EkStep Foundation for their steadfast partnership in the development of this Casebook. Together, these collective efforts reaffirm that AI, when developed and deployed responsibly, ethically and inclusively, can serve as a powerful instrument for strengthening education systems and advancing global knowledge equity.



Shri Abhishek Singh

CEO, IndiaAI Mission

Additional Secretary, Ministry of Electronics and Information Technology

Government of India



Technical Note: AI in Education Casebook

Artificial intelligence is increasingly shaping how education systems design, deliver, and assess learning at scale. Across diverse contexts, education stakeholders are exploring AI-enabled tools to address long-standing challenges such as access gaps, teacher capacity constraints, delayed feedback cycles, language diversity, and administrative burden. The real question is no longer whether AI will enter education, but how it will be integrated in ways that improve teaching and learning outcomes and reduce inequality. Therefore, there is a need for grounded, practical examples that move beyond conceptual discussions and demonstrate how these technologies function within real education ecosystems, particularly with respect to underserved and low-income populations. Documenting practical implementations helps clarify both the opportunities and limitations of AI in improving learning experiences and system efficiency.

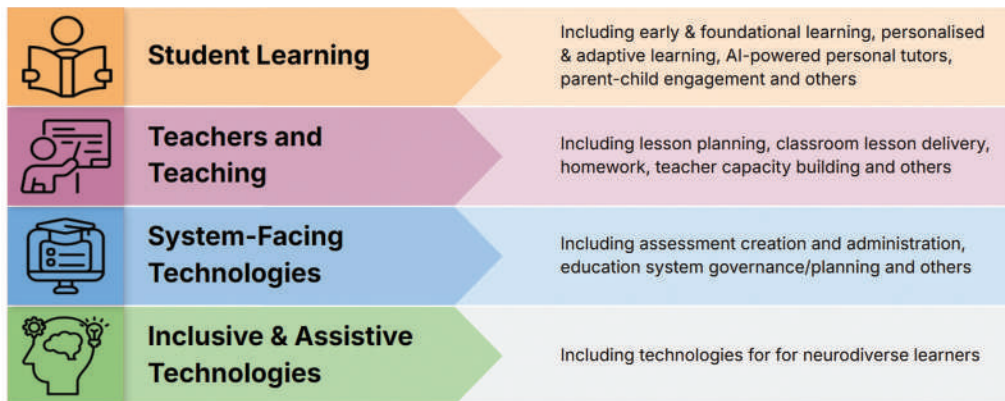
The Casebook on the Real-World Impact of AI in Education, led by the IndiaAI Mission in collaboration with Central Square Foundation and the EkStep Foundation, aims to showcase real-world applications of AI in education, highlight its transformative potential in democratising access to quality education, and facilitate cross-sector and cross-country learning through a globally benchmarked reference document. By bringing together implemented use cases from across different geographies, grade groups, and education applications, it provides practical insights into how AI is being integrated into existing workflows and delivery models in different parts of the globe. The focus on deployed solutions allows policymakers, education practitioners, and academic researchers to understand emerging trends better, identify scalable approaches, and reflect on the evolving role of AI in supporting more accessible, responsive, and inclusive education systems.

The compendium was developed through an open and structured outreach and evaluation process, resulting in 271 global applications across multiple education use cases and grade groups. Submissions were assessed using a weighted scoring framework covering relevance, problem definition, AI application, implementation maturity, reach and scale, evidence of impact, data privacy, inclusion, and overall clarity. The process resulted in a final shortlist of 36 case studies, ensuring balanced representation across use cases as well as Pre-K, K-12, and higher education segments.

Evaluation criteria	Indicators for evaluation
Relevance and alignment	Demonstrates strong fit with key themes and addresses real challenges in low-resource or government school contexts.
Background & problem definition	Clearly defines a meaningful education problem grounded in systemic barriers and user needs.
AI use and technical strength	Explains how AI is applied in a feasible, scalable, and innovative way.
Implementation	Shows a realistic, well-designed approach for operational rollout and sustainability.
Reach and scale data	Provides clear evidence of adoption, user numbers, and growth across regions.
Evidence of impact	Presents credible results demonstrating measurable educational outcomes, based on external or internal evaluations.
Data and privacy	Outlines safeguards for privacy, consent, and child protection.
Inclusion	Ensures accessibility through local adaptation, low-tech compatibility, and support for diverse learners.

Analysis of the submissions and shortlisted use cases highlights several patterns in how AI is currently being applied across education contexts. Some of these are outlined below:

- 1. Distribution across use cases:** Student-facing learning tools formed the largest share of shortlisted solutions, followed by system-facing and administrative technologies. Inclusive and assistive technologies represented the least represented segment, with four shortlisted solutions. Across student learning solutions, personalised learning, personal tutors, and AI literacy tools were among the most commonly represented themes.
- 2. Grade-level coverage:** Shortlisted submissions were analysed across grade groups, including both K-12 and higher education. Student learning solutions showed a relatively even distribution between these age groups, while teaching-focused tools were more concentrated in specific segments. System-facing technologies demonstrated a similar pattern, with representation across both K-12 and higher education as well as a subset of grade-agnostic solutions.



3. Geographic representation and scale: Shortlisted submissions originated from countries such as India, Australia, Germany, Japan, Syria, South Africa, Spain, the United Kingdom, and the United States. The strongest representation was from India, highlighting active experimentation and implementation of AI tools across diverse education settings in the Indian context. International submissions provided additional perspectives, enabling cross-country comparison and learning within the casebook.

4. Responsible design and data governance: A recurring theme across applications was the emphasis on data protection and responsible deployment. Many submissions referenced privacy considerations or incorporated elements aligned with a “privacy by design” approach, reflecting growing awareness of data governance in AI-enabled education tools.

5. Human oversight remains central: Across many submissions, AI systems were positioned as tools that support, rather than replace, educators, reviewers, or institutional processes. Human involvement often remained essential for quality assurance, handling edge cases, maintaining trust, and ensuring contextual judgment where automated outputs may be insufficient.

Taken together, the cases documented in this Casebook underscore that AI’s value in education lies in how well tools are embedded within real-world systems, aligned to user needs, and governed responsibly. While the diversity of applications reflects experimentation at different stages of maturity, common threads emerge around the importance of contextual design, human oversight, and evidence-informed scaling. As AI capabilities continue to evolve, such grounded evidence will be essential to ensuring that innovation advances learning outcomes, equity, and system resilience, rather than exacerbating existing gaps.

¹ Individually tailored instruction

² AI-powered learning assistants

³ Tools that help understand AI and how to use it



Glossary

Term	Definition
AI	Artificial Intelligence
GenAI	Generative AI
OCR	Optical Character Recognition
RAG	Retrieval Augmented Generation
Personalised learning	Technology-enabled approaches that adapt content, pace, or feedback based on individual learner needs.
AWW	Anganwadi Worker
Anganwadi Centre	Community-based early childhood centres under India's ICDS programme that provide early learning, nutrition, and health services.
K-12	The full span of school education from Kindergarten through Grade 12.
FLN	Foundational Literacy and Numeracy
ECCE	Early Childhood Care and Education
DAU	Daily active users
Adoption rate	The proportion of target users who begin using a new tool or system.
NPS	Net Promoter Score
Engagement rate	A measure of how actively users interact with a platform or content over time.

AI NESTHAM: Democratizing AI Literacy Through Gamified Learning for Classes 3-9 in Indian Government Schools

Author: Chintalapally Harsha Vardhan Reddy, Chandana Naga Surya

K-12

Summary

AI NESTHAM is a government-school-focused AI-literacy platform designed to democratise access to foundational artificial-intelligence education for students in Classes 3–9 across Telangana and Andhra Pradesh. Built as a gamified, curriculum-aligned system, AI NESTHAM introduces core concepts such as prompt engineering, computer vision, natural-language processing, machine learning, and robotics through interactive story-driven challenges rather than lecture-based instruction. Priced at a fraction of commercial offerings, the platform aims to close India's emerging "AI divide" by making 21st-century skills accessible to learners from low-income households.

Context

As India accelerates investment in artificial intelligence, most AI-education programmes remain concentrated in private urban schools or premium EdTech platforms charging tens of thousands of rupees per year. Government schools in southern India serve more than twenty-two lakh students from predominantly low-income families, many of whom have little exposure to advanced digital tools. Connectivity constraints, limited hardware budgets, and English-centric curricula further restrict participation.

Compounding these access barriers, most imported AI-education materials are designed for Western schooling systems and do not align

Indian state-board syllabi, creating pedagogical mismatches that limit adoption by public schools. Without early exposure to computational thinking and AI concepts, large cohorts of students risk being excluded from emerging labour markets.

AI NESTHAM was created to change this for government-school students in Telangana and Andhra Pradesh. It teaches core AI ideas like machine learning and robotics - through story-based games where children solve problems instead of listening to lectures. Costing about ₹500 per student each year, it works on basic phones, needs little internet, supports Telugu, and follows state standards, making AI learning affordable and accessible for underserved communities.

Challenge

India faces a growing digital divide in AI education. Government schools in Telangana and Andhra Pradesh serve more than 22 lakh students from low-income families who have little exposure to emerging technologies. Most AI courses are offered by expensive private platforms, costing ₹30,000–₹1,80,000 a year which is far beyond the reach of most households. Existing solutions also assume fast internet, costly devices, and English fluency, which excludes many rural and semi-urban learners.

In addition, most AI curricula are designed for Western education systems and do not align with



Case Study 1

Indian state-board standards, making them difficult for government schools to adopt. Without early exposure in Classes 3-9, students risk missing critical opportunities to build computational thinking and problem-solving skills.

This widening gap threatens India's ambition to become a global AI leader, highlighting the urgent need for affordable, curriculum-aligned, low-infrastructure solutions that can reach public-school students at scale.

Solution & Impact

AI NESTHAM delivers instruction through 35 narrative-driven interactive games featuring characters that guide students through AI concepts using typed responses and puzzle-solving. The system adapts difficulty based on learner inputs and encourages voluntary re-engagement through quest-based progression.

Designed to serve 22 lakh students across 26,000 schools, pilots indicate improved conceptual understanding, higher completion rates, and strong student enthusiasm. At approximately ₹500 per learner annually, the platform offers 60-300× cost reductions relative to commercial alternatives. Early results mainly show qualitative gains, including stronger student interest in AI topics, higher lesson completion, and more students choosing to continue learning on their own. Learners also show better understanding of basic AI ideas, while the programme achieves major cost savings-between 60 and 300 times cheaper than traditional options.

AI NESTHAM delivers instruction through 35 narrative-driven interactive games featuring characters that guide students through AI concepts using typed

Governance frameworks include encryption in transit and at rest, minimal personal-data capture limited to learning metrics, parental-consent protocols, age-appropriate content filters, and restricted educator access. Transparency requirements ensure that automated feedback is explainable rather than opaque. AI NESTHAM illustrates how large-scale AI literacy programmes can be embedded into public systems responsibly and affordably.

Amira Learning: AI-powered oral reading tutor accelerates foundational literacy growth

Authors: Adam Porsch, Ashish Gupta, Ankur Dahiya

K-12

Summary

Amira Learning addresses gaps in foundational literacy support for early-grade students in low-resource schools, where teachers often lack time for individual reading practice. Amira Learning is grounded in more than 30 years of research at Carnegie Mellon University on intelligent tutoring systems for literacy, and is deployed across all 50 U.S. states and 18 countries, supporting 4 million students across diverse linguistic and socioeconomic contexts. In India, the program is implemented with Khushii NGO across four municipal schools in Delhi. The AI-powered oral reading tutor listens to students read aloud and provides immediate feedback. Among 135 students, reading scores increased from 1.79 to 2.78 on the Amira Reading Measure, representing roughly ten months of learning growth within five months.

Context

Foundational literacy remains a major challenge in early-grade education, particularly in low-resource K-12 systems where large class sizes and limited instructional time reduce opportunities for personalised reading practice. Oral reading with timely correction is widely recognised as critical for developing fluency and comprehension, yet it is difficult to deliver consistently at scale. In many municipal schools, literacy instruction relies on whole-class teaching, with limited tools for tracking individual progress. The Khushii NGO operates within public-school environments serving learners

from low-income communities, where device availability, attendance variability, and infrastructure constraints shape implementation choices. Digital learning tools, when used, must integrate into existing school schedules and operate with minimal disruption to classroom workflows.

Challenge

Early-grade learners require frequent, individualised reading practice to build decoding skills, fluency, and comprehension. However, in classrooms with large student-teacher ratios, educators often struggle to provide one-on-one support or monitor progress consistently. Without timely feedback, students may continue practising incorrect reading patterns, slowing literacy development and widening learning gaps.



Students using Amira Learning's AI oral reading support. (Source: Amira Learning website)



Case Study 2

Traditional approaches rely heavily on teacher-led assessment or periodic evaluations, which can be time-intensive and difficult to sustain. Limited access to instructional resources and competing classroom demands further constrain the ability to deliver personalised support. In low-resource settings, infrastructure limitations and inconsistent attendance can also affect learning continuity. These systemic constraints create a gap between recommended literacy practices and what is feasible at scale. As a result, education providers are exploring AI-enabled tools that can offer consistent reading practice and actionable insights without increasing teacher workload or requiring extensive new infrastructure.

Solution & Impact

Amira Learning is an AI-powered oral reading tutor designed to support early literacy through short, structured reading sessions. Students read aloud into a device while the system uses speech processing to identify decoding errors, mispronunciations, and signs of comprehension breakdown, and delivers immediate, targeted prompts aligned with the Science of Reading. In addition, the system generates dashboards for teachers that help translate students' reading behaviors into actionable instructional insights. The AI models underlying this tool are trained on diverse voice datasets, supporting effective use across varied accents, dialects, and multilingual environments.

In partnership with the Khushii NGO, the solution was deployed across four municipal schools in Delhi, reaching 135 students. Over a five-month instructional period, average scores on the Amira Reading Measure increased from 1.79 to 2.78 — equivalent to approximately ten months of learning growth. Students who engaged consistently with the tool demonstrated the strongest improvements, while lower gains were primarily linked to irregular attendance or limited device access.

Amira Learning is an AI-powered oral reading tutor designed to support early literacy through short, structured reading sessions.

These results align with findings from earlier studies in the United States and Sub-Saharan Africa, where AI-supported reading practice produced outcomes comparable to one-on-one tutoring.

Amira Learning developed and operates the platform, and it is already deployed globally across multiple countries and education systems. The Delhi implementation represents a contextual adaptation within a low-resource setting. While the local pilot operated at a modest scale, the broader platform supports millions of learners internationally, demonstrating potential for expansion. Scaling considerations include device availability, structured scheduling within the school day, and ongoing human oversight to interpret data and support instruction. Short session design and compatibility with standard devices help reduce infrastructure requirements, making the model adaptable to diverse education environments while maintaining a focus on equitable access and responsible data practices.

Answerr.ai: Designing Trusted AI Infrastructure for Education

Authors: Simran Suri, Md Qaiser Mallik, Virajeet Sandu

Higher-education

Summary

Answerr.ai is an AI learning infrastructure built around trust, learning, and governance through “Learning Provenance,” a governance-first architecture that documents how learning unfolds by tracking inquiry patterns, revision behaviour, and conceptual development over time. Answerr.ai reframes responsible AI adoption not as a policing problem, but as an infrastructure problem requiring new technical and institutional layers.

Context

As generative AI becomes central to student work, education systems face a growing trust deficit: faculty struggle to assess authentic learning, students worry about over-reliance on AI tools, and institutions lack governance frameworks that ensure equity and compliance. Projections suggest AI systems may soon contribute to a significant share of academic outputs, shifting the challenge from tool adoption to institutional readiness.

Answerr.ai emerged from this governance gap: the absence of infrastructure capable of supporting pedagogical AI use while maintaining transparency, accountability, and regulatory compliance. Rather than viewing AI as a threat to integrity, the initiative frames the issue as a systems design problem requiring new layers of trust, learning validation, and administrative oversight.

Challenge

As AI spreads through classrooms, universities face a growing trust problem. Teachers struggle to judge whether students are truly learning, students remain uncertain about their own skill development, and administrators lack oversight of how AI is being used across courses and departments. This creates risks around fairness, compliance, and academic integrity, especially as institutions attempt to scale AI adoption.

Common responses such as banning tools or relying on detection software have proven adversarial and misaligned with learning goals. Institutions therefore face a deeper question: how can generative AI be integrated into classrooms while preserving trust in assessment, protecting student privacy, and meeting regulatory requirements? Any solution must avoid constant surveillance, minimise data collection, and provide reliable audit trails that accreditation bodies and institutional leaders can depend on. Without trusted infrastructure, AI adoption risks widening inequities and eroding confidence among educators and students.



