BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

PRIVACY-PRESERVING MULTIMODAL AI FOR REAL-TIME PERSONALIZED LEARNING OF STUDENTS WITH DISABILITIES

Aarvee Vijay Nimbalkar

Student, B.E. (Artificial Intelligence and Data Science)
Department of Artificial Intelligence & Data Science Engineering
Ajeenkya DY Patil School of Engineering, Pune, India

Dr. Vishal Puranik

Assistant Professor
Department of Artificial Intelligence & Data Science Engineering
Ajeenkya DY Patil School of Engineering, Pune, India

Abstract

Classrooms that rely on a single, general-purpose AI tool to support students with disabilities tend to cover very little of what those students actually need, because such tools are built for one input or output channel and cannot adjust a lesson in response to how a specific learner is doing. This paper reviews fourteen recent studies (2021-2026) spanning three research strands that bear on this problem: multi-agent AI systems for teaching, disability-focused assistive and agentic AI, and Trust, Risk and Security Management (TRiSM) for agentic systems. The review shows that coordinated multi-agent teaching platforms exist but are evaluated on mainstream learners; disability-focused agentic systems exist but target daily-living routines rather than classroom teaching; and governance frameworks for agentic AI exist but have not been adapted to education or to disability-linked learner data. Building on this gap, we propose a conceptual four-agent framework, comprising a Learning Agent, a Teaching Agent, an Accessibility Agent and a Privacy & Security Agent, that reads a learner's needs from speech, gesture and gaze, adapts the lesson and its presentation in real time, and governs sensitive disability data throughout. We describe the agents' responsibilities, the data flow between them, applicable governance controls, and an evaluation plan. The contribution is a literature-grounded design and research agenda rather than an implemented or empirically validated system.

Keywords: multi-agent systems; large language models; inclusive education; assistive technology; accessibility; agentic AI; AI TRiSM; privacy

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

I. INTRODUCTION

A student with low vision may need enlarged text one day and audio description the next. A student with a motor impairment may switch between gaze, voice and switch input depending on fatigue. No two disabled learners need identical support, and one learner's needs can change within a single session. Software built to perform one function, such as reading text aloud, cannot notice these shifts and has no representation of the lesson it is meant to support.

Three research strands bear on this problem, and this review finds that they have developed largely apart from one another. Multi-agent teaching systems built on large language models (LLMs) coordinate role-specialised agents to plan and refine lessons and materials [1], [4], [6]. Agentic AI has been applied to disability support outside the classroom, for daily routines, communication and mobility [5], [7], [14]. Trust, Risk and Security Management (TRiSM) has matured into a structured governance approach for agentic multi-agent systems generally [2], [3], but has not been adapted to education or to disability-linked learner data specifically.

This paper makes three contributions. First, it reviews fourteen studies published between 2021 and 2026 across the three strands above, comparing their methods, findings and limitations (Section III). Second, it derives a set of research gaps from that comparison (Section IV). Third, it proposes a conceptual four-agent framework, separating understanding the learner, teaching the content, real-time accessibility adaptation and data governance, together with a comparative analysis and an evaluation plan (Sections V-VII). The review is purposive rather than exhaustive, and the proposed framework is a design contribution: it has not been implemented or tested with learners.

II. PROBLEM STATEMENT AND OBJECTIVES

A. Problem Statement

Current AI tools for students with disabilities mostly perform single tasks, such as speech-to-text or gesture detection, without any teaching logic or coordination between them. Multi-agent teaching systems, in turn, target mainstream learners and do not build disability-specific needs, real-time accommodation or protection of sensitive data into their architecture. This work therefore examines whether a four-agent architecture that reads student input through speech, gesture and eye-tracking, and acts on it through coordinated teaching and accessibility agents, can close the gap left by single-function tools, while a dedicated agent keeps sensitive disability data governed under a defined privacy standard.

B. Objectives

1. Lay out a four-agent architecture aimed at disability-first learning rather than mainstream learning with accommodations layered on afterwards.
2. Survey current multi-agent teaching systems, disability-focused agentic AI, and agentic-AI governance frameworks to ground the design in a documented literature gap.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

3. Define a Learning Agent that interprets student needs from speech, gesture and eye-tracking signals.
4. Define Teaching and Accessibility Agents that turn that reading into an adapted lesson and real-time accommodations.
5. Define a Privacy & Security Agent that keeps disability-related data governed and agent behaviour accountable, drawing on TRiSM principles.

III. RELATED WORK

Fourteen studies published between 2021 and 2026 were selected from IEEE, Elsevier, Springer, MDPI and arXiv sources to cover three strands: multi-agent AI for teaching, disability-focused assistive and agentic AI, and AI governance. The review below reports each source's central idea, method, finding and limitation as stated by its authors; the classification into strands and the drawbacks are this review's own reading.

