

Human-Guided Agentic AI for Skill Formation in Indian Higher Education: A Framework-Centred Study for Learner Modelling, Skill-Gap Diagnosis and Responsible Evaluation

Manjunath D R¹, Preetha S^{2,*}, Anil Kumar B³, Siva Sumanth Reddy⁴, Krupa K S⁵

^{1,2,5}B.M.S. College of Engineering, Bengaluru, India

³R.V. Institute of Engineering and Management, Bengaluru, India

⁴Dayananda Sagar Academy of Engineering and Management, Bengaluru, India

Email ID: ¹manjunthdr.cse@bmsce.ac.in, ^{2,*}preetha.ise@bmsce.ac.in, ³anilkumarb.rvitm@rvei.edu.in,

⁵sumanthdsatm@gmail.com, ⁵krupaks.csi@bmsce.ac.in

HOW TO CITE:

Manjunath D R, Preetha S, Anil Kumar B, Siva Sumanth Reddy, Krupa K S (2026). Human-Guided Agentic AI for Skill Formation in Indian Higher Education: A Framework-Centred Study for Learner Modelling, Skill-Gap Diagnosis and Responsible Evaluation. International Journal of Special Education, 41(7s), 1123-1140.

COPYRIGHT STATEMENT:

Copyright: © 2026 Authors.
Open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

ABSTRACT:

Artificial intelligence in education has a longer research history than the current wave of large language models. Intelligent tutoring systems, knowledge tracing, learning analytics and educational data mining established the methodological basis for adaptive support; generative AI subsequently made natural-language feedback, question generation, coding assistance and assessment redesign widely accessible; and agentic AI now promises multi-step workflows that can plan, retrieve, monitor and adapt. The difficulty is that agentic educational AI is advancing faster as a technology than as a validated classroom practice, and controlled evidence is especially scarce in Indian higher and engineering education, where scale, multilingual learners, Outcome-Based Education (OBE) requirements and data-protection duties shape what responsible deployment can look like. This paper addresses that gap through design-science reasoning: it synthesises foundational and recent evidence only to the extent needed to derive design requirements, and then develops SAKSHAM-AI, a human-guided agentic framework built on the principle of bounded autonomy. The framework integrates learner profiling, skill-gap diagnosis, knowledge tracing, retrieval-augmented content grounding, formative assessment, misconception-level feedback, self-regulated learning support, faculty dashboards, CO/PO outcome mapping and governance monitoring into a single auditable workflow in which agents recommend, retrieve, explain and flag while teachers retain authority over high-stakes academic decisions. We specify a trial-ready evaluation protocol—a 12–16-week pilot with a delayed-retention test-reporting knowledge gain, skill-gap closure, retention, independent transfer, equity, hallucination rate, explainability, privacy minimisation, escalation and OBE coverage. The contribution is therefore a bounded-autonomy, human-guided architecture together with its measurement and governance design, rather than a deployment claim; the framework is proposed and trial-ready but requires empirical validation in live institutional settings before any effectiveness claim can be made.

Keywords: Artificial intelligence in education; intelligent tutoring systems; Agentic AI; Indian higher education; learning analytics; knowledge tracing; skill formation.

1. Introduction

The current discussion on artificial intelligence in education is often framed as if it began with ChatGPT. That view is historically narrow. The deeper research lineage begins with intelligent tutoring systems, student modelling, knowledge tracing, learning analytics and educational data mining. These traditions established a central proposition that remains relevant: educational technology improves learning only when it can represent learner state, interpret errors, provide timely feedback and support teachers in making better decisions (Corbett et al., 1994, Filamant et al., 2024, Adesope et al., 2014, Kulik et al., 2016, Roll 2016, Kuzilek et al., 2017, Choi et al., 2020, Liu et al., 2022).

Large language models shifted this adoption curve sharply. Because the interaction is conversational, students no longer need a dedicated tutoring environment to obtain an explanation, a corrected program, a summary or a practice quiz; they simply ask (Guduru et al., 2026). The convenience is genuine, but it unsettles assessment, since a fluent answer obtained on demand weakens the link between a submitted task and a student's own competence (Amaresh et al., 2025). A fluent response is not necessarily a pedagogically sound response; a completed task is not necessarily evidence of independent competence; and a personalised answer is not necessarily equitable or privacy-preserving (Kasneci et al., 2023, Yan 2024, Yusuf et al., 2024, Wang et al., 2024, Gruenhagen et al., 2024, Al Shloul et al., 2024, Strzelecki 2024, Kofinas et al., 2025).

The emerging agentic-AI literature adds a further layer. An agentic system may decompose tasks, retrieve evidence, use tools, maintain memory, reflect on prior actions and coordinate specialised agents. In education this could support personalised remediation, teacher co-design, formative assessment and institutional analytics (Ellikkal et al., 2025, Lachheb et al., 2025, Sergeeva et al., 2025, Deng et al., 2025, Guan, 2024, Chen et al., 2025, Coban et al., 2025, Looi & Jia, 2025, Alfredo et al., 2025). Yet the same autonomy raises questions about accountability, hallucination, over-delegation and hidden learner profiling (Gangothri &

Chandrappa, 2026). For that reason, this paper treats agentic AI as a promising research direction, not as a mature classroom intervention (Guo, et al., 2025, Jiang et al., 2024, Liu et al., 2025, Sudarshan et al., 2026).

India is an especially important context because the problem is not only technological. Indian higher education combines scale, regional language diversity, uneven infrastructure, engineering-employability pressure, Outcome-Based Education requirements and national digital platforms. AISHE reports enrolment on the scale of crores; SWAYAM and NPTEL provide large-scale online and engineering-oriented learning infrastructure; and DIKSHA, Bhashini, BharatGen and AIKosh indicate a public-platform route for multilingual and data-centred AI. The question is whether these infrastructures can be connected to rigorous learning evidence, not simply whether AI tools can be procured (Wu et al., 2026, Liang et al., 2026, Rashmi et al., 2026, Doo & Park, 2026, Aktar, 2021, Swayam 2026, PIB 2025, AIKosh 2025, Bhashini 2026, BharathGen Consortium 2025).

The problem this creates is easy to misread. In Indian higher and engineering education the constraint is not a scarcity of AI tools: capable models, multilingual resources and national platforms are already in place. What is missing is a disciplined way of organising these capabilities into a teaching and learning process - one that yields defensible evidence of skill formation while keeping faculty in charge of academic decisions, protecting the validity of assessment, and handling learner data responsibly.