Case Study 3

Solution & Impact

Answerr.ai embeds governance directly into classroom workflows, using Learning Provenance to shift institutional AI use from restriction toward confidence and accountability. Instead of evaluating final outputs alone, the platform captures how students interact with AI over time, including inquiry patterns, revision cycles, and conceptual development, allowing educators to assess learning processes rather than authorship.

Students access multiple AI models within a unified environment, supporting equitable access and comparative AI literacy. Educators receive dashboards that surface process-level insights into learning progression without monitoring individual content, while institutions gain aggregated analytics that inform policy decisions and ensure compliance without surveillance. Governance mechanisms follow FERPA- and COPPA-aligned practices, including encryption, role-based access, and retention limits that preserve student ownership of learning records.

During the Babson College pilot, faculty confidence in AI-assisted teaching rose from 2.0 to 4.5, and student participation in critical-thinking activities increased from 10% to 33%. Institutions gained system-level visibility into AI usage patterns while maintaining privacy safeguards. The deployment demonstrates how governance-first AI architectures can enable scalable, trustworthy AI adoption in higher education and offers a potential framework for large-scale education systems seeking responsible AI integration.

Answerr.ai embeds governance directly into classroom workflows, instead of evaluating final outputs alone, the platform captures how students interact with AI over time, including inquiry patterns, revision cycles, and conceptual development, allowing educators to assess learning processes rather than authorship.

Aural AI : AI-Powered Listening-First Foundational Literacy for Rural K-12 Learners

Authors: Prajwal G Thanjavur, Ramana Tadepalli

K-12

Summary

Aural is an AI-powered audio tutor for schoolchildren that tackles poor English listening skills in rural India using a screenless, offline-first device paired with cloud-based analysis. It personalises practice, gives teachers clear diagnostics, integrates with government systems, and uses speech recognition and adaptive sequencing - delivering learning gains while embedding strong child-safety, governance, and multilingual support in low-infrastructure settings.

Context

English is part of the curriculum in government primary schools across India, but teaching is mostly textbook-based, with little structured spoken practice. Rural schools serve many first-generation learners and often face unreliable electricity, limited internet access, and large class sizes. Teachers must cover broad syllabi with few supplementary tools. India has created national digital platforms such as DIKSHA to support education at scale, alongside civil-society efforts to improve foundational learning. However, classroom instruction remains largely paper-driven, with limited ways to track individual listening progress. Any new intervention must therefore work within low-connectivity environments, fit existing school workflows, integrate with public systems, and remain affordable for large-scale government adoption.

Challenge

Rural primary-school children in India receive very little meaningful exposure to spoken English, even though it is part of the formal curriculum. Many teachers themselves lack confidence in English, classrooms are overcrowded, and students rely heavily on paper-based instruction. As a result, children may learn to read or write basic words but still struggle to understand simple spoken sentences, which undermines confidence and limits future educational and job opportunities.

Most existing education technologies are poorly suited to these settings. They assume access to smartphones, stable internet connections, and digitally literate adults - conditions that are often absent in government schools in rural areas. Large class sizes make it difficult for teachers to give individual listening practice or feedback, and weak infrastructure restricts the use of audio or interactive tools.

These combined barriers create a persistent “listening gap” that prevents students from building the foundational language skills needed for later literacy and learning. Without affordable, low-infrastructure ways to provide regular spoken-English practice and monitor progress, rural schools struggle to address this gap at scale, deepening inequality between urban and rural learners.



Case Study 4

Solution & Impact

Aural was designed as an offline-first, classroom-based audio learning system that enables regular, personalised English-listening practice in government primary schools without relying on smartphones, screens, or continuous internet connectivity. At the centre of the model is a low-cost, durable audio device that teachers use during class to play curriculum-aligned listening activities several times each week. Students respond verbally, and the device records their answers locally, allowing whole-class participation even in crowded classrooms.

When connectivity becomes available—typically through a teacher’s hotspot or school Wi-Fi—the device securely synchronises recordings to the cloud for processing. Automatic speech-recognition models fine-tuned on Indian children’s speech analyse pronunciation, while natural-language models evaluate comprehension. An adaptive-sequencing engine then determines the next set of activities for each learner, targeting gaps and reinforcing emerging skills. Teachers receive simple dashboards showing which students are struggling, common error patterns, and suggested groupings for follow-up instruction. These insights are delivered through existing state and national education platforms so that Aural fits naturally into established classroom workflows rather than creating parallel systems.

Child safety and responsible AI practices are embedded throughout deployment. Audio data is encrypted in storage and transit, personal information is minimised, and informed consent is required from schools and families. Raw recordings are retained only briefly for quality assurance, while de-identified performance features are preserved for learning analysis.

To support large-scale adoption, Aural integrates directly with India’s Digital Public Infrastructure for education, including DIKSHA, NDEAR-compliant systems, and state dashboards.

Regular fairness audits check for differences in system performance across gender, region, and home language, and all activities remain low-stakes, used for practice rather than grading.

To support large-scale adoption, Aural integrates directly with India’s Digital Public Infrastructure for education, including DIKSHA, NDEAR-compliant systems, and state dashboards. Distribution occurs through partnerships with education departments and civil-society organisations, enabling the model to reach rural classrooms at population scale. By combining infrastructure-aware hardware, adaptive AI, and strong governance, Aural provides a practical pathway to improving foundational listening skills in low-resource school environments.

Aural has reached tens of millions of learners, reporting learning gains above 0.5 SD, over 70% weekly teacher usage, and minimal participation gaps by gender or ability. The system achieves 90% speech-recognition accuracy and 94% agreement with teacher scoring, while data-minimisation, encryption, consent processes, and fairness audits ensure responsible operation at population scale. It also serves 260 million students due to integration with DIKSHA platform.

BharatGen Yojaka: AI-assisted formative assessment for spoken language learning

Author: Ganesh Ramakrishnan, Venkatapathy Subramanian, Nitish Kamal Singh

K-12

Summary

BharatGen Yojaka is a human-in-the-loop AI assessment platform developed to scale spoken-language evaluation in Indian public-school programmes. By automating initial transcription, rubric-conditioned scoring, and diagnostic feedback for oral responses — while requiring teacher verification before results are released — the system reduces grading burdens without undermining pedagogical trust. The project reflects India's broader effort to integrate generative models into public services under strong accountability regimes.

Context

India faces a serious learning crisis, with many students unable to read or do basic maths. Crowded classrooms and a lack of good teaching materials in local languages make the problem worse and leave teachers overwhelmed with planning work. BharatGen Yojaka is an AI teaching assistant that helps create high-quality lessons, stories, quizzes, and worksheets in multiple Indian languages. Built using Google's AI tools and the open-source Sunbird platform, it supports grade-level learning and encourages understanding rather than rote memorisation. Early trials show teachers spent 30% less time planning and students were twice as engaged. Strong data protections and ethical design make the system suitable for large-scale use in low-resource schools.

Challenge

Assessing spoken language at scale is difficult in schools, especially in large programmes serving first-generation learners. Teachers still rely on manual recording and grading, which takes time, adds to workload, and limits how often students are assessed. BharatGen Yojaka, as an AI-supported system, helps with spoken-English assessment while keeping teachers in control. Built with the Kotak Education Foundation, it records student audio, uses AI to suggest scores and feedback based on rubrics, and requires teacher approval before final results.

BharatGen Yojaka needed to reduce grading time while preserving rubric integrity and avoiding fully automated final decisions. The system also had to ensure transparency in scoring, protect sensitive voice data, and align with national governance frameworks governing student assessment and data use.

Trials in semi-urban schools showed grading time fell by about 30%, with strong agreement between teachers and AI, though classroom noise remained a challenge. Overall, the system reduced teacher burden without compromising educational quality.



Case Study 5

Solution & Impact

Students submit recordings through school devices. Speech-to-text engines transcribe responses; scoring models trained on rubric-labelled examples generate preliminary marks and qualitative diagnostics. Teachers review, edit, or override all outputs.

Pilots across municipal schools reduced grading time by 30%, enabling more frequent speaking assessments.

Safeguards include anonymised audio identifiers, deletion of recordings after grading, audit logs for every automated decision, DPDP-aligned consent protocols, and mandatory human verification. The case demonstrates how assessment automation can expand feedback frequency without relinquishing professional judgment.

BharatGen Yojaka is a human-in-the-loop AI assessment platform developed to scale spoken-language evaluation in Indian public-school programmes. It supports grade-level learning and encourages understanding rather than rote memorisation.

Bookbot: Offline AI-Powered Reading Tutor for Rural Indonesian Schools

Author: Adrian DeWitts

K-12

Summary

Bookbot is an offline AI-powered reading tutor deployed across Indonesian primary schools to address early-literacy gaps in low-connectivity environments. The platform combines phoneme-level speech-recognition models running entirely on local devices with a library of more than 1,800 levelled phonics books in Bahasa Indonesia. By delivering immediate pronunciation feedback without relying on continuous internet access, Bookbot enables guided oral-reading practice in classrooms that would otherwise be excluded from cloud-based EdTech solutions.

Context

Indonesia's rural and island communities face persistent shortages of trained reading instructors, classroom books, and reliable broadband connectivity. Early-grade teachers are often responsible for large classes, making it difficult to provide the one-to-one reading support essential for developing fluency. While digital tools could help, most AI-based reading tutors depend on cloud processing, placing them out of reach for schools with unstable networks or limited budgets.

The Indonesian Ministry of Education sought scalable literacy interventions that could function offline, integrate with national curricula, and be deployed through standard Android devices rather than specialised hardware. Indonesia struggles with early reading, especially in rural areas where children lack books, trained teachers, and reliable internet.

Most AI reading tools depend on constant connectivity and costly devices, which keeps many students out. Bookbot was created to solve this by acting as an offline digital reading tutor that listens to children read aloud and gives instant pronunciation feedback on low-cost devices. Its speech models are trained on diverse Indonesian accents and classroom noise, and all processing runs locally after download.

Teachers were trained to use the system, books were aligned with national curricula, and schools were closely involved during pilots. Results showed higher reading speed and accuracy, strong weekly use, and growing support from parents and teachers. Because it works offline, uses open materials, and keeps costs low, Bookbot can scale easily. The project shows that locally trained, offline-first AI paired with teacher support is key for literacy tools in low-resource settings.

Challenge

Bookbot needed to operate fully offline after installation while delivering fine-grained, phoneme-level pronunciation feedback. The system also had to function reliably in noisy classrooms and enable teachers to track student progress without requiring advanced digital skills. At the same time, child-data protection and ministry-level governance requirements had to be satisfied for nationwide deployments.



Case Study 6

Solution & Impact

Bookbot runs compact speech-recognition models locally on Android tablets, Chromebooks, and Windows laptops, analysing student speech in real time and flagging mispronunciations word-by-word. Automatic noise-classification systems filter classroom interference, while teacher dashboards aggregate progress data for instructional planning.

The system has been deployed across 19 Indonesian provinces, training hundreds of teachers and operating on thousands of devices. Usage analytics show strong engagement, millions of books read annually, and reading-speed gains correlated with time spent on the platform.

By delivering immediate pronunciation feedback without relying on continuous internet access, Bookbot enables guided oral-reading practice in classrooms that would otherwise be excluded from cloud-based EdTech solutions.

Governance frameworks include parental consent obtained through ministry channels, device-level data storage, and the ability for families and schools to delete accounts at any time. Only anonymised datasets are shared with research partners under ethics approvals. Bookbot demonstrates that sophisticated AI can function effectively even in low-infrastructure contexts when engineered for local constraints.

Chimple: Generative AI Enables Contextually-Relevant Learning Games for Foundational Learning

Authors: Srikanth Talapadi, Manisha Kukreja, Ashay Agrawal

K-12

Summary

Chimple addresses the shortage of engaging, local-language foundational learning content in India's early education system. Through Chimple's GenAI Studio, educators can create curriculum-aligned learning games using simple prompts, without technical skills. This enables faster content creation tailored to local languages and contexts. The solution is currently reaching around 500 students and is planned to scale to over 30,000 learners across 7+ Indian states.

Context

Foundational literacy and numeracy are critical for long-term learning outcomes, yet many children in India struggle due to limited access to engaging and locally relevant content. While national curricula such as NCERT provide structure, teachers often lack the time and technical capacity to convert textbook material into interactive digital learning experiences.

These challenges are more pronounced in multilingual, rural, and low-resource settings, where mother-tongue instruction and cultural relevance are essential for early learners. Creating high-quality digital learning games typically requires design and technical expertise, restricting educators' ability to adapt content to their students' linguistic and social contexts. Chimple operates within this landscape as an open-source, gamified learning platform used across India and Africa.

Challenge

At the foundational level, learning outcomes are closely tied to engagement, language familiarity, and cultural relevance. However, most digital content is produced centrally, often in limited languages, and cannot be easily adapted by teachers. This results in low engagement, especially for first-generation learners and children in non-English or non-Hindi settings.

Teachers and curriculum developers face significant constraints in time, skills, and resources. Manual creation of digital games or interactive content is slow and requires specialised expertise, making it difficult to respond to local needs or rapidly update materials. As a result, textbooks remain underutilised, and learning experiences fail to reflect students' lived realities. The challenge is to



Students engaging with Chimple's gamified learning modules. (Source: Chimple website)



Case Study 7

enable educators to quickly create high-quality, curriculum-aligned, and localised digital content at scale, without increasing workload or requiring technical training.

Solution & Impact

Chimble addresses this challenge through its GenAI Studio, a generative AI-powered content creation tool integrated into the Chimble learning platform. GenAI Studio allows non-technical educators and curriculum developers to generate interactive, curriculum-aligned learning games using simple text prompts.

Educators can transform textbook concepts into engaging games adapted to local languages, cultural references, and learner needs. The platform supports foundational literacy, numeracy, and socio-emotional learning, and enables rapid iteration and customisation of content. All AI-generated outputs are reviewed by educators before use, ensuring alignment with learning outcomes, age appropriateness, and cultural sensitivity.

GenAI Studio has been deployed across multiple use cases, including NCERT Grade 3 content in government primary schools, socio-emotional learning modules developed with a Juvenile Justice Home in Mumbai, and Portuguese-language content for Grades 4-6 in Africa. These deployments currently reach approximately 500 students.

In its initial phase, GenAI Studio has been used by 5 educators and curriculum developers, demonstrating strong time savings and increased flexibility in content creation. Based on this foundation, Chimble plans to expand the studio to over 50 educators and reach more than 30,000 active learners across 7+ Indian states within six months.

Chimble's GenAI Studio allows non-technical educators and curriculum developers to generate interactive, curriculum-aligned learning games using simple text prompts.

The platform is designed for inclusion and scale. It supports multilingual classrooms, first-generation learners, anganwadis, childcare institutions, and low-resource settings through adaptable visual and audio formats. As an open-source system, it avoids licensing costs and enables cost-effective expansion. Chimble's GenAI Studio demonstrates how generative AI, combined with human oversight, can strengthen foundational learning by empowering educators to create locally meaningful learning experiences at scale.

CrazyGoldfish Technologies: AI-powered assessment engine reduces evaluation time by up to 60% in subjective learning assessments

Author: Rahul Khandelwal

K-12

Summary

India's schools often struggle to evaluate handwritten responses quickly, slowing feedback and remediation. This shared AI assessment engine helps teachers review subjective answers faster by interpreting handwriting, suggesting rubric-based scores, and generating structured feedback. In a PM Shri pilot with 240 students, evaluation time dropped by 50-60%, while AI-teacher score differences stayed within a margin of 3-5 marks, enabling faster learning support.

Context

Across India's K-12 system, handwritten assessments remain central to evaluating conceptual understanding. However, manual evaluation is slow and varies across teachers, making it difficult to provide timely feedback aligned with NEP 2020's focus on competency-based learning. Many schools, nonprofits, and learning platforms lack the technical capacity to build automated evaluation systems. Assessment workflows are often fragmented, paper-based, and difficult to scale. In low-resource environments, limited connectivity and heavy teacher workloads further delay feedback cycles. Organisations such as government schools, NGOs, and exam-preparation platforms rely on subjective assessments but struggle to manage high volumes efficiently while maintaining scoring consistency and transparency.

Challenge

Evaluating subjective handwritten responses is one of the most time-intensive tasks for educators. Teachers must manually read, interpret, and score answers, which often leads to long delays before students receive feedback. These delays reduce opportunities for timely remediation, particularly when learning gaps need to be addressed quickly.

Scoring can also vary across evaluators, making it harder to ensure fairness and alignment with competency-based learning goals. As education systems move toward more formative assessment models, the need for reliable, consistent evaluation becomes more urgent. Yet many institutions lack the engineering capacity to build scalable AI solutions on their own.