An unconstrained conversational tutor cannot supply that evidence, because it has no stable model of the learner, no binding link to course outcomes, no faculty-approved content boundary and no audit trail. What is required instead is a framework that converts established learning-science foundations and recent generative and agentic capabilities into a bounded, inspectable workflow. This paper develops such a framework. Using design-science reasoning, it derives design requirements from the evidence base and proposes SAKSHAM-AI-Skill-Aware Knowledge Support and Holistic Adaptive Mentoring through Agentic AI-a human-guided,

bounded-autonomy architecture intended for empirical validation in Indian institutions rather than for immediate deployment.

2. Research Questions and Contribution

Table 1. Research questions and analytical focus of the proposed agentic AI framework

RQ	Focus	Question
RQ1	Design requirements	What design requirements should guide a human-guided agentic AI framework for skill formation in Indian higher and engineering education?
RQ2	Integration	How can learner profiling, skill-gap diagnosis, knowledge tracing, RAG-based content grounding, formative assessment and feedback be integrated into one bounded-autonomy workflow?
RQ3	Oversight and governance	What human-in-the-loop, privacy, explainability and governance controls are required for responsible classroom use?
RQ4	Validation	What trial protocol and evaluation metrics can validate the framework in Indian institutions?
RQ5	Outputs	What student-, faculty-, department- and governance-facing outputs should the framework generate?

Abbr.: AG=agentic AI; GenAI=generative AI; HE=higher education; RAI=responsible AI; DPI=digital public infrastructure; OBE=Outcome-Based Education.

The contribution of this paper is primarily architectural rather than bibliometric. First, it formulates an explicit set of design requirements for human-guided agentic learning support in Indian higher and engineering education, derived from learning-science foundations, generative and agentic

capabilities, and responsible-AI constraints. Second, it presents SAKSHAM-AI, a layered bounded-autonomy architecture in which ten specialised agents perform defined academic functions-learner profiling, skill-gap diagnosis, knowledge tracing, retrieval-augmented content grounding, formative assessment, feedback, self-regulated-learning support, faculty dashboards, OBE evidence mapping and governance auditing-under continuous faculty oversight. Third, it specifies the human approval points, student correction rights, source-approval gates and escalation rules that keep autonomy bounded. Fourth, rather than describe the architecture and leave it at that, we attach a measurable trial metric to every agent, so that each component can be put to an empirical test rather than merely asserted. Fifth, it positions the literature only as the evidence base that motivates these design choices, keeping the manuscript framework-centred from the methodology onward. The supporting evidence synthesis is therefore in service of the architecture, not the other way around. The study was guided by five framework-oriented research questions, as shown in Table 1. These questions organize the paper around design needs, workflow integration, governance controls, validation strategy, and expected outputs of SAKSHAM-AI.

3. Methodology

3.1 Design-science orientation

The framework is developed through design-science reasoning, which treats an architecture as a research artifact to be constructed, justified against a knowledge base, and evaluated against explicit requirements. The aim is not to summarise the literature but to translate four bodies of knowledge-learning-science foundations, generative and agentic capabilities, Indian institutional and Outcome-Based Education (OBE) requirements, and responsible-AI constraints-into a coherent, testable design. Design-science research alternates between a relevance cycle that grounds the artifact in a real institutional problem and a rigour cycle that grounds it in established theory and evidence. Following this logic, the manuscript first fixes the problem (verifiable skill formation under faculty authority and data-protection duties), then derives

design requirements, then specifies the SAKSHAM-AI artifact, and finally defines how the artifact would be evaluated; this workflow is summarized in Figure 1. No live deployment is reported; the

methodological commitment is to make every design choice traceable to a learning or governance need and to a measurable trial outcome.




Figure 1. Design-science workflow

3.2 Design inputs

Design inputs were drawn from a focused, traceable evidence base rather than an exhaustive survey. Five strands informed the requirements: (i) intelligent tutoring systems, knowledge tracing, learning analytics and educational data mining, which supply the logic of learner modelling and mastery estimation; (ii) generative AI and retrieval-augmented generation, which supply the natural-language interaction and source-grounding layer; (iii) agentic and multi-agent orchestration, which supply task decomposition, tool use and coordination; (iv) Indian higher-education requirements, including scale, multilingual learners, OBE and accreditation evidence, and faculty workload; and (v) responsible-AI and Digital Personal Data Protection (DPDP) requirements, which make consent, minimisation, explainability and auditability first-class constraints. The development decisions that follow from these inputs are summarised below, and the evidence strands that grounded them are shown in the accompanying mapping. Corpus-level detail is deliberately confined to the short evidence-base section so that it supports, rather than supplants, the framework. To ensure methodological traceability, the framework development process was organized into a sequence of design-science activities. To keep the evidence base both grounded and current, the search was organised along five strands: foundational work on learner modelling and tutoring systems; recent generative AI; agentic and multi-agent systems; India-specific platforms and policy; and governance and responsible-AI studies.

3.3 Principles of Framework Design

Eleven design principles constrain the artifact, and each is later enforced by a particular agent or approval point. The first, and the one the rest depend on, is bounded autonomy: an agent may recommend, retrieve, explain, draft, summarise, flag or escalate, but it may do nothing beyond this. Authority over grading, certification and accreditation stays with faculty. Students, in turn, can see the model the system holds about them and can ask for it to be corrected; and any weakness the system claims to have found must trace back to an observable item, rubric or task rather than to an opaque score. Source-grounded content requires recommendations to be drawn from a faculty-approved corpus rather than open generation. Formative-first assessment treats AI-generated items as practice unless faculty explicitly approve summative use. Explainable recommendations require each suggestion to carry a reason a student or teacher can inspect. Faculty oversight routes class-level inferences through the teacher before action. OBE evidence mapping links learner artifacts to course and programme outcomes. Privacy minimisation collects only the fields a function needs. Auditability records data use, AI actions, approvals and escalations so that the workflow can be reviewed after the fact.

3.4 Agent design procedure

Each agent was derived from a single learning or governance need rather than from a technical capability looking for a use. Starting from the design requirements, every requirement was mapped to the smallest function that could satisfy it and to the human checkpoint that should govern it. The need to represent the learner produced the learner-profile

and knowledge-tracing agents; the need to make weaknesses actionable produced the skill-gap agent; the need to keep content trustworthy produced the retrieval-augmented content agent and its source-approval gate; the need to practise without distorting grades produced the assessment and feedback agents under a formative-first rule; the need to develop independent learners produced the self-regulated-learning agent; the need to make the system useful to teachers and to accreditation produced the faculty-dashboard and OBE agents; and the need to remain lawful and inspectable produced the governance agent. This one-need-one-agent discipline keeps the architecture legible and ensures that no agent acquires authority that the design did not intend.

3.5 Evaluation design

Because the central claim is architectural, evaluation is defined at the level of agents and outcomes rather than of overall satisfaction. Ten indicators are specified and each is bound to the agent whose value it tests: knowledge gain (KG) and skill-gap closure (SGC) for the diagnostic and remediation agents; retention and independent transfer for durability of learning; equity gap (EqGap) for fairness across language and access groups; hallucination rate (HallR) and explanation usefulness (XAI-U) for the content and feedback agents; privacy-minimisation ratio (PrivR) and escalation rate (EscR) for the governance agent; and OBE coverage (OBE-Cov) for the outcome-mapping agent. Defining the metrics together with the agents means the framework can be falsified: if an agent does not move its indicator in a trial, that component—not the idea of agentic AI in general—is what needs revision.

3.6 Validation boundary

The boundary of the present contribution is stated plainly. SAKSHAM-AI is a proposed, trial-ready architecture; it has not been validated through a live institutional classroom trial, and no effectiveness result is claimed. The evidence base motivates the design and the evaluation model specifies how the design would be tested, but the step from a justified artifact to a validated intervention requires the pilot described in Section 9. All later references to outputs and effects are therefore expressed as

expected outputs, possible trial evidence and future validation, not as demonstrated results.

4. Evidence Base Supporting the Framework

The framework is grounded in a focused evidence base rather than an exhaustive review; corpus detail is kept brief here so that it supports the design without becoming the subject of the paper. A traceable set of records was profiled by period, evidence type and India relevance, with the temporal spread, evidence-type distribution, dominant themes, outcome coverage and India-specific signal shown in Figs 2–5. Three observations follow that directly justify the design. First, the foundational intelligent-tutoring and knowledge-tracing literature provides the strongest basis for measurable learning claims, which is why learner modelling is placed at the core of the architecture. Second, recent generative-AI evidence is useful but heterogeneous, which is why the framework treats the model as a grounded, formative-first interaction layer rather than an autonomous tutor. Third, India's policy and platform readiness is strong while classroom-level learning evidence is weak, which is why the framework is offered as a trial-ready design and not as a deployment claim. These figures therefore motivate the requirements; the remaining sections develop the architecture itself.

5. Framework Design Rationale

5.1 Learner modelling rationale

The framework places learner modelling at its Centre because the most credible evidence for measurable learning comes from intelligent tutoring systems, knowledge tracing, learning analytics and educational data mining, not from usage or perception data. This is the design reason for two agents. A learner-profile agent is needed because adaptive support is only as good as its representation of the student's background, goals and consented baseline; without an explicit, correctable profile, personalization degrades into generic answers. A knowledge-tracing agent is needed because skill formation unfolds over time and must be estimated continuously against course outcomes rather than judged from a single score. For Indian engineering education in particular,

mastery cannot be inferred from chatbot interaction; it requires a model that tracks concept-level competence across a semester so that

remediation targets the right gap at the right moment.




Figure 2. India relevance coding.




Figure 3. Evidence type distribution after expansion.




Figure 4. Dominant technical and policy themes in the expanded corpus.




Figure 5. Outcome and concern coverage in the expanded corpus.

5.2 Skill formation rationale

The framework treats skill formation, not content delivery, as the object of measurement, and this dictates how weaknesses are detected and reduced. The evidence on generative AI is genuinely mixed. Gains tend to appear when the tool sits inside a structured task that demands reflection and a redesigned assessment, and they fade when access is left unstructured. Our response is a skill-gap agent that infers weaknesses only from what we can actually observe - quiz responses, coding and design work, viva answers, rubric-scored assignments - and ties each inferred gap to a remediation step and, finally, to a transfer task the student must complete unaided. The coding tasks, rubrics, viva and transfer problems are not decorative. They are what allow us to speak of skill formation rather than familiarity,

because they show what a learner can still do once the support is removed.

5.3 Rationale for Content Grounding

Recommendations are only as trustworthy as their sources, which is why the framework relies on retrieval-augmented, faculty-approved content. Left to generate freely, a model can produce fluent but unsupported answers, drift away from the syllabus, or encourage dependence. The content agent therefore works from a teacher-approved corpus - course notes, NPTEL and SWAYAM material, worked examples and curated exercises - and releases nothing that has not passed a source-approval gate. Grounding the content this way keeps suggestions aligned with the syllabus, makes them auditable, and allows delivery in Indian languages without the teacher losing control of what

is taught. In practice it turns a general-purpose model into a course-specific tutor whose recommendations can be inspected.

5.4 Rationale for the Assessment Approach

Assessment validity has become a first-order concern, so AI-generated assessment is treated as formative unless a teacher decides otherwise. Once answers and code can be produced on demand, a conventional take-home task may no longer reveal what a student can do unaided; in an Outcome-Based programme that directly weakens the

CO/PO evidence on which accreditation rests. The design response is an assessment agent that generates practice items, keys and rubrics for low-stakes use, while any move to summative use requires explicit faculty approval. Process evidence-drafts, version history, viva and reflective declarations-is retained alongside outcomes. This formative-first rule protects the credibility of attainment evidence and ensures that the framework strengthens, rather than erodes, the assessment basis on which Indian programmes are accredited.

Table 2. Comparative role of AI families in learning improvement.

Stream	Key refs.	Core contribution	Strength	Main gap
ITS, KT and LA	Liu (2022), Liang, Rashmi, Doo et al., (2026)	Learner modelling, mastery tracking and learning prediction	Strong modelling base	Limited Indian learner-trace evidence
GenAI in learning	Kasneci, (2023), Yan, Yusuf, Wang et al., (2204), Sergeeva, Deng et al., (2025)	Explanation, feedback, writing, coding and tutoring support	Natural interaction and fast feedback	Weak evidence for durable learning
AI assessment	Kofinas, Chen et al., (2025)	Automated scoring, formative assessment and integrity concerns	Useful for feedback and evaluation redesign	Risky for high-stakes use without faculty approval
Personalisation	Mazhar, Strzelecki, Guan et al., (2024), Ellikkal, Looi et al., (2025)	Adaptive learning paths and engagement support	Improves learner-specific support	Often lacks transparent learner modelling
Agentic AI	Alfredo (2024), Guo (2025)	Multi-step tutoring, planning and workflow orchestration	Promising for structured remediation	Mostly prototype-level evidence
Indian education context	Sudarshan, Swayam (2026)	Policy, platform, language and scale readiness	Strong deployment context	Limited classroom-trial evidence
Responsible AI	Choi (2020), Yan (2024), Gruenhagen (2024)	Privacy, fairness, explainability and auditability	Strong governance direction	Rarely built into learning-system design

Abbr.: ITS = *Intelligent Tutoring System*; KT = *Knowledge Tracing*; LA = *Learning Analytics*; GenAI = *Generative AI*.