Traditional workflows also struggle to handle large volumes of scripts, especially in programmes serving underserved learners or exam-focused settings where written reasoning is essential. Low-connectivity environments, non-standardised marking schemes, and varied handwriting quality further complicate evaluation processes. Without automated support, teachers face heavy workloads, backlogs accumulate, and students often receive feedback too late to improve learning outcomes effectively.



Case Study 8

Solution & Impact

CrazyGoldFish Technologies developed a shared AI assessment engine designed to support evaluation workflows across schools, nonprofits, and learning platforms. The system interprets handwritten responses, suggests rubric-aligned scores, enables rapid re-evaluation, and generates structured personalised feedback. Rather than replacing teachers, it functions as a decision-support layer: educators retain control over final grading through human-in-the-loop verification and editable scoring.

The engine was piloted in a PM Shri school in Srikakulam, Andhra Pradesh, where 12 teachers evaluated handwritten responses from 240 students over one month. During the pilot, evaluation time decreased by 50-60%, allowing teachers to focus more on remediation. AI-teacher score deviation remained within a margin of 3-5 marks, indicating strong alignment between automated suggestions and educator judgment. Faster feedback cycles enabled small-group corrections within 2-3 days, compared to previous delays of 10-14 days.

The solution has since been validated for broader deployment. Integration with WhatsApp workflows and vendor registration has prepared the system for rollout in Educate Girls' NEEV programme, which supports 6,000 out-of-school girls, demonstrating readiness for low-resource contexts. The engine has also been used by Arivu, a professional exam-preparation platform, to evaluate more than 650 pages from 45 learners in three days, showing applicability in dense, reasoning-heavy assessments.

This shared AI assessment engine helps teachers review subjective answers faster by interpreting handwriting, suggesting rubric-based scores, and generating structured feedback.

Designed for scalability, the engine supports offline-first workflows, secure data handling, and auditable override logs to maintain transparency. Procedural safeguards, including educator audits and role-based access controls, help manage risks and maintain trust. By reducing manual workload while preserving teacher oversight, the platform offers a scalable pathway to improving assessment quality, accelerating feedback, and strengthening competency-based learning across diverse education settings.

Darsel: Personalized and Adaptive Math Learning through AI chatbot

Author: Abdulhamid Haidar

K-12

Summary

Darsel is a global education-technology nonprofit that uses AI-powered chatbots delivered through messaging platforms such as WhatsApp and Telegram to provide personalised mathematics practice to students in low-resource environments. By meeting learners on simple household phones, Darsel bypasses device and connectivity barriers that limit most EdTech tools. Its adaptive questioning, hint systems, and gamified interactions deliver curriculum-aligned remediation at extremely low marginal cost, enabling governments to scale numeracy interventions across thousands of schools.

Context


Worldwide, six in ten children fail to reach minimum proficiency in mathematics by the end of primary school. In many government-run systems, classrooms are overcrowded, teachers lack time for individual remediation, and households cannot afford specialised learning devices. Most digital interventions assume smartphones, apps, or stable broadband - excluding precisely the communities most in need of foundational support.

Darsel emerged to exploit a simple but powerful fact: even in low-income households, messaging apps are ubiquitous. If adaptive learning could be delivered through these channels, personalised tutoring could reach populations long excluded from traditional EdTech deployments.

Challenge

Across the world, many children fail to reach basic proficiency in mathematics, and early gaps in numeracy grow larger over time if not addressed. Government school systems in low-resource settings face severe constraints: large class sizes, limited budgets, and shortages of trained teachers make it difficult to provide personalised support at scale. While technology could help, most existing EdTech platforms assume reliable internet access, smartphones, and high data usage - conditions that are often absent in the communities most in need.

Another major barrier is curriculum misalignment. Many digital tools are built for international markets and do not match local syllabi or languages, making them difficult for public education systems to adopt. Teachers also lack real-time insight into where students are struggling, reducing their ability to target instruction effectively.



A student engages with Darsel's AI-powered Math learning chatbot. (Source: Darsel website)



Case Study 9

These structural challenges leave millions of learners without the tailored practice needed to master foundational maths skills. Without affordable, low-bandwidth, curriculum-aligned solutions that fit into existing school systems, public education authorities struggle to close learning gaps at population scale, deepening inequality between students in well-resourced and underserved environments.

Solution & Impact

Darsel's chatbot sends students tailored questions, follows up with hints and explanations, and adjusts difficulty based on performance patterns. Gamification features encourage persistence, while teachers and administrators monitor progress through dashboards.

The platform now serves over 350,000 students across 2,700 schools in India, Jordan, Nigeria, and beyond. Randomised evaluations report 0.25 standard-deviation improvements in mathematics performance at an annual marginal cost of roughly one dollar per learner.

Darsel bypasses device and connectivity barriers that limit most EdTech tools. Its adaptive questioning, hint systems, and gamified interactions deliver curriculum-aligned remediation at extremely low marginal cost, enabling governments to scale numeracy interventions across thousands of schools.

Data governance relies on end-to-end encryption through messaging platforms, strict role-based access controls, anonymised identifiers, and parental or school consent obtained through government partnerships. Outputs are constrained to mathematical domains and reviewed through validation pipelines. Darsel illustrates how AI tutoring can be delivered at scale in low-resource settings when engineered around everyday technologies.

Drishti: On-Device AI Boosts Foundational Learning in Multilingual Classrooms

Author: Madhavi Kaivalya Kandalam, Sailaja Raman

K-12

Summary

Drishti tackles the challenge of low foundational literacy among children aged 3-10 in low-resource public schools. Using on-device AI and “Draw-to-Learn” technology, children receive instant, personalised feedback in their native language without needing the internet. This edge-first approach improves engagement by 40% and delivers accurate sketch recognition across 2,000 objects with 95% top-5 accuracy, enabling scalable, inclusive learning in rural schools.

Context

Early childhood education is critical for long-term learning outcomes, yet many children in India's rural public schools start without sufficient literacy or numeracy skills. Classrooms are overcrowded, teachers have limited time for one-on-one guidance, and internet access is often unreliable. Traditional EdTech relies on cloud-based tools and English-centric content, leaving children in rural areas underserved. Learning often occurs via printed worksheets or rote classroom activities. Drishti operates within this context, using low-cost tablets to deliver interactive, bilingual learning experiences offline. Its design fits the realities of public schools, supports teacher oversight, and requires minimal infrastructure, enabling children to learn in their home language.

Challenge

Children in foundational years require immediate, contextual feedback to reinforce learning, but overcrowded classrooms and limited teacher

capacity prevent this. Traditional digital tools often rely on cloud AI, which is costly, slow, and internet-dependent, excluding children in low-connectivity regions. Manual feedback through worksheets is slow and inconsistent, reducing engagement and retention. Educators face constraints in time, training, and resources, while existing EdTech rarely provides culturally or linguistically relevant content. There was a need for a solution that could provide instant, personalised feedback, operate offline, respect children's languages, and scale across thousands of schools without adding extra burden on teachers.

Solution & Impact

Drishti delivers an edge-first, multimodal AI solution for foundational literacy and early learning. Children interact with the system by drawing objects on tablets, and on-device AI instantly recognises sketches and provides positive reinforcement. AI-enabled speech-processing offers pronunciation feedback in both the child's native language (e.g., Telugu) and English. Teachers monitor progress via a simple dashboard without requiring technical expertise.

The AI runs primarily on the local device (i.e., edge-first), eliminating internet dependency and ensuring data privacy. During PoC testing, the system achieved 95% top-5 accuracy in recognising 2,000 object classes, while engagement increased by 40% compared to passive screen-viewing. Children respond



Case Study 10

positively to instant, multilingual feedback, and teachers can focus on learners who need the most support.

Drishti is deployed in Telangana public schools under the “TG Tots” initiative, pre-installed on low-cost tablets. It is validated by the Telangana AI Mission and shortlisted for international recognition. With scalability in mind, Drishti is being expanded to include maths and science modules and reviewed for pilots across other linguistic regions. Its offline, teacher-centric, and culturally relevant design makes it a replicable model for the Global South.

Using on-device AI and “Draw-to-Learn” technology, children receive instant, personalised feedback in their native language without needing the internet.

Directorate of Higher Education, Government of Goa: State-led AI capacity building improves teaching quality across higher

Authors: Dr. Kavita Asnani

Higher Education

Summary

The Government of Goa addressed gaps in meaningful AI use among higher education faculty through a state-led capacity-building model. While many faculty were already experimenting with AI tools, use was largely informal and unstructured. The Directorate of Higher Education (DHE), Government of Goa, introduced a discipline-specific AI pedagogy programme focused on classroom application rather than technical expertise. The initiative trained 258 faculty across 52 institutions, reduced course preparation time by around 20%, and increased the use of activity-based and student-centred teaching practices.

Context

Higher education institutions in India are under pressure to improve teaching quality, student engagement, and learning outcomes while implementing the National Education Policy (NEP) 2020. Faculty often face large class sizes, diverse student needs, and limited time for course design and feedback. Although AI tools have become widely accessible, their use in higher education is uneven and often disconnected from pedagogy.

The Directorate of Higher Education, Government of Goa, oversees 65 higher education institutions, with over 3,000 faculty members & more than 40,000 students. As the state's nodal agency for NEP implementation, DHE operates within a regulated, system-wide environment where changes in teaching practices must be scalable, ethical, & aligned with academic standards.

Challenge

A baseline survey of 330 faculty members revealed that while over 70% were already using AI tools informally, fewer than 52% were integrating them meaningfully into teaching, assessment, or feedback. Faculty expressed concerns about student misuse of AI, academic integrity, and lack of clarity on ethical boundaries. Many were unsure how to use AI to improve learning outcomes rather than simply save time.

Traditional faculty development programmes often focus on tools or technologies in isolation, require significant time investment, and do not translate into sustained classroom practice.

There was also wide variation across disciplines, with science, law, humanities, and language departments facing very different pedagogical needs.

The core challenge was to build AI capacity across the higher education system without turning faculty into technical experts, while ensuring ethical use, discipline relevance, and consistent adoption across institutions.



Case Study 11

Solution & Impact

DHE designed a structured, state-led AI capacity-building model focused on pedagogy rather than technology. The model followed a closed-loop approach: training, classroom implementation, structured follow-up, and scaling. AI tools such as ChatGPT, Gemini, Claude, LLaMA-3, PhET simulations, Quizizz AI, and GitHub Copilot were mapped directly to teaching tasks like course planning, lesson design, assessment creation, feedback, and personalised learning support.

In the first phase, DHE launched a foundational AI pedagogy programme that emphasised responsible and ethical AI use. Faculty received a structured AI Pedagogy Handbook with ready-to-use prompts aligned to common academic tasks, including syllabus design, outcome-based education mapping, assessments, and feedback for slow or struggling learners. This ensured immediate classroom usability and reduced the learning curve.

To enable scale, DHE adopted a Train-the-Trainers model. Two faculty members from each discipline were trained as Master Trainers, who then trained peers within their institutions. Disciplines were grouped into six clusters to ensure contextual relevance, including science, commerce and management, applied sciences, health sciences, social sciences, and law, humanities, and languages.

After training, faculty were formally instructed to deploy AI tools in live courses. DHE conducted structured follow-up sessions during the semester to resolve issues, monitor adoption, and reinforce responsible AI use. This helped shift AI use from experimentation to sustained practice.

The Government of Goa addressed gaps in meaningful AI use among higher education faculty through a state-led capacity-building model by introducing a discipline-specific AI pedagogy programme focused on classroom application rather than technical expertise.

The programme trained 258 faculty across 52 institutions in the first phase, covering 69 disciplines. Faculty reported a 20% reduction in course preparation time, increased use of activity-based pedagogy, and higher levels of student engagement, including in traditionally non-technical disciplines such as humanities and law. The initiative is being expanded to train 500 faculty members, with the goal of full system-wide coverage.

The model embeds strong governance, ethical guidance, and inclusion, and is being further scaled through online courses under the state's digital education initiative. It demonstrates how a higher education regulator can systematically deploy and govern AI-enabled pedagogical transformation at scale.

Duolingo English Test: Expanding access to English proficiency assessment through AI-enabled testing

Authors: Rashi Dhanani, Tara Kapur

Higher education

Summary

The Duolingo English Test (DET) addresses access barriers in English proficiency testing by replacing centre-based, paper-driven models with a secure, fully digital, AI-enabled assessment. By enabling remote testing at lower cost and in shorter duration, the DET reduces the logistical and financial constraints faced by learners in Tier-2, Tier-3, and rural regions in India. Accepted by over 6,000 institutions globally and taken by candidates across more than 740 Indian cities, the DET demonstrates how AI-enabled assessments can scale efficiently while maintaining standardisation and reliability.

Context

English proficiency assessments play a central role in higher education admissions and international mobility. In India, most established English Proficiency Tests (EPTs) are conducted through physical test centres and fixed schedules. Candidates are often required to travel long distances, incur accommodation and registration costs, and test in unfamiliar environments. These constraints limit participation, particularly for learners outside major urban centres.

As demand for English proficiency certification has increased, especially for international education pathways, the limitations of centre-based testing models have become more pronounced. There is a growing need for assessment systems that can reach learners across geographies, operate at scale, and

deliver results quickly, while remaining aligned with institutional requirements for standardisation and comparability.

Challenge

Traditional English proficiency tests rely on in-person infrastructure and manual processes that are difficult to scale equitably. Test centre availability, scheduling rigidity, and administrative overhead create friction for candidates and institutions alike. For learners in smaller cities and towns, access is constrained by travel time, cost, and limited testing slots.

From an operational perspective, conventional assessments require significant coordination for test delivery, proctoring, scoring, and result processing. These factors increase costs, slow turnaround times, and restrict the ability of assessment providers to expand reach without proportionate increases in physical infrastructure.

The challenge is to deliver a standardised English proficiency test that reduces logistical dependence on physical centres, lowers cost and time burdens for candidates, and can be administered consistently at scale.

Solution & Impact

The Duolingo English Test addresses this challenge through a fully online assessment platform that integrates digital delivery and AI-enabled evaluation. Candidates can take the test remotely using a laptop and a stable internet connection, eliminating the need for travel to physical centres and allowing flexible scheduling.



Case Study 12

The DET employs AI for test design, scoring, and quality assurance. It uses a computer adaptive testing (CAT) model that adjusts question difficulty in real time based on candidate performance, efficiently estimating proficiency, reducing test length to under an hour, and giving each candidate a unique set of questions. The system scores speaking and writing responses at scale, analysing features such as vocabulary, grammar, coherence, pronunciation, and fluency, while human reviewers handle quality assurance and edge cases. It maintains test integrity by monitoring video, audio, and interaction data and combining AI review with asynchronous human oversight, including dual-camera recording, which ensures secure and consistent administration. Together, these mechanisms deliver reliable and consistent results in high-stakes assessments.

The solution delivers these functions through a secure online platform accessible via standard digital devices. It incorporates digital onboarding, identity verification, secure test delivery, and AI-supported evaluation with human oversight. Candidates can access the platform on a laptop or desktop with a stable internet connection. Where needed, the platform can be deployed in partnership with educational institutions, government agencies, and authorised testing bodies, helping integrate the system into existing assessment frameworks and build institutional trust.

The DET also helps candidates save time and reduce costs. Results are typically available within 48 hours, reducing waiting times compared to traditional assessments. The test costs approximately one-third of many conventional English proficiency exams, lowering the overall cost of participation.

Accepted by over 6,000 institutions globally and taken by candidates across more than 740 Indian cities, the DET demonstrates how AI-enabled assessments can scale efficiently while maintaining standardisation and reliability.

In India, the DET has reached candidates in over 740 cities, including Tier-2 and Tier-3 locations, and in 2023 India became the largest source of DET test-takers globally. The test is accepted by more than 6,000 higher-education programmes worldwide, including all eight Ivy League universities. Institutions report increased applications from a wider range of geographies, reflecting the role of remote, digital testing in expanding access to formal credentialing pathways. Overall, the solution expands access to EPTs by providing consistent, scalable assessments that eliminate reliance on physical infrastructure and manual scoring, while significantly reducing time and cost for candidates.

Eedi Labs: AI-Powered Personal Tutors for Foundational Education

Authors: Dr Bibi Groot, Ben Caulfield

K- 12

Summary

Eedi Labs tests whether generative AI can safely scale personalised tutoring by embedding teacher-supervised language models into classroom maths instruction. In UK randomised trials, its LearnLM system generated Socratic prompts that teachers reviewed before students saw them. AI-supported learners matched or outperformed peers, especially on new problems, showing that carefully governed human-in-the-loop AI can deliver high-quality teaching at scale.

Context

One-to-one tutoring is widely regarded as the most effective form of instructional support, often generating learning gains equivalent to several months of additional schooling. However, its cost makes it inaccessible for most families and education systems. Generative AI has been proposed as a scalable alternative, yet concerns remain about hallucinations, pedagogical quality, bias, and student safety-particularly in classroom settings involving minors.

Most early deployments of AI tutoring tools have occurred outside rigorous research environments, leaving policymakers and school leaders without credible evidence about effectiveness or risk. Eedi, already operating a mathematics platform used by thousands of schools, sought to move beyond speculation by conducting controlled classroom trials.

Challenge

While pedagogically fine-tuned AI systems show promise in classroom tutoring, important limitations remain. Generative models still struggle to recognise the emotional state of learners and to judge when persistence becomes frustration rather than productive challenge. In some cases, students became disengaged when the AI continued questioning for too long, requiring human tutors to step in and adjust tone, pacing, or instructional approach.

These gaps highlight a broader challenge: even well-trained instructional models lack the human intuition needed to manage motivation, encouragement, and classroom dynamics. Educators must therefore remain closely involved, which raises questions about how to balance automation with human oversight at scale. Systems must be designed so that AI supports teaching without undermining relationships between tutors and learners.

Another unresolved issue is how to translate AI-generated instructional strategies into consistent professional development for teachers. While some tutors learned from observing the AI's questioning style, this benefit is informal and uneven. Creating structured ways for educators to gain value from AI tools - without increasing workload - remains a key design and implementation challenge for future deployments.



Case Study 13

Solution & Impact

The Eedi platform integrated LearnLM, a generative AI model trained specifically for teaching dialogue, into its mathematics tutoring workflows using a strict “human-in-the-loop” approach. LearnLM drafts Socratic-style prompts based on each student’s question, likely misconception, and predicted ability level, encouraging learners to reason through problems rather than receive direct answers. Expert tutors review every AI-generated message before it reaches students, retaining full authority to approve, edit, or replace the content.

This approach was tested through a randomised controlled trial in five UK secondary schools, where Year 9 and 10 students were assigned either static hints or interactive tutoring. The pilot involved 165 students and 17 tutors. Although exploratory in scale, the trial ran on Eedi’s existing infrastructure, which already serves hundreds of schools, supporting future expansion. Post-trial simulations suggested that the AI could increase tutor capacity by enabling more sessions per hour without reducing instructional quality.

Results showed that students supported by LearnLM performed at least as well as peers receiving human-only tutoring on the immediate re-try and were more likely to solve the next problem in the study unit. Tutors approved three-quarters of the AI’s drafts with no edits,

LearnLM drafts Socratic-style prompts based on each student’s question, likely misconception, and predicted ability level, encouraging learners to reason through problems rather than receive direct answers.

Educators also highlighted the system’s effectiveness in generating reflective questions that improved classroom practice.

Ethical safeguards were central to deployment. The trial received formal approval and required informed consent, with continuous oversight ensuring safety. Audits of thousands of messages found no harmful content and extremely low factual error rates. Participating schools represented diverse socio-economic backgrounds, and the AI adapted its language to different ability levels, supporting inclusive use.

In conclusion, the LearnLM pilot demonstrates that carefully governed, teacher-supervised generative AI can deliver personalised learning at scale while maintaining educational quality, safety, and trust.

EIDU: AI-Enhanced Teaching: Evidence on the Impact of an LLM-Driven Remedial Programme in Kenya

Author: Aidan Friedberg, Amar Lalwani, Harriet Crisp

Pre-K

Summary

EIDU enhances remedial literacy instruction in Kenyan pre-primary classrooms by combining diagnostic learning analytics with LLM-generated lesson plans that are aligned to national curricula. The hybrid system identifies students most in need of intervention and produces short, structured remedial activities for teachers to approve and deploy. Designed for low-connectivity environments, the platform demonstrates how generative AI can augment teacher capacity without replacing professional judgment.

Context

Early-grade literacy remains a major concern across many low- and middle-income countries, where large numbers of children enter school without strong reading foundations. Governments and NGOs increasingly use structured-pedagogy programmes and digital learning tools to support classrooms, but teachers still carry heavy workloads and have limited time for diagnosing individual learner needs. Many schools operate with only a few shared devices and unreliable connectivity, especially in rural areas. At the same time, education authorities are exploring how generative AI might support teaching, provided systems align with national curricula, protect child data, and function safely within existing classroom routines and low-resource environments.

Challenge

Many low- and middle-income countries continue to face large gaps in foundational literacy, especially in early grades. Although structured teaching and remedial instruction can improve outcomes, these approaches depend on frequent diagnosis of student needs and significant teacher time to design targeted lessons, both difficult to sustain in overcrowded, resource-constrained classrooms. Teachers often lack tools to systematically identify which children are falling behind or to group learners with similar gaps for focused support.

Digital systems could help address these problems, but most generative-AI models are built for high-connectivity environments and risk producing content that does not match national curricula or classroom realities. In rural schools, internet access is intermittent, devices are limited, and educators have little capacity to experiment with complex new tools. Any new system must therefore work offline or asynchronously, fit within existing teaching practices, and preserve teacher autonomy.



An educator uses EIDU's platform to deliver targeted remedial literacy activities.
(Source: EIDU website)



Case Study 14

There are also serious governance concerns. Schools must avoid inaccurate or inappropriate AI-generated material, protect young children's data, and ensure that teachers remain in control of instructional decisions. Finally, education systems need credible evidence that AI-supported remediation actually improves learning before scaling widely. Balancing instructional quality, safety, infrastructure limits, and proof of effectiveness remains a central challenge for applying generative AI in early-grade classrooms.

Solution & Impact

Learners interact with EIDU's digital activities, generating anonymised performance data that feeds into knowledge-tracing algorithms estimating mastery across foundational-literacy skills. The system flags up to five students per class for targeted remediation and triggers LLM pipelines to draft two alternative lesson plans per topic, grounded in Kenya's Tayari pedagogy programme.

Teachers review and approve plans before classroom use, ensuring pedagogical control. In 12-week pilots across 40 government schools, teachers completed an average of 1.48 remedial lessons per week and approved 91% of generated plans.

The hybrid system identifies students most in need of intervention and produces short, structured remedial activities for teachers to approve and deliver in the classroom. Designed for low-connectivity environments, the platform demonstrates how generative AI can augment teacher capacity without replacing professional judgment.

Data governance relies on device-level encryption, anonymisation, restricted analytics access, and informed consent processes. Prohibited-topic filters, bias audits, and teacher-in-the-loop safeguards underpin responsible deployment.

Government of Rajasthan: Improving Mathematics outcomes through AI-enabled personalised tutoring in Rajasthan

Authors: Dr. Saumya Jha, Sushil Kumar Agrawal, Bhaskar Mishra

K-12

Summary

PadhaiWithAI addresses low Class 10 Mathematics performance among government school students in Tonk district, Rajasthan. It supports nearly 12,000 students and their teachers through an AI-powered bilingual tutoring platform combined with academic dashboards. Implemented over six weeks, the initiative contributed to a 95.4% Mathematics pass rate and a 5 percentage point rise in first-division scores.

Context

Secondary education outcomes strongly influence students' access to science streams and future career options. In Rajasthan's government school system, many students struggle in Mathematics due to teacher shortages, uneven instruction quality, and foundational learning gaps. Instruction is largely textbook-driven, with limited scope for personalised practice or timely feedback. Schools operate in low-resource settings with varied connectivity and large class sizes. District administrations oversee academic performance but often lack real-time visibility into student progress. Digital tools, where used, are typically limited to basic portals and manual data entry, offering little support for adaptive learning or rapid academic decision-making.

Challenge

Class 10 Mathematics performance in Tonk district had remained persistently weak, limiting students' access to science-based academic pathways. While nearly 90% of students aspired to science careers, only around 12% were able to enter science streams, primarily due to low Mathematics scores. Fewer than one in four students achieved first-division results in the subject.

Several constraints contributed to this challenge. Teacher shortages and frequent diversion of teachers for administrative and election duties reduced consistent instructional time. Students entered secondary grades with significant foundational gaps; more than half showed weaknesses at Class 5-8 Mathematics levels. Seasonal absenteeism further disrupted learning continuity.

Traditional teaching methods relied heavily on uniform classroom instruction and textbook problem-solving, leaving little room for personalised explanations or targeted practice. Teachers had limited capacity to diagnose individual learning gaps or provide differentiated support at scale. Administrative monitoring relied on delayed and fragmented data, making it difficult to intervene early or coordinate district-wide academic support. Together, these factors created a need for a scalable solution that could strengthen instruction, personalise practice, and improve system-level visibility without increasing teacher workload.



Case Study 15

Solution & Impact

PadhaiWithAI was designed as a district-led, AI-enabled personalised tutoring and academic governance solution to strengthen Mathematics learning where instructional capacity was constrained. The platform uses generative AI to provide bilingual, step-by-step explanations of textbook problems, generate unlimited practice questions with adjustable difficulty, and support targeted revision aligned to the Class 10 curriculum. AI acts as an assistive instructional layer, complementing teachers rather than replacing classroom pedagogy, while enabling students to practise independently at their own pace and reinforce classroom learning.

Alongside student-facing support, the platform includes dashboards for teachers, school leaders, block officers, and district administrators. These dashboards provide near real-time insights into attendance, assessment performance, and common learning gaps. A standardised academic calendar was implemented, including two daily Mathematics periods, AI-supported homework, and weekly district-wide tests. Marks were uploaded within fixed timelines, enabling faster feedback loops and coordinated academic oversight.

The initiative was developed and operated by the District Administration of Tonk and deployed through a customised web portal, without requiring major changes to existing school processes. All 353 government senior secondary schools in the district were onboarded, reaching 11,977 Class 10 students over a focused six-week period.

Despite the short duration, the initiative delivered measurable gains. Tonk district recorded a 95.4% Mathematics pass rate in the 2025 RBSE Class 10 examinations, representing a 3 percentage point improvement over the previous year.

PadhaiWithAI was designed as a district-led, AI-enabled personalised tutoring and academic governance solution to strengthen Mathematics learning where instructional capacity was constrained.

The share of students achieving first-division scores increased from around 23% to over 28%, while the proportion of high achievers scoring above 80% rose by nearly 2 percentage points. These improvements exceeded state-level trends, suggesting impact beyond general exam-year variation.

The model demonstrates strong potential for scale, as it operates on existing infrastructure, supports low-bandwidth environments, and retains human oversight in academic decision-making. It offers a replicable approach for district-led AI adoption in public education systems seeking rapid, equitable learning gains.

SATHEE : AI-Powered Multilingual Learning Platform

Authors: Dhruv Garg, Rahul Garg, Prof. Amey Karkare

Higher-education

Summary

SATHEE (Self-Assessment Test and Help for Entrance Exams) is a multilingual AI platform that uses AI to personalise learning and support teachers through multilingual content creation, voice-based assessments, and knowledge graphs aligned to major government exams such as JEE, NEET, and SSC. It offers adaptive tutoring for students, automated materials for teachers, and analytics dashboards for administrators, promoting education equity in rural India across seven regional languages.

Context

India's linguistic diversity and exam-focused coaching system leave many rural students without vernacular materials or communication-skills training. To address this, the Ministry of Education partnered with IIT Kanpur to build SATHEE, a multilingual AI platform that transforms content creation, delivery, and assessment across seven Indian languages, with interfaces in Hindi and Punjabi. SATHEE uses computer vision, speech systems, NLP, knowledge graphs, and RAG chatbots to generate teacher notes and slides, run voice-based assessments with rubric scoring, create interview-practice modules, and convert lectures into multi-format learning resources. Deployed across 569 schools and serving over 170,000 students, the system has produced thousands of lesson plans, handled more than 220,000 voice interactions, reduced teacher

preparation time by 70%, and driven major exam gains - 48.4% in NEET and 73.5% in JEE - while maintaining strong governance through anonymisation, consent, and linguistic-bias audits.

Challenge

The central challenge was delivering holistic, multilingual exam preparation - covering subject mastery, interview readiness, and teacher support - without imposing new infrastructure burdens or compromising data privacy across language groups.

Solution & Impact

SATHEE was designed as a modular, multimodal AI platform. Computer-vision pipelines tag lecture videos into concept-level segments; speech-to-text engines transcribe student responses; OCR systems digitise handwritten notes and exam papers; and text-to-speech modules generate vernacular study materials. Knowledge graphs encode syllabus structures and inter-topic dependencies, enabling adaptive diagnostics and revision planning.

It is built as a multilingual AI platform so students and teachers can learn and work in their own languages. Its website runs in Hindi and Punjabi, while video lessons and learning materials are available in seven Indian languages. AI tools turn video lectures into teacher notes, worksheets, and slide decks, cutting teacher preparation time by about 70%.



Case Study 16

The system has already processed hundreds of hours of content and created thousands of lesson plans and presentations.

It also helps students prepare for interviews and exams by generating practice questions in different fields and languages. Voice-based assessments read questions aloud and record spoken answers, while AI checks understanding and language use and gives feedback. A smart chatbot searches textbooks and lectures to give exact answers with video references. Knowledge maps link school concepts so learning follows the right sequence.

In Punjab alone, SATHEE has served over 170,000 students across 569 schools, reporting 48.4% improvements in NEET scores and 73.5% gains in JEE outcomes.

The platform now supports large numbers of users, with strong exam results, high preference for local languages, and big improvements in mother-tongue understanding. Privacy protections, consent, and bias testing are built in. Overall, SATHEE shows how multilingual, locally designed AI can expand access to quality education across India at scale.

SATHEE is a multilingual AI platform that transforms content creation, delivery, and assessment across seven Indian languages, with interfaces in Hindi and Punjabi.

Governance mechanisms include anonymised language-preference tracking, deletion of raw audio after transcription, federated learning pipelines for language models, and encrypted storage of student data. These design choices allow SATHEE to scale while remaining compliant with public-sector data obligations and child-safety standards.

Indra AI: AI-Powered Teaching Assistant for Content Creation and Multilingual Localization

Author: Shivam Sharma

K-12

Summary

Indra AI develops generative-AI teaching assistants that reduce teacher workload while expanding access to high-quality multilingual learning materials in India's K-12 classrooms. Integrated with open-source digital-education platforms, the system produces curriculum-aligned lesson plans, worksheets, quizzes, and stories in multiple Indian languages enabling activity-based learning at scale while preserving public-sector governance and data sovereignty.

Context

India's K-12 education system serves highly diverse linguistic communities while operating under high student-teacher ratios, often exceeding 40 students per classroom. Many teachers spend large amounts of time on lesson planning and administrative work rather than individual instruction. High-quality digital learning content is still concentrated in English, creating a major disadvantage for students who study primarily in regional languages. At the same time, government platforms such as DIKSHA and Sunbird form the backbone of national digital education efforts. Any innovation must therefore integrate with public infrastructure, function in low-bandwidth settings, and support multilingual classrooms at scale.

Challenge

India faces a severe "learning-poverty" crisis, with over half of Grade-3 students lacking basic literacy and numeracy. Overcrowded classrooms make it difficult for teachers to provide individual attention, while shortages of regionally relevant teaching materials force heavy reliance on rote instruction. Teachers are overburdened with repetitive planning tasks and manual content creation, leaving little time for remediation or creative pedagogy.

Language further deepens inequity. Most digital resources are English-first, excluding millions of rural and semi-urban learners who are more comfortable in state languages. Schools also lack scalable ways to generate grade-appropriate, culturally relevant materials aligned to national curricula.

Education systems face additional governance concerns. AI-generated materials must be accurate, bias-free, curriculum-aligned, and safe for classroom use. Teachers need to trust such tools rather than fear over-automation, while administrators require oversight and compliance with responsible-AI frameworks.

Without new approaches, public systems struggle to personalise instruction, localise content, and reduce teacher workload simultaneously - widening learning gaps rather than closing them.



Case Study 17

Solution & Impact

The AI-Powered Teaching Assistant was built as a “co-curator” for teachers, using generative AI to transform official textbooks and curriculum documents into multilingual, classroom-ready learning materials. Developed on Sunbird’s open-source digital-public-infrastructure layer and powered by large language models deployed through cloud AI platforms, the system generates stories, real-world examples, quizzes, worksheets, and lesson plans aligned to grade levels and foundational-learning competencies. Automated translation and transliteration localise English resources into culturally relevant regional languages, enabling “Teaching at the Right Level” pedagogy while reducing preparation time.

The platform integrates directly with DIKSHA and Sunbird user systems, avoiding new accounts or parallel workflows. A central content repository stores approved materials for teachers and students, while a secure teacher workspace allows educators to prompt AI, review outputs, and provide AI-assisted feedback on student work. All generated content is grounded in curriculum-mapped sources and must be verified by teachers before classroom use. Micro-learning modules, guided prompts, and dashboards support onboarding and sustained adoption, while telemetry data allows states to monitor engagement and reuse across schools.