5.5 Faculty and OBE rationale

The framework foregrounds faculty dashboards and outcome mapping because, in Indian engineering education, learning support must connect to teaching decisions and to accreditation evidence. National platforms supply scale and content, but the

missing link is learner-level evidence that a teacher can act on and that a department can defend to NBA or NAAC. The design response is a faculty-dashboard agent that summarises class-level misconceptions, at-risk students and module-wise attainment into intervention-ready views, and an

OBE agent that maps learner artifacts to course and programme outcomes. Together they reduce routine feedback burden while keeping the teacher in authority, and they produce the CO/PO evidence that turns personalized learning into institutional quality assurance.

5.6 Governance rationale

Finally, the framework makes governance a built-in layer rather than an afterthought because responsible classroom use is a precondition, not a feature. Learner traces are personal data, recommendations can be biased, and autonomous action without records is unauditable. Under India's

Digital Personal Data Protection regime, consent, purpose limitation, minimization and correction are obligations. The design response is a governance agent that maintains consent and data-use records, enforces field-level minimization, samples AI outputs for hallucination and bias, logs every recommendation and faculty approval, and routes low-confidence or high-impact cases to human escalation. Students keep the right to inspect their profile and to have it corrected. Governance, in this design, is not paperwork bolted on at the end; it is what actually holds the system's autonomy in check and leaves a record that can be examined afterwards.

Table 3. Trial metric mapping: each metric, its purpose, data source and the risk if it is ignored.

Metric	Definition	Data	Caution
KG	(Post-Pre)/(Max-Pre)	Concept inventory	Delayed retention required
SGC	(Gap_pre-Gap_post)/Gap_pre	Skill rubric/code/viva	Independent transfer task required
Ret	Delayed score / immediate score	Follow-up test	Needs 2-6 week delay
Eng	$z(\text{time}) + z(\text{active days}) + z(\text{submit rate})$	LMS/ERP trace	Engagement may not equal learning
EqGap	max group score - min group score	Language/access groups	Requires sufficient subgroup n
HallR	Unsupported AI outputs / audited outputs	Expert audit sample	Needs domain rubric
XAI-U	Mean explanation usefulness	Student/faculty survey	Subjective but actionable
PrivR	Necessary fields / collected fields	Data audit	Higher ratio means less over-collection
EscR	Escalations / high-risk decisions	System logs	Human oversight indicator
OBE-Cov	COs with valid evidence / total COs	Rubric map	Requires faculty validation

Abbr.: KG=knowledge gain; SGC=skill-gap closure; Ret=retention ratio; Eng=engagement index; EqGap=equity gap; HallR=hallucination rate; XAI-U=explanation-usefulness score; PrivR=privacy-minimisation ratio; EscR=escalation rate; OBE-Cov=Outcome-Based Education coverage.

6. Comparative Analysis

It is more useful to group the reviewed work by what it contributes than to march through it paper by paper. Table 2 presents the research streams line up against the framework: intelligent tutoring and knowledge tracing supply the mastery modelling, generative models supply the interaction layer, assessment and agentic methods supply scale and

orchestration, and explainable, responsible-AI work supplies trust and governance. SAKSHAM-AI adds no new family of AI; it arranges existing ones with some discipline, assigning each strength to an agent and answering each limitation with a human checkpoint. The second comparison (Table 3) maps every evaluation metric to its purpose, data source, interpretation and the risk incurred if it is omitted,

so that the framework's value can be tested component by component rather than asserted as a whole. A third comparison-between an ordinary chatbot and the framework-is developed where it is most concrete, in the implementation discussion (Table 5).

7. Research Gap Matrix

The gaps identified in the evidence base are not left as open problems; each is matched to a concrete design response in the architecture and to the trial evidence that would show the response works (Table 4). Read together, the rows make the logic of

the paper explicit: weak agentic validation motivates a bounded pilot, the absence of Indian learner-trace data motivates consented logging, assessment-validity risk motivates formative-first design, multilingual inequity motivates retrieval-augmented support, and privacy and trust concerns motivate a governance agent and explainable dashboards. The relative intensity of these gaps-heaviest in data, deployment and governance for agentic and India-specific concerns-is shown in Figure 6, which is also why those dimensions receive the most design attention.

Table 4. Research-gap matrix: each gap mapped to its SAKSHAM-AI design response and the trial evidence required.

Gap	Problem in current practice	SAKSHAM-AI design response	India expression	Trial evidence needed
G1	Agentic education work is mostly prototypes and preprints, with little controlled classroom validation	Bounded-autonomy architecture with defined agents and human checkpoints, built for a 12–16 week pilot	Few Indian agentic pilots in HE/engineering	Pre/post, delayed and transfer results from a supervised trial
G2	Short perception surveys dominate; durable learning is rarely measured	Knowledge-tracing and skill-gap agents that track mastery over a semester	Limited delayed-retention evidence in Indian HE	Delayed-retention and transfer scores by course outcome
G3	Learner-trace datasets are mostly non-Indian and not consent-cleared	Consented learner-profile and OBE evidence logs, minimised and correctable	No broad Indian HE learner-trace corpus	Anonymised, consent-cleared learner traces and OBE evidence
G4	Tools are English-heavy; regional-language support is uneven	Retrieval-augmented content agent grounded in multilingual, faculty-approved resources	Indic-language learners under-served	Equity-gap (EqGap) results across language groups
G5	AI tools rarely map to course/programme outcomes	OBE agent that maps learner artifacts to CO/PO/PSO with faculty validation	NBA/NAAC evidence burden on departments	OBE-coverage (OBE-Cov) of outcomes with valid evidence
G6	Detector-only responses are weak; take-home tasks no longer prove competence	Formative-first assessment agent plus process logs, viva and rubric transfer tasks	Authentic process assessment uneven across institutions	Validity evidence: independent transfer without AI help
G7	Personalisation needs telemetry, but minimisation and consent are unclear	Privacy-by-design governance agent enforcing minimisation, consent and correction	DPDP duties make minimisation and consent mandatory	Privacy-minimisation ratio (PrivR) and consent audit
G8	Human-centred design is limited; faculty workload and trust are barriers	Explainable faculty-dashboard agent that reduces routine feedback load	Faculty time and trust constraints in large sections	Faculty actionability and explanation-usefulness (XAI-U)

G9	Platform scale is taken as impact, but access constraints persist	Access-sensitive deployment and equity monitoring across cohorts	Rural and tier-2/3 access constraints	Subgroup equity analysis by access and device
G10	Opaque models limit trust and institutional audit	Inspectable learner model, source-approval gate and full audit logs	Institutions require audit trails for AI use	Escalation rate (EscR) and hallucination-rate (HallR) audit

Abbr.: HE=higher education; OBE=Outcome-Based Education; CO=course outcome; PO=programme outcome; PSO=programme specific outcome; NBA/NAAC=Indian accreditation bodies; DPDP=Digital Personal Data Protection; EqGap=equity gap; OBE-Cov=OBE coverage; PrivR=privacy-minimisation ratio; XAI-U=explanation-usefulness score; EscR=escalation rate; HallR=hallucination rate.