Strong governance is embedded through human-in-the-loop review, curriculum anchoring, bias audits, and compliance with national responsible-AI guidelines. The system is optimised for low-bandwidth environments, offline access, and multiple Indian languages, ensuring inclusion for teachers and learners in rural and semi-urban settings.

The system generates stories, real-world examples, quizzes, worksheets, and lesson plans aligned to grade levels and foundational-learning competencies. Automated translation and transliteration localise English resources into culturally relevant regional languages, enabling “Teaching at the Right Level” pedagogy while reducing preparation time.

Early pilots demonstrate meaningful gains. Teachers reported 30-40% reductions in non-instructional workload, lower stress, and more time for one-to-one mentoring. Students using AI-localised, activity-based materials showed roughly double the engagement compared with rote methods and better alignment with learning outcomes. These results suggest that embedding governed generative AI within India’s public digital infrastructure can strengthen instructional capacity at scale while advancing equity in foundational learning.

Khanmigo: AI-powered tutoring and teacher support improves learning at scale

Author: Deepak Agarwal, Akash Burlawar

K-12

Summary

Khanmigo addresses gaps in personalised learning and teacher capacity in K-12 education. Students often need step-by-step support, while teachers face time constraints in planning and tracking progress. Khanmigo is an AI-powered teaching and learning assistant integrated into Khan Academy. It supports students through guided tutoring and helps teachers with lesson planning and classroom insights. Globally, 80,000-100,000 students and 20,000-25,000 teachers use Khanmigo weekly. In surveys, 76% of students report enjoying learning with Khanmigo, and teachers report saving 2-4 hours per week.

Context

K-12 education systems globally, including in India, face persistent challenges related to large class sizes, uneven access to quality tutoring, and limited teacher time. Teachers are expected to cover the curriculum, support diverse learning needs, assess progress, and manage administrative tasks, often with minimal support. Students, especially those without access to private tutoring, struggle to get timely help when they are confused or make mistakes.

Digital learning platforms are increasingly used to supplement classroom instruction and home learning. Khan Academy is a widely trusted platform already embedded in many school systems and learning programmes. Khanmigo operates within this existing ecosystem, extending current digital learning models rather

than replacing them, and is accessible across multiple languages and geographies, including low-resource settings.

Challenge

Effective learning requires timely feedback, personalised explanations, and opportunities for students to reason through problems. However, in most classrooms, teachers do not have the time or capacity to provide one-on-one support to every student. This results in gaps in understanding, reduced confidence, and over-reliance on rote memorisation.

At the same time, teachers face increasing demands on their time. Lesson planning, creating assessments, tracking student progress, and adapting instruction to different learning levels are often manual and time-consuming. Past digital tools have struggled with adoption because they added complexity, required extensive training, or operated outside existing teaching workflows.

The challenge is to provide personalised academic support to students and meaningful productivity gains to teachers, at scale, without increasing workload, disrupting classroom practices, or introducing unsafe or untrusted technology into schools.

Solution & Impact

Khanmigo addresses these challenges through an AI-powered assistant, built on state-of-the-art large language models (LLMs), that supports both students and teachers within the Khan Academy



Case Study 18

platform. For students, Khanmigo acts as a patient, always-available tutor. It guides learners step by step using questions and hints rather than direct answers, helping students identify and correct mistakes on their own. Khanmigo's Socratic approach encourages students to reason through problems rather than jump to answers, supporting deeper understanding instead of rote learning. Khanmigo is able to engage even when student questions are incomplete or contain spelling errors, including in Hindi and English, the languages used in Indian pilots.

For teachers, Khanmigo functions as a classroom assistant and productivity tool. It offers ready-to-use prompts for common tasks such as lesson planning, generating quizzes, creating lesson hooks, summarising topics, and reviewing class progress. These tools require minimal input and little training, reducing cognitive load and making adoption easier. Usage data shows that teachers frequently use tools like Class Snapshot, Refresh My Knowledge, Multiple Choice Quiz, and Lesson Hook. Surveys indicate that these tools save teachers approximately 2-4 hours per week, time that is redirected toward teaching and student engagement.

Khanmigo is developed and operated by Khan Academy and is fully deployed globally, with recent pilot programmes launched in India. It is currently available in over 30 languages, including Hindi, Odia, and Marathi, and has been rolled out across several Indian states such as Uttar Pradesh, Karnataka, Maharashtra, and Odisha. More than 2,000 teachers and students in India use Khanmigo regularly.

For students, Khanmigo acts as a patient, always-available tutor. It guides learners step by step using questions and hints rather than direct answers, helping students identify and correct mistakes on their own.

Trust and safety are central to Khanmigo's design. Student interactions are linked to classroom accounts, allowing teacher oversight. Inappropriate content triggers alerts, the system remains grounded in age-appropriate curriculum, and student data is not used to train AI models. In surveys, 80% of teachers report that Khanmigo is safe to use in classroom settings.

Khanmigo demonstrates how AI can scale personalised support in education while keeping teachers central to the learning process. By integrating into existing platforms and prioritising safety, usability, and human oversight, it offers a model for responsible AI adoption in schools.

Kibo: AI-Powered Technology for Inclusive Education

Authors: Akshita Sachdeva, Bonny Dave, Akhil Battula

K-12

Summary

Kibo (“Knowledge in a Box”) is an AI-powered assistive-technology device developed by Trestle Labs to enable visually impaired and print-disabled learners to independently access printed educational materials in real time. The tabletop device integrates computer vision, multilingual optical character recognition, document-structure analysis, and neural text-to-speech engines to convert textbooks, exam papers, classroom handouts, and handwritten notes into spoken output across more than sixty languages. Designed for deployment in schools, libraries, and testing centres, Kibo replaces slow centralised accessibility pipelines with point-of-use conversion, allowing learners to participate alongside peers without waiting for specialised formats. Its adoption across hundreds of institutions demonstrates how hardware-embedded AI systems can mainstream accessibility when paired with strong institutional governance and privacy controls.

Context

In many low- and middle-income countries, fewer than five percent of published textbooks are available in accessible formats. Braille production requires specialised printers and long lead times, while audiobook creation depends on studio recording and manual editing. These delays force visually impaired learners to rely on teachers, volunteers, or classmates for reading assistance, undermining autonomy and limiting participation during lessons or exams.

The problem is compounded in multilingual education systems where materials must be converted into several languages and scripts. Rural schools often lack trained special-education staff, and ministries struggle to maintain central repositories of accessible content that keep pace with curriculum changes.

Trestle Labs founded Kibo to decentralise accessibility - shifting conversion from distant production centres to the classroom desk itself.

Challenge

Kibo’s designers faced a range of operational constraints. The system had to instantly convert printed and handwritten documents, support dozens of scripts and languages, and function reliably without constant internet access. It also needed to protect highly sensitive materials such as exam papers, remain affordable for public procurement, and ensure that learner data



Learners using Kibo’s assistive-reading device. (Source: Trestle Labs website)



Case Study 19

stayed under institutional control. Any weakness in accuracy or privacy could undermine trust and prevent adoption in government systems.

Solution & Impact

Kibo integrates document-intelligence pipelines that detect layout, tables, equations, and handwriting before routing content through script-specific OCR models.

Language-identification engines select appropriate neural voices for text-to-speech output, while on-device caching allows continued operation during outages.

Deployed across 850+ schools, universities, and libraries, Kibo now supports 200,000+ learners, converting millions of pages annually.

Independent audits report high recognition accuracy across Indian and international scripts.

Kibo (“Knowledge in a Box”) is an AI-powered assistive-technology device developed by Trestle Labs to enable visually impaired and print-disabled learners to independently access printed educational materials in real time.

Governance is central to the design: documents are processed locally by default, encrypted at rest, and never uploaded without institutional permission. Consent workflows govern individual use, audit logs track every conversion, and procurement contracts preserve school ownership of all data. Kibo illustrates how assistive AI can scale when accessibility and governance are treated as inseparable design requirements.

Lenovo Leap NextGen Scholar Program: AI for Workforce Readiness

Authors: Poulamee Dey, Rajat Chandra, Shrilakshmi Nair ——— Higher-education

Summary

Lenovo Leap responds to India's widening gap between formal education and workforce readiness by delivering AI-enabled, industry-aligned learning pathways for college students and persons with disabilities. The programme combines adaptive digital platforms, hands-on innovation labs, and mentored project-based learning to build practical skills in emerging technologies. By embedding inclusion and accessibility into programme design, Lenovo Leap demonstrates how private-sector capabilities can support equitable skill development at scale.

Context

India faces a major skills gap, with fewer than 5% of workers formally trained and only 40% of technical graduates considered industry-ready. The Lenovo Leap: NextGen Scholar Program responds through an AI-enabled education initiative reaching 10,000 college and polytechnic students, including 1,000 persons with disabilities. Built around project-based learning, a digital innovation lab, and inclusive training, the eight-week hybrid programme combines adaptive learning platforms, workshops, and mentored projects in AI, data analytics, and full-stack development. With over 2,000 student projects completed, 40% women participation, and accessibility features for PwD learners, the initiative partners with government institutions to bridge the education-employment gap while advancing inclusive digital transformation.

Challenge

India faces a serious skills gap in its workforce. Only a small share of workers receive formal training, and many technical graduates are not considered ready for industry roles. Colleges and polytechnics often struggle to keep curricula aligned with fast-changing technologies such as artificial intelligence, data analytics, and software development, leaving students without practical, job-relevant experience.

Students from smaller cities and disadvantaged backgrounds face even greater barriers, with limited access to modern labs, mentorship, and industry exposure. Women and persons with disabilities are particularly underrepresented in technical fields, reflecting deeper structural gaps in access to quality training.

Education institutions also lack scalable ways to personalise instruction, connect learning with real-world problem-solving, and track student readiness for employment. Without stronger links between classrooms and industry needs, graduates risk entering the workforce unprepared, slowing innovation and widening inequality. Addressing these challenges requires new approaches that can expand access to future-ready skills while supporting diverse learners across India.



Case Study 20

Solution & Impact

The Lenovo Leap platform delivers adaptive digital-skills training in artificial intelligence, data analytics, and full-stack development through a blended AI-enabled learning system.

Discriminative AI models track learner progress and generate performance analytics, while generative AI creates customised learning materials and project simulations.

Natural-language tools allow students to ask questions in real time, and computer-vision systems assist with evaluating project work.

Together, these components personalise learning paths based on individual pace, competency levels, and interests, strengthening both skill retention and real-world application.

The programme has reached over 10,000 learners, including 40% women and 10% persons with disabilities, with more than 2,000 industry-aligned projects completed. Thousands of projects have been completed, demonstrating applied skill acquisition. Data governance is managed through secure learning platforms with consent-based enrolment, role-based access controls, and accessibility safeguards.

The Lenovo Leap platform delivers adaptive digital-skills training in artificial intelligence, data analytics, and full-stack development through a blended AI-enabled learning system.

The initiative was anchored by a central Digital Innovation Lab and supported by more than a thousand employee mentors. It reached institutions across urban and semi-urban areas, with a focus on underserved public colleges. Primary beneficiaries included thousands of students, alongside teachers, institutional leaders, and state education departments.

Overall, Lenovo Leap demonstrates how carefully governed, inclusive AI-driven education - combined with public-sector partnerships and industry engagement - can prepare students for emerging digital careers while advancing national goals around employability and equity.

Luma Learn: AI Powering Quality Education in Rural and Underserved Communities

Authors: Louis James Barnard

K-12

Summary

Luma Learn is a mobile-based education platform created to reduce education gaps in rural and underserved communities in Africa and other emerging markets. Many students lack resources like tutors, and steady academic support, but most families still have basic smartphones and use WhatsApp daily. Luma Learn uses this reality to deliver affordable, curriculum-aligned learning through chat-based lessons, study guides, and teaching plans powered by an AI assistant that sticks closely to official syllabi. The platform supports students with personalised help, gives teachers ready-made digital resources, and helps parents stay involved in learning at home.

Context

In many African regions, access to textbooks, tutors, and consistent academic support remains limited, while mobile-phone penetration - particularly through WhatsApp - is widespread. Traditional EdTech platforms requiring app downloads, high-bandwidth connections, or specialised devices exclude precisely the learners most in need of support.

Education ministries and NGOs increasingly seek low-friction digital interventions capable of operating across multiple languages and infrastructural contexts without overwhelming teachers or families.

Challenge

Education systems serving rural and underserved communities face persistent structural barriers that limit learning outcomes and widen inequality. Many learners lack access to textbooks, tutors, and consistent academic support, while teachers manage overcrowded classrooms and heavy administrative workloads that reduce time for individual instruction. Although mobile phones are widely available, most digital-learning tools are designed for high-bandwidth environments, standalone apps, or English-first curricula, making them poorly suited to local conditions.

Language diversity further complicates delivery. Students often learn in second or third languages, yet teaching materials and assessments are rarely adapted to local linguistic contexts, lowering comprehension and engagement. Parents, who play a critical role in supporting learning at home, have limited visibility into what children are studying or how to help them improve.

From an institutional perspective, schools and NGOs struggle to scale personalised tutoring without significant cost, training, or infrastructure investments. Traditional approaches-printed worksheets, after-school classes, or generic digital platforms-cannot provide real-time feedback or adapt to individual needs at population scale. These constraints



Case Study 21

create a widening gap between learners who can access high-quality support and those in low-resource settings, driving demand for new, scalable models that improve learning without adding complexity for teachers or families.

Solution & Impact

The platform combines conversational AI tutors with structured content repositories drawn from approved textbooks and exam materials. AI assistants guide learners through concepts, generate revision schedules, and support exam preparation while drawing exclusively from vetted sources. Luma Learn is built with privacy-by-design at its core, using domain-driven architecture and strict data boundaries to protect learner information. Data is encrypted both at rest and in transit, stored within isolated network environments, and accessed only through role-based permissions tailored to learners, parents, and teachers. Integrations with WhatsApp Business and OpenAI APIs operate through encrypted connections and secured keys, while onboarding controls ensure authorised access and account-level isolation.

The platform embeds strong governance practices, including informed-consent capture, clear data-retention policies, and anonymisation of all personally identifiable information. Data residency requirements are respected in South Africa and Angola, and responsible-AI principles are enforced through syllabus-aligned outputs and expert validation processes. Together, these safeguards allow Luma Learn to operate at scale across multiple countries while maintaining regulatory compliance and institutional trust.

Luma Learn delivers affordable, curriculum-aligned learning through chat-based lessons, study guides, and teaching plans powered by an AI assistant that sticks closely to official syllabi. The platform supports students with personalised help, gives teachers ready-made digital resources, and helps parents stay involved in learning at home.

Within a month of launch, Luma Learn served 125,000+ users across 11 languages. Governance is embedded through encryption, anonymisation, role-based account structures for parents and teachers, consent capture, retention limits, and compliance with national data-residency requirements. Expert review panels validate generated content before deployment, ensuring pedagogical quality alongside rapid scaling.

NavGurukul Foundation for Social Works: AI-powered adaptive assessments improve learner mastery & reduce mentor workload

Authors: Nilesh Yadav

Higher Education

Summary

NavGurukul Foundation's AI-powered adaptive assessment system addresses the limits of traditional exams that rely on fixed difficulty and manual feedback. Designed for learners in skills and finishing-school programmes, the platform evaluates MCQs, coding tasks, and open-ended responses while adjusting question difficulty in real time. By turning assessments into personalised learning experiences, the system reached over 3,000 learners, achieving 35% mastery gains & reducing mentor workload by 60%.

Context

Assessment remains a core part of learning across higher education and skilling ecosystems, yet many programmes continue to rely on static tests that measure performance without guiding improvement. In large-scale training environments, mentors often lack the capacity to provide detailed feedback to every learner, especially in cohorts that include first-generation students or those from low-income backgrounds. Learning platforms increasingly combine digital coursework with remote mentorship, but feedback loops remain slow and inconsistent. As skilling programmes expand in the Global South, there is growing demand for evaluation models that maintain quality while supporting diverse learners using mobile-first, low-cost infrastructure.

Challenge

Traditional assessments typically rely on binary scoring, fixed question difficulty, and mentor-dependent evaluation. While this approach may measure correctness, it rarely captures reasoning quality or conceptual understanding. As learner cohorts scale, mentors struggle to provide timely and personalised feedback, leading to disengagement and shallow learning outcomes.

These challenges are particularly acute for first-generation learners, who may require more structured guidance and reinforcement. When assessments are either too easy or too difficult, learners lose motivation, resulting in higher dropout rates. Institutions also face operational constraints, as manual grading and mentoring require significant time and resources. Without adaptive feedback, learners often progress with unresolved knowledge gaps, reducing long-term mastery.