Figure 6. Research-gap heatmap. The highest intensities occur in data, deployment and governance dimensions for agentic and India-specific concerns.

8. The SAKSHAM-AI Framework

SAKSHAM-AI-Skill-Aware Knowledge Support and Holistic Adaptive Mentoring through Agentic AI-is the central contribution of this paper. It is not an autonomous tutor that replaces teachers but a human-guided architecture in which ten specialised agents perform defined academic functions under faculty authority. Its organising principle is bounded autonomy: agents may recommend, retrieve, explain, draft, summarise, flag and escalate, while grading, certification and accreditation judgements

remain with teachers and institutional committees. The architecture is organised into five layers-institutional input, student modelling, learning action, faculty and OBE, and governance and audit-through which learner data flows under explicit human checkpoints. The overall structure is shown in Figure 7, and the role, input, output and oversight of each agent are summarised in Table 9. The subsections below describe the layers, the data flow and weekly learning cycle, and the human controls that keep autonomy bounded.

Table 5. SAKSHAM-AI agent roles.

Agent	Input	Output	Oversight
LP	Consent, profile, goals	Correctable learner model	Student view/correct
SKG	Quiz, code, viva, rubric	Skill-gap map	Faculty confirms high-stakes inference
KT	Item attempts, mastery traces	CO-wise mastery timeline	Teacher validates weak concepts
RAG-C	Curriculum + trusted content	Source-grounded learning path	Source approval
ASM	CO, Bloom, rubric	Formative test and key	Faculty approval
FB	Response + misconception	Feedback/remediation plan	Low-confidence escalation
SRL	Goal/progress/time	Reflection prompts	Student privacy control
F-Dash	Class evidence	Intervention dashboard	Teacher authority
OBE	CO-PO-PSO map	Outcome evidence	Dept. review
GOV	Policy/log/bias metrics	Audit record	AI governance committee

Abbr.: LP=learner profile; SKG=skill gap; KT=knowledge tracing; RAG-C=RAG content; ASM=assessment; FB=feedback; SRL=self-regulated learning; F-Dash=faculty dashboard; OBE=Outcome-Based Education; GOV=governance; CO=course outcome; PO=programme outcome; PSO=programme specific outcome.

8.1 Layered architecture

The first layer, institutional input, encodes the academic context that constrains everything downstream: the course and its outcomes, the syllabus, the faculty-approved content pool, assessment rubrics, consent records and institutional policy. Nothing the system recommends or assesses is permitted to fall outside this approved boundary, which is what keeps the architecture course-specific rather than generic.

The second layer, student modelling, builds and maintains the representation of the learner. The learner-profile agent records consented background, goals and baseline performance into a profile the student can view and correct; the knowledge-tracing agent updates concept-level mastery after each interaction; and the skill-gap agent converts that mastery picture into an actionable list of weaknesses, each linked to the observable evidence that produced it. This layer is the analytical core, because every later action depends on an accurate, transparent model of where the learner stands.

The third layer, learning action, is where support is delivered. The retrieval-augmented content agent assembles a personalized learning path from approved resources; the assessment agent generates

formative practice with keys and rubrics; the feedback agent returns misconception-level guidance; and the self-regulated-learning agent prompts goal-setting and reflection. These agents act only on the student model and only within the approved content boundary, so their outputs are both personalized and inspectable.

The fourth layer, faculty and OBE, turns individual learning into teaching and accreditation evidence. The faculty-dashboard agent aggregates class-level misconceptions, at-risk students and module attainment into intervention-ready views, while the OBE agent maps learner artifacts to course, programme and programme-specific outcomes for departmental validation. This layer keeps the teacher in authority and produces the CO/PO evidence that Indian programmes require.

The fifth layer, governance and audit, wraps the others. The governance agent maintains consent and data-use records, enforces field-level minimization, samples outputs for hallucination and bias, logs every recommendation and approval, and escalates low-confidence or high-impact cases. It is this layer that makes bounded autonomy verifiable rather than merely asserted.




Figure 7. SAKSHAM-AI human-guided bounded-autonomy architecture for Indian higher and engineering education.

8.2 Data flow and the weekly learning cycle

Data flows in one principled direction. Institutional input defines the boundary; the student-modelling layer reads learner evidence and produces a current profile, mastery estimate and skill-gap list; the learning-action layer consumes that model to generate paths, practice and feedback; the faculty and OBE layer aggregates the resulting artifacts for teaching and accreditation; and the governance layer observes and records the entire flow. In a typical week the cycle runs: a short diagnostic mapped to a course outcome; an updated skill-gap profile; a personalised set of approved resources and practice tasks; misconception-level feedback; a brief reflection; a faculty review of the class dashboard; an update to OBE evidence; and a governance audit of data use and AI actions. It is the repetition of this cycle across a full semester, rather than any single interaction, that lets the framework demonstrate skill formation instead of momentary engagement.

8.3 Human validation, content correction, source verification and escalation mechanisms

Four controls hold this autonomy in place, each at a fixed point in the cycle. No AI-generated item becomes summative, and no attainment record is finalised, until a teacher signs it off; the high-stakes call is made by a person, not an agent. Learners can

examine their own profile and mastery record and request changes to it, so the model stays answerable to the student it describes.

Source approval requires every recommended resource to originate from the faculty-approved pool, with the content agent unable to introduce ungrounded material. Escalation rules route low-confidence feedback, suspected hallucination, bias signals and high-impact cases to a human reviewer, and the governance agent records each escalation. Together with consent and minimisation, these controls and their audit logs are the operational meaning of bounded autonomy: the system can act usefully and at scale, but only inside boundaries that faculty set and can inspect.

9. Implementation and Evaluation Protocol

The framework is designed to be tested, and this section specifies how. A credible evaluation is a mixed-method, semester-long study in a single course or laboratory module rather than a one-time perception survey, with the minimum protocol summarized in Table 6 and detailed below. Sections or lab batches are assigned to standard support and SAKSHAM-AI support with baseline equivalence checks; where randomisation is not feasible, a matched quasi-experimental design with ANCOVA or mixed-effects modelling is more defensible than

a simple before-after comparison. The subsections set out the setting, participants, duration, design, conditions, weekly cycle, data, analysis and expected outputs of the pilot.