Learner completing a digital assessment activity via NavGurukul's platform. (Source: NavGurukul Foundation website)



Case Study 22

The sector therefore faces a need for assessment models that can provide personalised support at scale while maintaining fairness, transparency, and efficiency.

Solution & Impact

The solution introduces an AI-powered adaptive evaluation system that transforms assessments into continuous learning experiences rather than one-time measurements. The platform evaluates multiple formats, including MCQs, coding challenges, and open-ended responses, using specialised AI evaluators that assess not only correctness but also reasoning, clarity, and conceptual understanding. AI acts as a decision support layer, providing instant feedback while mentors retain oversight of learner progress.

Assessment difficulty dynamically adapts to maintain an optimal success range, ensuring learners remain appropriately challenged. Following an adaptive engagement mechanism, around 40% of each assessment is generated based on prior performance, enabling mastery-based progression where learners advance only after demonstrating competency. Learners receive clear explanations of mistakes and targeted practice recommendations, strengthening engagement and reducing knowledge gaps.

The solution introduces an AI-powered adaptive evaluation system that transforms assessments into continuous learning experiences rather than one-time measurements.

Deployed through the Zuvy Finishing School, the system has supported more than 3,000 learners, achieving an 85% adoption rate, a 35% improvement in mastery outcomes, significant reductions in dropout rates, and approximately 60% savings in mentor time. Shared analytics dashboards provide visual mastery maps and cohort-level insights, helping mentors and programme managers identify learning gaps quickly.

Designed for scale across Global South contexts, the platform features mobile-first access, low-cost infrastructure requirements, offline readiness, and planned multilingual support. While AI enables rapid evaluation and personalised feedback, human oversight remains central to ensuring fairness, contextual judgement, and responsible use. The model demonstrates how adaptive assessment systems can improve both learning quality and operational efficiency without requiring major changes to existing training workflows.

NetPractice: Using AI for Personalized Practice and Adaptive Assessment

Authors: Harshit Agrawal, Radhika Agrawal

Higher-education

Summary

NetPractice uses AI to personalise competitive-exam preparation by analysing practice questions, past papers, and timed tests to create daily study plans that target weaknesses and reinforce strengths. It tracks accuracy, speed, and memory to time revisions, adds gamified features for motivation, and gives teachers cohort insights - turning mass coaching into efficient, data-driven, individual learning at scale.

Context

India's competitive-exam ecosystem is characterised by high pressure, dense curricula, inconsistent practice and uneven access to quality coaching. Traditional models rely on fixed timetables and one-size-fits-all content, leaving many students underprepared and overwhelmed. Educators also lack granular insight into individual learning gaps across large cohorts. NetPractice tackles this through an AI-powered mobile platform that diagnoses proficiency at onboarding and continuously adapts micro-quizzes, mocks, and revision plans using accuracy, speed, and error patterns. Its engine combines competency-mapped question banks, adaptive difficulty, and spaced-repetition scheduling, supported by daily personalised feeds, analytics dashboards, peer challenges, and educator portals. Serving over 200,000 users in Tier-2 and Tier-3 cities, NetPractice shows how AI can democratise exam preparation by turning mass coaching into precise, data-driven learning journeys.

Challenge

The key challenge was to convert mass exam preparation into individualised learning journeys without increasing costs or requiring intensive human intervention, while maintaining fairness and data protection at scale.

Students preparing for competitive exams in India face enormous pressure. Syllabi are vast, competition is intense, and many learners struggle to keep up with daily practice. Knowledge fades quickly without regular revision, yet most coaching centres still rely on fixed schedules and standard question banks that treat all students the same. As a result, a large share of aspirants remain poorly prepared for actual exam patterns and have trouble managing time and retaining concepts. Teachers and coaching institutes also face limitations. With large batches of students, it is difficult to track individual weaknesses or see where many learners are going wrong. Without detailed data on performance, educators cannot easily design targeted interventions or focus class time on the topics students need most.

These problems are sharper in Tier-2 and Tier-3 cities, where access to elite coaching and personalised guidance is limited. Students from these areas often depend on generic materials that do not adapt to their learning pace or gaps. The combination of inconsistent practice, lack of personalised feedback, and weak diagnostic insight creates an uneven playing field, making it hard for motivated students to compete fairly in high-stakes national exams.



Case Study 23

Solution & Impact

NetPractice uses machine-learning models to track accuracy, speed, and retention, dynamically generating personalised daily practice sessions and revision schedules. Educators receive analytics on cohort-level weaknesses, enabling targeted support.

Serving over 200,000 learners, the platform demonstrates improved engagement and study efficiency. A five-pillar AI governance framework ensures anonymisation, audit trails, fairness checks, and compliance with India's data-protection laws. NetPractice shows how adaptive AI can improve outcomes in high-stakes, large-scale learning environments.

NetPractice uses AI to personalise competitive-exam preparation by analysing practice questions, past papers, and timed tests to create daily study plans that target weaknesses and reinforce strengths.

Paperpal: AI Support Improves Research Writing Quality at Scale

Authors: Nishchay Shah, Nolita Coelho, Riddhi Shah, Shardool Nair

Higher Education

Summary

Paperpal addresses persistent academic writing barriers faced by Indian researchers, especially non-native English speakers. Using AI tools designed specifically for scholarly writing, Paperpal helps researchers improve clarity, structure, and efficiency across the research lifecycle. With over 700,000 users in India and more than 9.7 billion words corrected globally, users report faster revisions, fewer reviewer comments, and shorter publication timelines.

Context

India is one of the world's largest producers of academic research by volume, yet citation impact and global visibility remain below the global average. Many researchers operate with limited institutional writing support, heavy teaching or clinical workloads, and linguistic barriers that affect manuscript quality and acceptance rates. Academic writing support has traditionally relied on manual editing, peer feedback, or costly professional services, which are often inaccessible or slow. At the same time, AI adoption in higher education is accelerating, creating opportunities to support researchers at scale, provided tools are aligned with academic standards, ethics, and real-world workflows. Paperpal operates within this ecosystem as a research-focused writing assistant designed specifically for academic use rather than general-purpose writing.

Challenge

High-quality academic writing is essential for publication, impact, and career progression, yet many researchers, particularly in developing countries, struggle with language clarity, structure, and time-intensive revisions. These challenges contribute to high rejection rates, repeated review cycles, and delayed dissemination of research.

Manual editing is costly and does not scale, while generic AI writing tools often lack academic context, raise concerns about misuse, and provide unreliable citations. Researchers and institutions need tools that can support writing and revision without compromising academic integrity, while fitting into existing workflows within Word, Google Docs, or Overleaf.

The challenge is to reduce language and writing barriers at scale, improve efficiency, and support responsible AI use, especially for early-career researchers and graduate students, without replacing human judgement or peer review.



Case Study 24

Solution & Impact

Paperpal is the flagship product of CACTUS Communications, a technology company that specialises in expert services and AI-driven products to improve how research gets funded, published, communicated, and discovered. Paperpal provides a guided, AI-powered academic writing toolkit built specifically for researchers. The platform supports grammar and language refinement, content structuring, idea development, faster reading, and manuscript readiness checks. Its AI models are fine-tuned for academic writing and grounded in over 23 years of publication expertise from CACTUS.

Paperpal is available across web browsers and as extensions on Microsoft Word, Google Docs, Overleaf, and Chrome, allowing researchers to work within familiar tools. The system emphasizes guidance rather than automation, with robust data security and safeguards to reduce hallucinations, citation errors, and misuse. Human-in-the-loop quality checks and academic-only reference retrieval help maintain trust and integrity.

Since its launch in 2022, Paperpal has grown to over 4 million global users, with adoption in India increasing from approximately 43,000 users in 2023 to more than 700,000 in 2025. Academic language correction alone has processed more than 9.7 billion words, and over 1 million documents have been checked for journal submission readiness.

Paperpal helps researchers improve clarity, structure, and efficiency across the research lifecycle.

Survey evidence shows strong impact: users report improved writing clarity (89%), increased confidence (86%), and faster revisions, saving 4-10 hours per manuscript. Over half report fewer language-related reviewer comments, and many experience shorter publication timelines.

Paperpal is built to scale through cloud-based delivery, multilingual support across 50+ languages, free and tiered pricing, and purchasing power parity-based plans. Its focus on responsible AI use, combined with accessibility and academic alignment, makes it a scalable solution for strengthening research quality across India and other emerging research ecosystems.

Pratham: AI-Powered Assessment Tools for Foundational Learning

Author Name: Siddhesh Mhatre, Neel Pathak, Suchitra Bapat, Rahee

K-12

Summary

PadhAI is an offline ASR-powered reading assessment tool using Hindi and Marathi speech recognition to measure fluency, errors, and reading levels in real time. Tested with 5,000+ learners, it delivers 95% accuracy and follows a privacy-first design.

Anytime Testing Machine (ATM) is an AI-powered assessment tool that generates curriculum-aligned questions, grades responses using pedagogical rubrics, and provides feedback in English and Hindi. Tested with 4,000 learners, it improves insight into learning gaps.

Context

Foundational literacy and regular assessment are critical for improving learning outcomes, yet many education systems struggle to measure student progress consistently in low-resource and low-connectivity environments.

Assessments are often infrequent, time-consuming, and difficult to administer in overcrowded or multi-grade classrooms. As a result, teachers lack timely diagnostic insights to guide instruction and respond to diverse learning needs, highlighting the need for scalable and reliable assessment tools that function within real classroom constraints.

Challenge

Traditional reading assessments are difficult to scale because they require trained evaluators, significant classroom time, and manual feedback processes. In large classrooms with varied learning levels, teachers struggle to monitor progress regularly or provide targeted support. Existing assessments rarely capture detailed indicators such as fluency, pronunciation, and comprehension, limiting their usefulness for guiding instruction.

Infrastructure constraints further complicate implementation. Many digital assessment systems depend on stable internet access and advanced devices, making them unsuitable for rural or underserved schools. This creates a need for offline, accessible tools that provide accurate diagnostics while reducing teacher workload.



Learners engaging with the PadhAI reading assessment application. (Source: Pratham website)



Case Study 25

Solution & Impact

Two AI-enabled tools — PadhAI and ATM — were developed to support foundational literacy and formative assessment in multilingual, low-resource classrooms.

PadhAI evaluates children's oral reading using automatic speech recognition. Learners read aloud while the application detects reading errors and calculates speed and accuracy, generating visual progress reports for educators. Designed as an offline-first mobile solution, it enables quick baseline and progress assessments even in low-connectivity environments. The system has reached more than 5,000 learners and 42 facilitators across India, demonstrating over 95% accuracy while supporting data-minimisation practices to protect student privacy.

ATM complements reading assessment by supporting question generation, answer digitisation, grading, and feedback. It produces curriculum-aligned questions mapped to Bloom's Taxonomy, evaluates handwritten responses using educator-validated rubrics, and generates multilingual feedback. Pilot deployments with 4,000 learners across several Indian states have reduced grading time and improved teacher insight into learning gaps.

PadhAI is an AI-powered reading assessment application that evaluates children's oral reading skills using automatic speech recognition.

Together, the tools enable faster diagnostics, personalised feedback, and more inclusive assessment practices while preserving teacher agency in low-resource classrooms.

QwiXGenie: AI-Powered Personalised Tutoring for Coding and Technical Skills

Authors: Dhadi Sai Praneeth Reddy

Higher-education

Summary

QwiXGenie is an AI-powered personalised-learning platform from QwikZen Group that delivers on-demand coding and technical tutoring to students in India's Tier-2 and Tier-3 cities. Through web and mobile chat interfaces, it provides step-by-step explanations, multi-path problem solving, debugging support, and adaptive revision plans, complementing classrooms and coaching centres while expanding access to quality technical education with strong privacy and safety safeguards.

Context

Engineering and IT career's aspirants in non-metro regions often face limited access to specialist teachers, peer communities, and high-quality coaching institutes. Even where online courses exist, learners struggle to resolve doubts in real time or receive personalised feedback, leading to disengagement and uneven learning outcomes. Simlantenously, widespread smartphone adoption and falling data costs have created demand for continuous digital support systems that operate outside fixed class schedules. Education providers and families increasingly seek AI-driven tutoring solutions capable of scaling to hundreds of thousands of learners while remaining affordable. QwiXGenie was developed in response to this combination of unmet demand and growing digital infrastructure across India's smaller cities.

Challenge

The QwiXGenie team needed to design a system capable of delivering reliable technical explanations across multiple programming languages while adapting to wide differences in learner backgrounds and skill levels. At the same time, the platform had to discourage rote copying or academic dishonesty, protect sensitive learner data at scale, and comply with evolving regulations on student-data handling. Balancing rapid growth with institutional trust and regulatory compliance was therefore central to QwiXGenie's long-term viability.

Solution & Impact

QwiXGenie's platform integrates LLM-based tutors fine-tuned on curated programming corpora and instructional datasets. The system decomposes learner queries into conceptual gaps, generates multiple solution pathways, and surfaces debugging strategies rather than only final answers. Adaptive engines track performance across topics and adjust future practice recommendations.

Within months of launch, QwiXGenie enrolled 600,000+ learners, with usage data showing high daily engagement and repeat sessions for doubt resolution. Partnerships with training institutes and schools expanded reach beyond individual subscriptions.



Case Study 26

Governance mechanisms include minimal-data-collection policies, encryption at rest and in transit, parental dashboards for minors, content-moderation layers to prevent unsafe outputs, and audit logs for institutional partners. Compliance teams monitor evolving Indian data-protection regulations to ensure continued alignment. QwiXGenie demonstrates how consumer-scale AI tutoring platforms can grow rapidly while maintaining safety and accountability as core design principles.

QwiXGenie is an AI-powered personalised-learning platform that delivers on-demand coding and technical tutoring to students in India's Tier-2 and Tier-3 cities.

Rabbitt AI: AI admission chatbot reduces helpdesk workload during university admissions

Authors: Harneet Singh

Higher Education

Summary

During peak admission cycles, universities face overwhelming volumes of student queries, slowing response times and increasing confusion. Guru Gobind Singh Indraprastha University (IP University), Delhi, deployed an AI-powered admission chatbot that provides 24/7 multilingual guidance using verified institutional data. The system supported 78,000+ users, answered over 2.35 lakh queries, and cut repetitive helpdesk workload by over half, improving access to accurate information.

Context

Large public universities in India manage admissions for tens of thousands of applicants each year, generating continuous queries about eligibility, counselling schedules, documentation, fees, and seat availability. Most institutions rely on human-operated helpdesks, which struggle to respond consistently during peak periods. Delays and conflicting information can increase anxiety among students and parents, particularly for first-generation learners and applicants from Tier-2 and Tier-3 regions who rely heavily on official guidance. IP University faced similar challenges, with administrative teams handling high volumes of repetitive questions while also managing complex admission processes and policy updates.

Challenge

Admission cycles are high-stress periods for students, parents, and universities alike. Applicants often need timely clarification on eligibility criteria, counselling rounds, documentation requirements, and fee timelines. However, traditional helpdesks rely on manual responses, which can be slow and difficult to scale when thousands of queries arrive simultaneously.

Repetitive questions consume a large share of administrative time, limiting staff capacity to handle complex cases that require personalised attention. Information gaps or delays may lead students to miss deadlines or rely on unofficial sources, increasing confusion and inequity. These challenges are particularly significant for learners from non-metro regions or those navigating higher education admissions for the first time.

Universities also face operational constraints. Frequent policy updates, changing cut-offs, and multiple counselling rounds require constant communication, yet updating responses across channels can be difficult. Without automated support, maintaining consistency and accuracy becomes challenging, especially when information must remain strictly aligned with official university sources.

The need is for a scalable solution that can provide accurate, real-time information while reducing administrative workload and improving access to reliable guidance during admission cycles.



Case Study 27

Solution & Impact

Rabbitt AI developed and deployed an AI-powered Admission Support Chatbot for IP University using a Retrieval-Augmented Generation (RAG) framework, which grounds responses exclusively in official institutional data. The system continuously indexes information from the university website, including admission brochures, fee structures, counselling notices, cut-offs, and FAQs, ensuring that responses reflect the latest published updates.

The chatbot provides 24/7 multilingual support in English and Hindi through a mobile-responsive interface integrated directly into the university website. Designed to minimise operational complexity, deployment required limited IT integration and allowed the university to maintain control over policies and content. Safety controls ensure that responses remain citation-backed and avoid speculative or unverified guidance.

Between June and December 2025, the system supported approximately 78,000 unique users and resolved over 2.35 lakh admission-related queries. Around 41% of interactions occurred outside office hours, highlighting the importance of always-available guidance. The chatbot reduced repetitive human helpdesk workload by 55-65%, enabling staff to focus on complex student cases. Average response time was approximately 1.8 seconds, and user feedback showed a 92% “helpful” rating.

Rabbitt AI developed an AI-powered admission chatbot that provides 24/7 multilingual guidance using verified institutional data.

Usage data indicates strong equity impact, with 28% of interactions coming from Tier-2 and Tier-3 regions. The system also handled traffic spikes of up to 6-9 times normal volume during counselling announcements and result days without disruption. By improving access to verified information while reducing administrative burden, the chatbot demonstrates how retrieval-based AI systems can strengthen transparency, student experience, and operational efficiency across large higher education institutions.

Rocket Learning: Empowering Early Learners & Anganwadi Workers through AI-driven Feedback and Guidance

Author: Kshitij Jain, Shivansh Nagi

Pre-K and early years

Summary

Rocket Learning addresses persistent gaps in India's early childhood education system, particularly in low-resource settings. Worksheet-based home learning is delivered at scale, but timely assessment is difficult due to manual processes. By deploying an AI-powered system that automatically reviews worksheet images submitted via WhatsApp and returns instant visual feedback, Rocket Learning enables timely assessment and improved learning engagement. At the same time, Anganwadi Workers (AWWs) manage teaching, nutrition, health, and administrative duties with limited real-time support, making consistent guidance challenging. To address this, Rocket Learning provides Shiksha Saathi, a voice-first AI coach delivering pedagogical guidance and administrative support, reaching 7,238 workers with sustained adoption and an NPS of +43.6

Context

Early childhood education for children aged three to six is delivered at scale through Anganwadi centres, which also coordinate nutrition, health, and parent engagement. Worksheets are widely used for structured learning, but timely feedback is difficult due to large class sizes, low digital literacy, and fragmented support for Anganwadi Workers (AWWs). Teachers and AWWs face inconsistent pedagogical guidance and administrative hurdles, while children often receive delayed feedback on learning tasks.

Access to low-cost mobile phones and WhatsApp offers a channel for scalable solutions. There is a critical need to strengthen last-mile teacher support without requiring new infrastructure or high-end devices.

Challenge

Timely and meaningful feedback is essential for early learning, yet manual worksheet assessment at scale is infeasible for Anganwadi Workers. Mistakes often go unaddressed, slowing learning and reducing motivation. Existing systems rely on text-heavy instructions that increase cognitive load for workers with low digital literacy and fail to provide structured guidance or real-time reporting. AWWs manage teaching, nutrition, health, documentation, and community engagement simultaneously,



An Anganwadi worker supports young learners using Rocket Learning's platform.

(Source: Rocket Learning website)



Case Study 28

leaving limited bandwidth for personalised support. Fragmented guidance and weak parent communication further limit program effectiveness, disproportionately affecting low-income, first-generation learners. The core challenge is to deliver instant, actionable feedback for children while providing on-demand pedagogical and administrative support for educators within low-resource digital constraints. Any solution must work on low-end devices, support intermittent connectivity, minimise cognitive load, and fit existing workflows without adding overhead.

Solution & Impact

Rocket Learning addresses these challenges through two complementary AI solutions supporting both learners and educators. For students, it offers an AI-powered worksheet auto-correction system embedded within its existing WhatsApp-based learning model. The solution provides instant annotated visual feedback on worksheet submissions and functions reliably across diverse geographies, languages, and low-bandwidth environments. During worksheet creation, the content team defines Areas of Interest (AOIs) around correct answer regions and maps them to predefined responses. Using computer vision, OCR (optical character recognition), and neural network-based techniques, the system analyses worksheet images to detect markings such as ticks or circles and determine whether they fall within correct regions. It then generates an annotated feedback image highlighting responses and returns it within seconds. Currently deployed across 11 states and 180 districts, the solution reaches about 4 million children and supports 300,000 Anganwadi workers. With multilingual support and primarily visual feedback, it reduces reliance on text literacy. Following deployment, Rocket Learning observed a 4% increase in daily active users and a 5% rise in average messages per user, alongside qualitative feedback showing greater child confidence and higher parental engagement.

Rocket Learning provides instant, annotated visual feedback on children's worksheets to boost engagement and learning, & delivers voice-first AI coaching that supports Anganwadi Workers with pedagogy & administrative tasks.

For Anganwadi Workers, Rocket Learning provides Shiksha Saathi, a voice-first, multilingual AI coach via WhatsApp that delivers on-demand guidance for pedagogy, classroom management, and administrative tasks. Using intent classification and a generative AI layer grounded in ECCE frameworks, AWWs receive structured troubleshooting, lesson plans, activity suggestions, and support for documenting observations. Voice-first interactions reduce reading load and support low-literacy users, while structured reporting improves supervisor visibility and parent communication. Across 18 districts over nine months, 7,238 AWWs engaged with the system, with 35% asking multiple questions per session. NPS survey results averaged +43.6, with 69% promoters, demonstrating improved confidence, consistency, and administrative efficiency.

Both solutions operate on low-end devices with intermittent connectivity, require no app installation, and rely on WhatsApp as a familiar platform. Human review, curriculum alignment, and safe interaction design ensure reliability and ethical use. By integrating AI for both learners and educators, Rocket Learning strengthens early learning outcomes, improves teacher support, and creates a scalable, inclusive model suited to low-resource contexts.

Sakhee: AI-Powered Teacher Coaching Strengthens Classroom Instruction in Government Schools

Authors: Abhimanyu Maheshwari, Karishma Shanghvi

K-12

Summary

Sakhee addresses gaps in teacher coaching within India's public education system, where classroom feedback is infrequent and limited in quality. Using AI to provide non-evaluative, personalised instructional feedback, Sakhee helps teachers reflect on and improve classroom practices. In an eight-week pilot with 50 government teachers, 74% showed improvement in instructional practices, and 87% found the feedback accurate and helpful.

Context

Strong global evidence shows that structured pedagogy programmes improve foundational literacy and numeracy when paired with regular, supportive teacher coaching. In India, however, large class sizes, limited mentor capacity, and administrative burdens make sustained instructional coaching difficult.

Classroom observations are often infrequent and perceived as evaluative rather than supportive, reducing their impact on teacher growth. Coaches and academic mentors have limited time to provide detailed, practice-based feedback to every teacher. Within the NIPUN Bharat Mission, there is a growing need for scalable tools that can support continuous teacher development while fitting into low-resource, government school contexts. Sakhee was developed to operate within this environment, complementing, not replacing, existing coaching systems.

Challenge

Improving classroom instruction at scale requires consistent, reflective feedback aligned with specific pedagogical practices. Yet most government school teachers receive limited qualitative feedback on their daily teaching, especially for foundational literacy in early grades.

Manual classroom observations are time-intensive and difficult to scale, and feedback is often delayed, generic, or focused on compliance rather than instructional quality. Teachers may also feel anxious about observation-based evaluations, reducing openness to feedback.

The challenge is to provide timely, actionable, and non-judgmental coaching that helps teachers reflect on their practice, without increasing workload for mentors, requiring frequent in-person visits, or introducing high-tech infrastructure that is unsuitable for low-resource settings.

Solution & Impact

Sakhee is an AI-powered Teacher Coach designed to provide on-demand, non-evaluative instructional feedback to government school teachers. Teachers audio-record their classroom lessons and receive AI-guided feedback aligned to a clear pedagogical rubric for teaching Hindi literacy in Grades 1 and 2. The AI processes classroom recordings using speech-to-text and multimodal analysis to identify which instructional practices occurred and which were missed. It then



Case Study 29

generates evidence-backed, descriptive feedback focused on prioritised teaching practices. The feedback is intended to support self-reflection and improvement, not evaluation or ranking.

Sakhee was developed through multiple proof-of-concept cycles to ensure accuracy, usability, and alignment with government school realities, including low digital fluency, limited connectivity, and low-cost smartphones.

An eight-week pilot was conducted with 50 government school teachers across three blocks in Jhansi district, with official approval from district education authorities. Teachers recorded lessons regularly, averaging one recording per week of about 14 minutes. Sixty-eight percent of teachers completed six or more recordings.

Results showed measurable instructional improvement: teachers implemented 53% of key pedagogical steps in their first recording and 74% by their final recording, with 74% of teachers showing overall improvement. Additionally, 68% of teachers reported feeling comfortable recording their lessons, 87% found the feedback accurate and helpful, and 72% said they would recommend Sakhee to other teachers.

Sakhee is an AI-powered Teacher Coach designed to provide on-demand, non-evaluative instructional feedback to government school teachers.

Sakhee is designed with strong safeguards. It does not make automated decisions or judgments, does not process student identity data, and allows user-controlled data deletion. The tool supports Hindi and English, offline use, and minimal data consumption.

Early evidence suggests that Sakhee can strengthen instructional practice by complementing existing coaching systems with scalable, reflective, and context-sensitive AI support. Larger pilots are planned over the next two years to enable rigorous evaluation and potential integration within state-level NIPUN programmes.

SocraticAI: AI-powered personal tutoring for foundational Computer Science

Author: Aalok Thakkar, Ayush Thonge

Higher education

Summary

SocraticAI is a research-driven generative-AI tutoring system developed at Ashoka University to support introductory computer-science students without encouraging shortcut-driven learning or academic misconduct. Rather than acting as an answer generator, the system is explicitly designed to function as a scaffolded learning partner: diagnosing misconceptions, posing targeted follow-up questions, and guiding students toward correct reasoning paths through structured dialogue. Integrated into coursework under faculty supervision, SocraticAI demonstrates how large language models can be reshaped into pedagogically aligned tools when paired with institutional governance and transparent instructional design.

Context

Universities worldwide are confronting the rapid uptake of generative AI among students, particularly in technical subjects such as programming and data structures. While such tools can accelerate debugging and exploration, faculty increasingly report patterns of surface-level learning, copy-paste coding, and erosion of conceptual understanding. Traditional plagiarism detectors struggle to keep pace with evolving models, while blanket bans are difficult to enforce and risk driving usage underground. Ashoka University launched SocraticAI as part of a broader effort to explore responsible AI use in higher education: not whether students should use generative systems, but how such systems

could be structured to reinforce - rather than undermine - core learning objectives. The project was embedded inside live courses, with instructors shaping prompt strategies, oversight rules, and assessment integration.

Challenge

The SocraticAI team aimed to address several interconnected challenges. The system needed to prevent generative models from simply providing complete solution code, while actively encouraging conceptual reasoning and debugging skills. It also had to align interactions with course objectives and assessment norms, protect student privacy and academic records, and give faculty meaningful visibility into how the tool was being used. Crucially, SocraticAI had to function in real classroom environments without adding administrative burden for instructors or creating adversarial monitoring systems for students..

Solution & Impact

SocraticAI integrates diagnostic engines that classify student errors and misconception types from partial code submissions or explanations. These signals feed into prompt templates that instruct the LLM to ask Socratic questions, offer hints, or suggest intermediate checks rather than final answers. Faculty configure topic boundaries and difficulty thresholds, while logging systems track interaction patterns at cohort level.



Case Study 30

Pilot deployments across introductory CS courses showed substantial increases in student-reported conceptual confidence, reductions in purely syntactic “fix my code” queries, and more iterative problem-solving behaviour. Instructors reported higher-quality classroom discussions and fewer last-minute assignment crises.

Governance frameworks include anonymised identifiers, strict data-retention limits, opt-in participation, institutional hosting of logs, and IRB-style ethics approvals for research use. Audit dashboards allow faculty to inspect aggregate patterns without surveilling individuals. SocraticAI illustrates how generative tutoring systems can be aligned with academic norms when pedagogy and governance shape model behaviour from the outset.

SocraticAI is a research-driven generative-AI tutoring system that supports introductory computer-science students without encouraging shortcut-driven learning or academic misconduct. Rather than acting as an answer generator, the system is explicitly designed to function as a scaffolded learning partner.

Technovation: Mobilizing a Global Ecosystem to Prepare Girls for an AI Future

Author: Rebecca Anderson, Juana Lorena Lara, Dagmawit Getahun

Other

Summary

Technovation brings an integrated ecosystem approach to AI education that prioritizes adaptive capacity, agency, and resilience for girls in the Global South. It addresses the persistent underrepresentation of girls in technology and AI by shifting the focus from narrow skill training to holistic capability building by mobilizing a global "village" of mentors and educating them. This system creates a safe environment for participants to solve complex problems, fostering empowerment, agency, and teamwork required to navigate future volatility in emerging economies. Rather than positioning AI education as workforce preparation alone, Technovation builds ecosystems that cultivate agency, confidence, collaboration, entrepreneurial leadership and problem-solving through project-based learning supported by global mentorship networks.

Context

As AI reshapes economies, education systems face pressure to "train for jobs that do not yet exist." For girls in low- and middle-income countries, barriers are compounded by social norms, lack of role models, and limited access to supportive learning environments. Traditional coding bootcamps or curriculum-only interventions often fail to sustain participation or translate into long-term confidence and leadership. Drawing on examples from India, Nigeria, and Honduras, Technovation shows how safe, supportive environments enable girls to build teamwork, confidence, and entrepreneurial

mindsets alongside AI technical skills. The analysis finds that programmes embedded in strong community infrastructure develop greater agency and persistence than standalone technical courses, concluding that inclusive AI leadership development requires mobilising entire social systems rather than simple workforce pipelines.

Challenge

The challenge was to design an AI education model that could scale globally while remaining safe for girls, culturally adaptable, and capable of fostering long-term agency rather than short-term technical outputs. The programme needed to work across diverse education systems and community



Girls showcase a project developed through Technovation's programme.
(Source: Technovation website)



Case Study 31

contexts, without relying on uniform infrastructure or formal school integration. Isolated curricula and stand-alone training programmes were insufficient to sustain participation or build long-term agency, including the ability to translate technical ideas into viable business ventures, particularly in gender-unequal contexts.

Solution & Impact

Technovation created a distributed ecosystem model combining curriculum, local champions, volunteer mentors, educators, parents, and community events. Girls can work in teams to identify real-world problems and build AI-enabled solutions developing comprehensive business canvases and pitches to bring their innovations to life. This collective scaffolding, supported by structured mentorship and peer networks, reduces dropout and builds confidence to tackle complex challenges.

It addresses the persistent underrepresentation of girls in technology and AI by shifting the focus from narrow skill training to holistic capability building by mobilizing a global "village" of mentors and educating them.

The programme has reached over 300,000 learners across multiple countries. The data indicate higher improvements in persistence, leadership, and entrepreneurial problem-solving skills compared to standalone technical training. Strong governance frameworks underpin operations, including parental consent, data minimisation, anonymisation, and compliance with international child-protection standards. Technovation demonstrates that scaling AI education responsibly requires social infrastructure and real-world applications as much as technical tools.

Top Parent: AI-powered vernacular learning companion strengthens parent-led foundational literacy and numeracy at home

Authors: Pradipt Singh Tomar, Krishna Kumari

K-12

Summary

Top Parent addresses foundational literacy and numeracy gaps by helping parents support daily learning at home, particularly in low-income communities facing language and literacy barriers. The vernacular, mobile-first platform uses AI to provide personalised nudges, adaptive learning pathways, and multilingual chatbot support. The app has reached over 1.8 million downloads and around 910,000 active users, with independent studies reporting higher parental engagement, a positive association between parent participation and children's learning progress and measurable gains in foundational learning outcomes.

Context

India's early childhood and primary education ecosystem faces persistent challenges in foundational literacy and numeracy, particularly among children from low-income households. Caregivers often want to support learning but may lack confidence, time, or access to appropriate resources, especially when language barriers exist. Learning support outside classrooms frequently relies on simple tools such as WhatsApp, lightweight mobile applications, and audio-based content that can function on low-end smartphones and low-bandwidth networks. Within this environment, Top Parent operates as a free, vernacular, mobile-first platform designed for parents and young children ages 3-8, combining structured activities with audio-first guidance and culturally relevant content suited to rural and semi-urban contexts.

Challenge

Low foundational learning levels are closely linked to limited opportunities for guided practice at home. Many parents struggle to support their children consistently due to low literacy levels, language differences, or uncertainty about what to do each day. Traditional interventions, such as printed worksheets or one-way messaging, often assume a level of reading ability or digital familiarity that may not exist in many households. Schools and educators have limited capacity to personalise guidance for every family, and most EdTech products focus directly on students rather than caregivers.



A caregiver uses Top Parent's mobile platform to support children's foundational learning.
(Source: NudgED website)



Case Study 32

Sustaining engagement over time is another significant barrier. Parents balancing work and household responsibilities may find it difficult to maintain daily routines without reminders or personalised encouragement. Content production in multiple languages also creates operational challenges for organisations working at scale. Without systems that simplify communication, adapt to individual usage patterns, and reduce the burden of creating and translating learning materials, scaling parent-led learning models remains difficult across diverse geographies and contexts.