Table 6. Minimum empirical protocol

Element	Recommendation	Why
Design	Cluster RCT or quasi-experiment	Controls selection bias
Duration	12-16 weeks + delayed test	Captures durable learning
Population	Indian HE/engineering; urban + tier-2/3	Tests contextual relevance
Primary outcomes	KG, SGC, Ret, transfer	Measures learning not mere usage
Secondary outcomes	Eng, SRL, trust, workload, EqGap	Explains mechanism and equity
Safety outcomes	HallR, bias audit, PrivR, EscR	Assesses responsible AI
Analysis	ANCOVA/mixed models + subgroup + qualitative coding	Balances rigour and context

Abbr.: RCT=randomised controlled trial; HE=higher education; KG=knowledge gain; SGC=skill-gap closure; Ret=retention ratio; Eng=engagement; SRL=self-regulated learning; EqGap=equity gap; HallR=hallucination rate; PrivR=privacy-minimisation ratio; EscR=escalation rate.

We propose to test the framework in a single course or laboratory module across a twelve-to-sixteen-week semester, with a delayed-retention check two to four weeks later to separate durable learning from a short-lived post-test gain. Where feasible we run a cluster randomised trial with whole sections as clusters; otherwise, a matched quasi-experiment with baseline equivalence is used. The intervention group follows the SAKSHAM-AI weekly cycle - diagnosis, skill-gap mapping, approved resources, formative practice, feedback, reflection, faculty review and governance audit - under the Section 8 controls, while a matched-time control group receives standard support. We collect pre-, post- and delayed scores, rubric-scored and AI-free transfer tasks, and consent and audit logs, storing only what each measure needs, and analyse outcomes with ANCOVA or mixed-effects models alongside subgroup-equity checks. The pilot should yield an inspectable evidence package - student and faculty dashboards, a skill-gap report, a CO/PO sheet, and privacy and escalation audits - reported as expected trial evidence rather than confirmed results.

10. Results Discussion

The central contribution of SAKSHAM-AI is not the use of a chatbot for educational support, but the

design of a controlled agentic learning workflow that can be implemented, audited, and evaluated inside Indian higher and engineering education. The framework is intended to function as a bounded-autonomy system in which different agents perform specific academic tasks under faculty supervision. The learner-profile agent captures consented academic background, learning preferences, course-outcome requirements, and baseline performance. The skill-gap agent identifies weak concepts from quizzes, coding tasks, viva responses, and rubric-based assignments. The knowledge-tracing agent updates the learner's mastery profile after every academic interaction. The recommendation, assessment, feedback, reflection, faculty-dashboard, OBE, and governance agents then operate on this learner model to support personalised remediation and institutional decision-making. The operational role of each SAKSHAM-AI agent and the type of evidence generated during the pilot are summarized in Table 7. The measurable academic, governance, and faculty-facing outputs expected from the trial are presented in Table 8.

In the proposed implementation, SAKSHAM-AI should be deployed first as a controlled pilot in one course or laboratory module rather than as a full institutional platform. The trial run for 12–16 weeks

and can be organized around weekly learning cycles. Each cycle begins with a short diagnostic activity mapped to course outcomes. From the pattern of a student's responses, the system assembles a skill-gap profile and points the learner towards targeted resources - a set of teacher notes, an NPTEL or SWAYAM video, a worked example, a coding or laboratory exercise, or a peer-learning activity. The feedback agent then offers guidance at the level of the specific misconception. What it does not do is decide grades: high-stakes assessment, marking and accreditation evidence still require the teacher's approval, which is how the system can support learning without taking on academic responsibility it should not hold.

The implementation is meant to produce outputs that are useful on both sides of the classroom. A

student sees a dashboard of current mastery, weak concepts, recommended practice, reflection prompts and progress over time. A teacher sees a class-level view instead: the misconceptions that recur, the students slipping behind, attainment broken down by module, and sensible groupings for remediation. The OBE agent rolls individual records up into course-outcome evidence, skill-gap summaries and anonymised attainment trends for the department, while the governance log captures what data was collected, what consent was given, which recommendations the AI made, which a faculty member approved, and where hallucination checks or escalations occurred. Outputs of this kind serve monitoring, mentoring and quality assurance as much as they serve personalised learning.

Table 7. SAKSHAM-AI implementation modules and expected outputs.

Agent	Function	Trial output	Faculty role
Learner Profile	Builds learner profile	Profile sheet	Verify/correct
Skill-Gap	Detects weak skills	Skill-gap map	Confirm gaps
Knowledge-Tracing	Updates mastery	CO-wise mastery curve	Review weak COs
Recommendation	Suggests resources	Personalised path	Approve sources
Assessment	Creates practice	Formative test bank	Approve before use
Feedback	Gives remediation	Misconception feedback	Escalate low-confidence
SRL	Supports reflection	Journal prompts	Review progress
Faculty Dashboard	Summarises class	Intervention dashboard	Plan remediation
OBE	Maps to outcomes	CO/PO evidence sheet	Dept. validation
Governance	Audits safety/privacy	Data/AI-action log	Governance review

CO=course outcome; SRL=self-regulated learning; OBE=Outcome-Based Education

A worthwhile trial has to yield academic outputs, not usage counts. The most basic is a pre/post comparison of concept mastery for each course outcome. Beyond that, a skill-gap closure score shows whether the weaknesses the system identified were actually reduced after remediation, and a follow-up test two to four weeks later checks whether those gains survived. We also ask students

to solve an unseen problem with no AI help, which tests transfer rather than recall, and we record whether the dashboard genuinely helped a teacher locate weak areas and act on them. These measures matter more than satisfaction ratings because they speak to whether learning behaviour changed, not merely whether students used the tool.

Table 8. Evaluation metrics and potential outputs of the SAKSHAM-AI pilot.

Output	Measurement method	Expected evidence
Knowledge gain	Pre/post concept test	CO-level improvement
Skill-gap closure	Rubric pre/post remediation	Fewer weak skills
Delayed retention	Test after 2–4 weeks	Learning stability

Independent transfer	New unseen task, no AI	Independent application
Engagement quality	Active days, submissions	Learning persistence
Faculty actionability	Dashboard log + feedback	Aids intervention planning
Hallucination control	Audit of AI responses	Feedback reliability
OBE evidence	CO/PO-mapped artifacts	Accreditation-ready evidence
Privacy compliance	Minimisation + consent audit	Responsible data use

CO=course outcome; PO=programme outcome; OBE=Outcome-Based Education

The design also lets us draw a sharper contrast between casual AI use and structured AI support. A general-purpose chatbot can answer a question, but it has no idea which course outcome the question serves, what the student previously got wrong, which sources the teacher has sanctioned, or what level the assessment demands. SAKSHAM-AI runs

through a defined academic pipeline instead - diagnosis, recommendation, assessment, explanation, monitoring and escalation - and because each of those steps leaves an observable output, students, faculty and reviewers can check what the system has done. Table 9 sets out the comparison in full.

Table 9. General chatbots and SAKSHAM-AI: unstructured AI interaction versus governed bounded autonomy.

Dimension	General chatbot	SAKSHAM-AI	Why it matters
Learner model	None/session-only	Persistent, correctable profile + KT	Personalisation needs a stable model
Course outcomes	Absent	CO/PO-aware (OBE agent)	Learning must map to attainment
Content source	Open generation	Faculty-approved RAG corpus	Provenance, syllabus alignment
Assessment	Ungoverned	Formative-first; summative needs approval	Protects assessment validity
Feedback	Generic	Misconception-level + escalation	Targets the real gap
Faculty role	Bypassed	Authority retained	Accountability for grades
Privacy	Opaque	DPDP-aware, minimised, consent-based	Lawful data use
Auditability	None	Full audit/approval/escalation logs	Makes autonomy inspectable

Abbr.: CO=course outcome; PO=programme outcome; KT=knowledge tracing; OBE=Outcome-Based Education; RAG=retrieval-augmented generation; DPDP=Digital Personal Data Protection.

If the design has one safeguard that matters above the others, it is bounded autonomy. The system is free to suggest a learning path, generate formative practice, summarise a learner's progress and raise a flag when something looks wrong. It is not free to award final marks, certify attainment or settle any other high-stakes academic question; those stay with the teacher, who keeps the authority while the system plays the part of a structured assistant. For Indian higher education the distinction is not cosmetic, since course-outcome attainment, internal

assessment, mentoring and accreditation all demand a clear line of academic accountability. The framework is built to support the teacher's judgement, not to replace it.

What the first trial should leave behind is a complete, inspectable package: anonymised learner traces, course-outcome-wise mastery records, skill-gap reports, the formative items used, feedback logs, faculty dashboard reports, OBE evidence summaries and governance audit records. With

these in hand, researchers can see where the framework earns its keep and where it does not. It may, for instance, do well at coding remediation and concept revision while struggling with open-ended design or the assessment of communication skills. A finding of that kind is far more useful than a blanket claim that agentic AI improves education, because it names the conditions under which the framework actually works.

The thrust of this discussion is that SAKSHAM-AI deserves to be evaluated as a research-grade intervention rather than demonstrated as a piece of software. Its worth lies in joining personalised learning, skill diagnosis, knowledge tracing, formative assessment, faculty decision-making, OBE evidence and governance into a single workflow that can be audited from end to end. Run through a carefully designed pilot, it can generate exactly the evidence Indian higher and engineering education now lacks: measurable learning gains, narrowing skill gaps, retention that lasts, transfer to new problems, dashboards faculty will actually use, and learner-data practices that can stand up to scrutiny.

12. Conclusion

This paper has developed SAKSHAM-AI, a human-guided agentic framework that converts established AI-in-education evidence into a trial-ready architecture for skill formation in Indian higher and engineering education. Its defining choice is bounded autonomy rather than unrestricted AI tutoring: ten specialised agents diagnose, recommend, assess, explain, monitor and escalate within a layered design, while faculty retain authority over grading, attainment and accreditation, students retain the right to correct their own model, and a governance layer keeps the whole workflow consented, minimised and auditable. The framework's value, however, is conditional. It rests on empirical validation, and the right criteria for that validation are not interface sophistication or usage volume but measurable learning: knowledge gain, skill-gap closure, delayed retention, independent transfer, equity across language and access groups, faculty trust, valid OBE evidence and demonstrable privacy governance. Stated this way, SAKSHAM-AI is best understood not as a finished product but as a serious, testable research agenda—a structure within which the promise of agentic AI can be confirmed or revised, component by component, under the conditions that actually matter for Indian higher and engineering education.

References

- Corbett, A. T., & Anderson, J. R. (1994). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4(4), 253–278.
- Filamant, T. C., Fulgence, R. A., & André, T. (2024). Adapting education for the digital era: Critical insights and reflections. *The International Journal of Technology, Innovation, and Education*.
- Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*, 106(4), 901.
- Amaresh, A. M., Gupta, S., Reddy, V. K., Kumar, R., Singh, T., & Utti, M. S. (2025). A Secure Web-Based Lab Examination System With AI-Driven Code Assist and Network Traffic Control. In *2025 International Conference on Communication, Computer, and Information Technology (IC3IT)* (pp. 01-08). IEEE. Doi: <https://doi.org/10.1109/IC3IT66137.2025.11340949>
- Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, 86(1), 42–78.
- Roll, I., & Wylie, R. (2016). Evolution and revolution in artificial intelligence in education. *International Journal of Artificial Intelligence in Education*, 26(2), 582–599.
- Kuzilek, J., Hlosta, M., & Zdrahal, Z. (2017). Open University learning analytics dataset. *Scientific Data*, 4(1), Article 170171.