Solution & Impact

Top Parent addresses these challenges through an AI-enabled learning companion designed to fit into existing mobile habits and communication channels. The platform combines NCERT-aligned literacy and numeracy activities with personalised nudges that guide parents through short daily learning routines. A multilingual chatbot, powered by a Retrieval-Augmented Generation (RAG) framework with a configurable base model, responds to common parent questions using curated knowledge sources. Integration with Bhashini and Jugalbandi enables voice-to-text and audio responses, making the platform accessible to caregivers who may not read or type comfortably. AI is also used internally to support faster development of multilingual learning content, helping draft worksheets, translations, and voice-overs while maintaining human review and approval.

The solution is currently deployed across multiple Indian states through direct downloads and partnerships with schools and community organisations. The app has surpassed 1.8 million downloads, with approximately 910,000 active users and over 100,000 parents using the AI chatbot. Independent studies cited in programme reports indicate improvements in foundational

Top Parent operates as a free, vernacular, mobile-first platform designed for parents and young children ages 3-8, combining structured activities with audio-first guidance and culturally relevant content suited to rural and semi-urban contexts.

learning outcomes, including an average 0.35 standard deviation gain in literacy and numeracy indicators. Parental engagement behaviours also improved, with the proportion of parents reporting reading with their children increasing from 74% to 87%. Early evidence suggests that personalised WhatsApp nudges are particularly effective in sustaining participation.

Designed for low-bandwidth environments and shared devices, the platform prioritises vernacular access, audio-first onboarding, and human oversight of AI outputs. Scaling the model across geographies may require adaptation to local dialects and continued partnerships with trusted community channels, alongside safeguards such as curated knowledge sources, validation workflows, and clear boundaries around the guidance provided to caregivers.

Torchit Electronics: AI-Powered Accessible Reading for Inclusive Education of Print-Disabled Learners

Author: Hunny Bhagchandani

K-12

Summary

Torchit Electronics develops AI-powered assistive-reading ecosystems that enable students with visual impairments and print disabilities to independently access textbooks, worksheets, exam papers, and classroom materials. By integrating optical character recognition (OCR), multilingual natural-language processing, summarisation, and text-to-speech systems, Torchit converts printed educational content into accessible digital and audio formats in real time. Torchit demonstrates how assistive AI can be operationalised at population scale while embedding privacy-first design and institutional data control.

Context

Globally, hundreds of millions of learners live with visual impairments or print disabilities, yet accessible educational materials remain scarce. In many developing contexts, fewer than one percent of textbooks are available in Braille or audio form. Producing accessible formats is slow and expensive; audiobooks can take weeks to prepare, while Braille printing requires specialised facilities that are rarely available in rural or under-resourced districts.

As classrooms increasingly rely on printed handouts, blackboard notes, and exam sheets, students who cannot read standard print often depend on human readers - peers, volunteers, or teachers - reducing autonomy and limiting classroom participation. These barriers are

magnified in multilingual environments like India, where content must be accessible across dozens of languages.

Torchit emerged in this landscape to provide point-of-use accessibility: allowing learners to convert any printed educational material into an accessible format at the moment it is needed, rather than waiting for specialised production pipelines.

Challenge

Students with visual impairments and print disabilities often face major barriers in education. Textbooks are rarely available in accessible formats, Braille or audio versions arrive late, and many learners must depend on classmates or volunteers to read materials aloud. These difficulties are even greater in low-resource countries, where schools lack trained support staff, assistive devices, and inclusive infrastructure. As a result, many students are excluded from everyday classroom activities, struggle to keep up with lessons, and face higher risks of dropping out before completing their education.

Torchit was created to address these gaps through an AI-powered accessible-reading ecosystem that allows learners to convert printed study materials into usable formats in real time. The system combines optical character recognition to read printed text, multilingual language processing to handle different scripts, and text-to-speech tools that read content aloud.



Case Study 33

Additional AI features summarise long passages and translate materials into other languages, giving students more control over how they learn.

Using Torchit Innovations like Jyoti AI, students can independently access textbooks, worksheets, exam papers, notices, and reference materials at school, at home, or in libraries. By placing accessibility directly at the point of learning, the solution reduces dependence on human readers and helps learners participate more fully in mainstream classrooms, strengthening confidence and continuity in their education.

Solution & Impact

Torchit's platform integrates computer-vision pipelines for OCR, multilingual language-identification systems, translation modules, summarisation engines, and neural text-to-speech models to deliver near-instant audio output from scanned pages. Learners can independently access textbooks, notices, worksheets, and handwritten notes in multiple Indian and global languages.

Rolled out across schools, colleges, training centres, and inclusive classrooms in India and other Global South regions, the solution tackles shortages of accessible textbooks and human readers. By delivering learning materials in usable formats at the point of need, it boosts participation, confidence, and continuity for print-disabled students. Designed to be affordable and scalable, it enables governments and CSR programmes to expand inclusive education, with Torchit already reaching over 1.4 million learners worldwide.

Torchit Electronics develops AI-powered assistive-reading ecosystems that enable students with visual impairments and print disabilities to independently access textbooks, worksheets, exam papers, and classroom materials.

Governance is built around a privacy-first architecture: only user-initiated materials are processed; no biometric data is retained without explicit consent; institutions retain ownership and control of documents; and encrypted transmission protects sensitive content. Torchit regularly audits model performance across languages to mitigate bias and ensure equitable accessibility. The case illustrates how assistive AI can be scaled responsibly within public-education systems.

TOY8: AI-Based Developmental Screening Expands Inclusive Early Education in ASEAN Countries

Authors: Masaki Ishibashi, Bruno Abrioux Takano, Chui Yee Chong

Pre-K and early years

Summary

TOY8 addresses the lack of early developmental screening for children aged 3-5 in the Association of Southeast Asian Nations (ASEAN), where many learning and developmental needs go undetected. Using an AI-powered mobile screening tool and digitalised intervention plans, TOY8 enables preschool teachers to identify developmental delays early and support inclusive education. This solution delivers strong value for money, with a government-commissioned cost-benefit analysis indicating that public investment yields almost double the economic return, alongside autism detection rates 18 times higher than conventional approaches. To date, Toy8 has screened over 15,000 children across three ASEAN countries, with children participating in interventions showing average developmental gains of 8.96 months within six months.

Context

Early childhood development strongly shapes long-term learning, health, and independence. Yet across many ASEAN countries, routine developmental screening typically ends after infancy, leaving preschool-aged children without systematic assessment. As a result, learning and developmental delays are often identified only after children enter formal schooling, when intervention is less effective and more costly.

In Malaysia and neighbouring countries, pre-schools serve children from diverse socioeconomic backgrounds, often with limited access to specialists such as developmental

paediatricians or speech therapists. Traditional assessments are expensive, clinic-based, and inaccessible for rural and marginalised communities. TOY8 operates within this ecosystem by equipping preschools with simple digital tools that allow early identification and support without relying on specialist availability.

Challenge

An estimated 10-15% of children experience developmental disabilities, but most remain unidentified during the critical preschool years. In Malaysia alone, nearly one in four children entering Year One lacks basic literacy or numeracy skills.

The challenge is most acute for children aged 3-5, a period when brain development is rapid and early intervention is most effective. However, there is no routine developmental surveillance during these years in most ASEAN countries. Existing systems depend on specialist-led assessments, which are costly, time-consuming, and geographically limited.

Preschool teachers often observe learning or behavioural concerns but lack validated tools to assess them objectively. Without early identification, children miss timely support, educators struggle to adapt classrooms, and systems face higher long-term costs related to remedial education, healthcare, and reduced independence. There is a clear need for scalable, affordable, and reliable screening that can be embedded into everyday preschool settings.



Case Study 34

Solution & Impact

TOY8 delivers an “Inclusive Education 2.0” model that combines AI-powered developmental screening with digitalised intervention planning. The solution includes a mobile screening application and a digital Individualised Education Programme (IEP) platform.

Children complete short, game-based activities guided by animated characters on a smartphone or tablet. The AI system uses voice recognition, image analysis, and behaviour tracking to assess performance across five developmental areas: gross motor, fine motor, language, cognitive, and personal-social skills. Behavioral patterns, response accuracy and speed are analysed during task performance to generate objective, bias-reduced screening results.

The assessment framework builds on Japan’s nationally mandated early childhood screening system and has been adapted for ASEAN contexts through partnerships with universities and medical experts. Preschool teachers administer screenings after completing certification training, addressing specialist shortages while strengthening teacher capacity.

As of 2025, TOY8 has screened over 15,000 children across Malaysia, Indonesia, and Singapore. A total of 157 preschool teachers and 16 healthcare professionals across 78 schools and 16 healthcare facilities have been certified.

Using an AI-powered mobile screening tool and digitalised intervention plans, TOY8 enables preschool teachers to identify developmental delays early and support inclusive education.

Clinical validation by Malaysia’s Ministry of Health showed high accuracy compared to gold-standard assessments. A government-commissioned analysis found the AI approach achieved 18 times higher autism detection than conventional methods and generated RM 1.85 in benefits for every RM 1 invested, with projected savings of RM 6.87 billion. Early intervention pilots showed an average developmental gain of nearly 9 months within six months.

TOY8 is designed for scale through smartphone delivery, multilingual support, and strong data privacy safeguards. Its teacher-empowerment model and cross-sector collaboration make it a replicable approach for inclusive early education across the Global South.

Transform Schools, People For Action: AI-powered platform uses school data to flag risks early & improve system-level decision-making

Authors: Tanushree Narain Sharma, Neetu Sahu, Prerak Shah

K-12

Summary

India's education system collects large amounts of data, but often identifies learning problems too late. Transform Schools, with Catalyst Management Services, developed Data For Decisions (DFD), an AI-enabled platform that turns routine school data into early warnings and actionable insights. The system predicts students at risk of learning decline with 84% recall accuracy and currently supports decision-making for 1.4 million students across four states.

Context

India's K-12 education system spans millions of schools and students, generating extensive assessment and administrative data through government and programme reporting systems. While this data is widely collected, it is typically used for retrospective reporting rather than timely action. State and district officials review results after learning gaps have already widened, limiting their ability to respond early. Teachers and school leaders often lack clear visibility into why learning outcomes vary across schools or which students need immediate support. Transform Schools operates within this system, working closely with state governments and education departments to strengthen learning outcomes through large-scale programmes that rely on existing school data and governance structures.

Challenge

Effective education planning depends on knowing not just what outcomes look like, but why they occur and who is at risk before learning declines. In practice, most monitoring systems in India are backwards-looking. Assessment data & school indicators are reviewed weeks or months later, by which time opportunities for early intervention are lost.

Administrators face several constraints. They manage thousands of schools with limited staff and time, making it difficult to analyse complex datasets manually. Teachers and school leaders may sense that certain students or schools are struggling but lack clear evidence on underlying causes or how risks differ across subjects. Traditional dashboards often present raw scores or averages without explaining patterns, drivers, or priorities for action.

These gaps reduce system efficiency and weaken alignment between policy intent and classroom practice. Without timely, interpretable insights, decisions around school visits, teacher support, and programme design are often reactive. The challenge is to move from data collection to early, actionable intelligence, using existing data, without adding complexity or burden for already stretched education systems.



Case Study 35

Solution & Impact

Data For Decisions (DFD) addresses this challenge by embedding AI-driven decision support directly into routine education workflows. Developed by Transform Schools in partnership with Catalyst Management Services, DFD converts assessment and administrative data into early warnings and practical insights for decision-makers at state, district, and school levels. DFD has the capability to analyse multi-source education data, generate predictive risk alerts, and translate analytics into role-specific, actionable recommendations across system levels.

The platform uses predictive models trained on approximately 110,000 student records across Classes 6–9, combined with government data sources such as UDISE+. These models identify students who are likely to experience learning decline or improvement, achieving 84% recall accuracy. DFD also analyses school-level data to highlight factors linked to stronger or weaker performance, such as attendance patterns, infrastructure, or subject-specific trends.

Insights are delivered through simple, role-based dashboards used by state officials, district administrators, school leaders, and educators. Rather than replacing human judgment, the system flags risks and priorities, enabling officials to focus attention where it is most needed. The platform is fully deployed and integrated into regular review meetings and planning processes.

DFD currently supports decision-making for around 1.4 million students across four states, with predictive modelling applied to large, real-world datasets. Its insights have directly influenced system action. In Karnataka, data-driven flags on low attendance informed state-level directives that increased official school

The platform predicts students at risk of learning decline and helps officials focus attention where it is most needed.

visits to 78%. In Odisha, the identification of delayed assessment data led to district orders that improved reporting timelines. Subject-level insights have also shaped programme priorities by revealing learning gaps and cross-subject dependencies.

Designed for scale, DFD works with both online and offline data and uses lightweight dashboards suitable for low-resource settings. Strong data access controls and a human-in-the-loop approach ensure responsible use. Overall, the platform demonstrates how AI, when used as decision infrastructure, can help education systems act earlier, target support better, and improve equity at scale. DFD is being strengthened through scalable, multilingual generative AI tools that enable teachers to access and use data for improved lesson delivery, contextual content curation, and adaptive assessment in the Indian classroom context.

Udhyam Saathi (AI Mentor): Democratising Entrepreneurial Mentorship at National Scale via Multimodal Generative AI on WhatsApp

Author: Shashanka Kundu, Sani Sabale

K-12

Summary

Udhyam Saathi is a generative-AI mentoring system embedded in WhatsApp that addresses one of India's most persistent education bottlenecks: the inability of teachers to provide individual entrepreneurial guidance to millions of government-school students. Built on retrieval-augmented generation pipelines and multimodal submission workflow, it provides personalised and adaptive learning for students through round-the-clock AI mentorship, including instant help, feedback on milestone submissions, and tailored reminders. Teachers are supported by an AI copilot that offers on-demand assistance, reviews student work, and sends timely prompts, helping reduce administrative workload.

Context

India faces a major youth opportunity gap, as millions enter the workforce without the skills or confidence to build their own livelihoods. While Udhyam Learning Foundation runs entrepreneurship programmes nationwide, teachers struggle to mentor the vast number of student business projects Udhyam's Entrepreneurial Mindset Curriculum reaches millions of secondary-school students, encouraging them to develop real-world ventures supported by government grants. These projects span agriculture, retail, services, and manufacturing, creating enormous diversity in mentorship needs. Teachers, often managing classes of sixty or more learners, cannot realistically provide detailed sector-specific feedback at the pace required.

Without structured guidance, many students abandon projects or submit low-quality plans, undermining the curriculum's long-term goals of livelihood creation and economic agency. Digital mentoring systems exist, but few are embedded in tools already used by students, and even fewer operate at the scale or governance level required for government deployment.

Challenge

Udhyam faced the problem of delivering individualised mentorship at population scale while preventing inappropriate AI interactions with minors and ensuring teachers retained authority. It also incurred challenges while avoiding the collection of unnecessary personal data and aligning with India's evolving digital-governance frameworks.

Solution & Impact

Udhyam Saathi allows students to get their queries addressed and to submit milestone work via WhatsApp in text, image, or video formats. Retrieval-augmented pipelines draw on curriculum repositories, sectoral examples, and government guidelines to generate contextualised feedback. Automated nudges encourage persistence, while teachers receive dashboards and grading assistance.

The programme now serves 3.9 million students across ten Indian states and uses cloud-based systems that can handle heavy traffic during submission periods, even in remote areas. The main users are government-school students in Grades 11–12, while teachers benefit from reduced grading work and more time for mentoring.



Case Study 36

After AI was added, student idea submissions rose by 60%, with 163% increase in novelty and 126% improvement in articulation of ideas. Overall submission quality also improved by about 30%. Further, some teams voluntarily revised their work, showing greater student initiative. Teachers reported lower administrative burden. Strong safety systems review AI responses, anonymise data, and block inappropriate content. The platform works on low-bandwidth WhatsApp connections and supports several Indian languages. Students were especially engaged by video-based feedback, though encouraging voluntary help-seeking remains a challenge. Lessons so far highlight the importance of easy access, friendly AI design, and constant monitoring. Overall, Udhya Saathi shows how AI can expand mentorship at scale and promote opportunity in low-resource settings.

Udhya Saathi provides personalised and adaptive learning for students through round-the-clock AI mentorship, including instant help, feedback on milestone submissions, and tailored reminders.

Safeguards include anonymised identifiers, role-based access controls, child-safety filters, negative-testing suites for harmful prompts, and DPDP-aligned consent protocols. Udhya Saathi shows how generative AI can expand mentoring capacity inside public systems without sacrificing institutional oversight.