- Choi, Y., Lee, Y., Shin, D., Cho, J., Park, S., Lee, S., . . . Heo, J. (2020, June). EdNet: A large-scale hierarchical dataset in education. In *Artificial intelligence in education* (pp. 69–73). Springer International Publishing.
- Liu, Z., Liu, Q., Chen, J., Huang, S., Tang, J., & Luo, W. (2022). pyKT: A Python library to benchmark deep learning based knowledge tracing models. *Advances in Neural Information Processing Systems*, 35, 18542–18555.
- Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., . . . Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, Article 102274.
- Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., . . . Gašević, D. (2024). Practical and ethical challenges of large language models in education: A systematic scoping review. *British Journal of Educational Technology*, 55(1), 90–112.
- Yusuf, A., Pervin, N., & Román-González, M. (2024). Generative AI and the future of higher education: A threat to academic integrity or reformation? Evidence from multicultural perspectives. *International Journal of Educational Technology in Higher Education*, 21(1), Article 21.
- Wang, N., Wang, X., & Su, Y.-S. (2024). Critical analysis of the technological affordances, challenges and future directions of generative AI in education: A systematic review. *Asia Pacific Journal of Education*, 44(1), 139–155.
- Gruenhagen, J. H., Sinclair, P. M., Carroll, J. A., Baker, P. R., Wilson, A., & Demant, D. (2024). The rapid rise of generative AI and its implications for academic integrity: Students' perceptions and use of chatbots for assistance with assessments. *Computers and Education: Artificial Intelligence*, 7, Article 100273.
- Al Shloul, T., Mazhar, T., Abbas, Q., Iqbal, M., Ghadi, Y. Y., Shahzad, T., . . . Hamam, H. (2024). Role of activity-based learning and ChatGPT on students' performance in education. *Computers and Education: Artificial Intelligence*, 6, Article 100219.
- Strzelecki, A., & ElArabawy, S. (2024). Investigation of the moderation effect of gender and study level on the acceptance and use of generative AI by higher education students: Comparative evidence from Poland and Egypt. *British Journal of Educational Technology*, 55(3), 1209–1230.
- Kofinas, A. K., Tsay, C. H. H., & Pike, D. (2025). The impact of generative AI on academic integrity of authentic assessments within a higher education context. *British Journal of Educational Technology*, 56(6), 2522–2549.
- Ellikkal, A., & Rajamohan, S. (2025). AI-enabled personalized learning: Empowering management students for improving engagement and academic performance. *Vilakshan – XIMB Journal of Management*, 22(1), 28–44.
- Gangothri, C. L., & Chandrappa, S. (2026). Plagiarism Detection System using Python with Text Similarity Analysis and Result Visualization. *International Journal of Computational Intelligence in Engineering (IJCIE)*, 1(2), 1-16.
- Lachheb, A., Leung, J., Abramenko-Lachheb, V., & Sankaranarayanan, R. (2025). AI in higher education: A bibliometric analysis, synthesis, and a critique of research. *The Internet and Higher Education*, 67, Article 101021.
- Sergeeva, O. V., Zheltukhina, M. R., Shoustikova, T., Tukhvatullina, L. R., Dobrokhoto, D. A., & Kondrashev, S. V. (2025). Understanding higher education students' adoption of generative AI technologies: An empirical investigation using UTAUT2. *Contemporary Educational Technology*, 17(2), Article ep571.
- Deng, R., Jiang, M., Yu, X., Lu, Y., & Liu, S. (2025). Does ChatGPT enhance student learning? A systematic review and meta-analysis of experimental studies. *Computers & Education*, 227, Article 105224.
- Guan, L., Li, S., & Gu, M. M. (2024). AI in informal digital English learning: A meta-analysis of its effectiveness on proficiency, motivation, and self-regulation. *Computers and Education: Artificial Intelligence*, 7, Article 100323.
- Chen, A., Zhang, Y., Jia, J., Liang, M., Cha, Y., & Lim, C. P. (2025). A systematic review and meta-analysis of AI-enabled assessment in language learning: Design, implementation, and effectiveness. *Journal of Computer Assisted Learning*, 41(1), Article e13064.

- Coban, A., Dzsotjan, D., Küchemann, S., Durst, J., Kuhn, J., & Hoyer, C. (2025). AI support meets AR visualization for Alice and Bob: Personalized learning based on individual ChatGPT feedback in an AR quantum cryptography experiment for physics lab courses. *EPJ Quantum Technology*, 12(1), Article 15.
- Looi, C. K., & Jia, F. (2025). Personalization capabilities of current technology chatbots in a learning environment: An analysis of student-tutor bot interactions. *Education and Information Technologies*, 30(10), 14165–14195.
- Alfredo, R., Echeverria, V., Jin, Y., Yan, L., Swiecki, Z., Gašević, D., & Martinez-Maldonado, R. (2024). Human-centred learning analytics and AI in education: A systematic literature review. *Computers and Education: Artificial Intelligence*, 6, Article 100215.
- Guduru, P., Lakshmi, K. V., & Gupta, S. (2026). Protecting individualized education and therapy records via two-tier integrity verification in distributed educational systems. *International Journal of Special Education*, 41(5s), 386–401.
- Guo, H., Johnson, M. S., Saldivia, L., Worthington, M., & Ercikan, K. (2025, October). Leveraging multi-AI agents for a teacher co-design. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Full papers* (pp. 25–34).
- Jiang, Y. H., Li, R., Zhou, Y., Qi, C., Hu, H., Wei, Y., . . . Wu, Y. (2024). AI agent for education: Von Neumann multi-agent system framework. *arXiv*. <https://arxiv.org/abs/2501.00083>
- Liu, Y., Song, Z., Lou, J., Wu, C., & Li, J. (2025). AgentTutor: Empowering personalized learning with multi-turn interactive teaching in intelligent education systems. *arXiv*. <https://arxiv.org/abs/2601.04219>
- Sudarshan, V. K., Sisodia, A., Ramachandra, R. A., Batra, S., & Leng, J. C. L. (2026). Agentic AI ecosystems in higher education: A perspective on AI agents to emerging inclusive, agentic multi-agent AI framework for learning, teaching and institutional intelligence. *arXiv*. <https://arxiv.org/abs/2605.14266>
- Wu, X., Zhu, P., Zhang, J., Yin, M., & Wang, Y. (2026). ChatGPT's impact on student learning outcomes: A meta-analysis of 35 experimental studies. *Humanities and Social Sciences Communications*.
- Liang, Z., Yang, K., Sha, L., Gašević, D., Yan, L., & Chen, G. (2026). A systematic review of generative AI in education: Empirical insights from a human–AI interaction perspective. *British Journal of Educational Technology*.
- Rashmi, G., Cemalettin, H., Hayri, U., da Silva, C. C. A., Santos, C. B., Patrícia, B., . . . Azevedo, G. A. L. (2026). AI in Indian education: Opportunities, challenges, and emerging paths in the Global South. *Education Sciences*, 16(2), Article 179.
- Doo, M. Y., & Park, Y. (2026). A meta-analysis of ChatGPT's influence on learning achievement. *International Review of Research in Open and Distributed Learning*, 27(1), 265–289.
- Ministry of Education, Government of India. (2020). National Education Policy 2020. https://www.education.gov.in/sites/upload_files/mhrd/files/NEP_Final_English_0.pdf
- SWAYAM. (2026). SWAYAM national online learning portal [Online platform]. <https://swayam.gov.in/>
- Ministry of Electronics and Information Technology / Press Information Bureau. (2025). Digital Personal Data Protection Rules, 2025 notified [Government notice]. <https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc20251117695301.pdf>
- IndiaAI / AIKosh. (2025). AIKosh datasets and AI artefact platform [Online platform]. <https://indiaai.gov.in/>
- Bhashini / Digital India. (2024–2026). Bhashini language technology initiatives [Online platform]. <https://bhashini.gov.in/>
- BharatGen Consortium. (2025). BharatGen sovereign multilingual AI initiative. <https://bharatgen.com/>