A. Multi-Agent AI for Teaching

Zhang et al. [1] present EduPlanner, in which an evaluator agent, an optimizer agent and a question-analyst agent work in adversarial collaboration to generate and refine mathematics lesson plans. A Skill-Tree structure models each student group's background ability, and a five-dimension evaluation module (clarity, integrity, depth, practicality, pertinence) scores plans and drives iterative optimisation. Experiments on the GSM8K and Algebra datasets show the framework outperforms single-LLM baselines and ablated variants. The demonstration is confined to mathematics and to benchmark datasets, and no accessibility or disability dimension is addressed.

Gonnermann-Müller et al. [4] present FACET, a teacher-facing four-agent system (learner simulation, diagnostic assessment, material generation, evaluation) that differentiates classroom worksheets for motivation, performance and learning differences such as dyslexia and ADHD (only a dyslexia profile was evaluated), while keeping the teacher in the loop. In workshops with 30 school principals and a survey of 70 K-12 teachers, scaffolded worksheets for low-motivation profiles were rated favourably (didactic structure M = 4.41-4.50 of 6), while the dyslexia-adapted worksheet scored lower on font readability (M = 3.30) and enrichment materials for high-performing, highly motivated learners scored lowest overall. The evaluation is limited to grade-7 mathematics, and learner agents simulate behaviour rather than model real students diagnostically.

Sudarshan et al. [6] offer a forward-looking perspective on agentic multi-agent AI ecosystems for higher education, arguing through a thematic literature analysis that current AI tools are fragmented and task-specific, that the field is transitioning from single agents to coordinated multi-agent platforms, and that inclusivity for learners with special educational needs is under-addressed. As a perspective paper it proposes no implemented system, but it explicitly names the same fragmentation-plus-inclusion gap this review targets.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

B. Disability-Focused Assistive and Agentic AI

Jan et al. [5] propose an agentic AI framework for people with disabilities and neurodivergence that coordinates four agents, a Meal Planner, a Reminder Agent, a Food Guidance Agent and a Monitoring Agent, through a shared Blackboard/Event Bus, with a hybrid rule-based and reinforcement-learning reasoning layer, explainable-AI justifications and a caregiver dashboard. On a synthetic dataset simulating eight weeks of activity for 500 users, the reported pilot results include a 27% gain in nutritional adequacy, adherence rising from 54% to 81% with adaptive reminders, and a 35% reduction in caregiver interventions. The demonstration is confined to daily-living routines (diet, scheduling) and uses synthetic rather than real user data; it does not address classroom teaching.

Panagiotou et al. [7] compare three gesture-recognition approaches, an IMU-based machine-learning pipeline, an on-device model, and a MoveNet-plus-CNN vision pipeline, for assisting people with disabilities in a smart-home setting. Piloted with the same volunteer group across methods, the combined MoveNet-and-CNN approach is reported as the most robust for real-time recognition, though each method has scenario-specific strengths and weaknesses. The work targets home automation rather than education and does not link gesture recognition to any teaching or accommodation logic.

Ahmed et al. [8] review AI- and ML-powered inclusive-education tools drawn from databases, blog and company sites, and propose a real-time personalised-learning system model that tracks engagement and development against Bloom's taxonomy. They identify emotional state, progress, preference, learning style and outcome as the key factors for personalisation. The contribution is a system model and tool survey, not an implemented or evaluated platform.

Kooli and Chakraoui [9] examine AI-powered screen readers, voice assistants and NLP interfaces for students with visual, physical and cognitive disabilities through original case analyses (e.g., a visually impaired student's use of screen-reading tools). They report gains in autonomy, engagement and content access, but also document algorithmic bias, data-privacy risk, uneven infrastructure and under-prepared educators, and end with policy recommendations for training, funding and standards. The analysis is qualitative and case-based rather than a proposed system.

Malviya and Rajput [14] describe how AI-driven gesture detection, emotion recognition and text recognition, together with augmentative and alternative communication (AAC) tools such as picture-exchange systems and speech-generating tablets, support communication for people with linguistic or cognitive disabilities, including AI chatbots and interactive robots aimed specifically at students. The chapter is a narrative synthesis of existing AAC and assistive techniques rather than a coordinated system.

C. Governance of Agentic AI (TRiSM)

Raza et al. [2] adapt Trust, Risk and Security Management (TRiSM) to LLM-based agentic multi-agent systems around five pillars, explainability, ModelOps, application security, model privacy and lifecycle governance, and propose a risk taxonomy spanning coordination failures to prompt-based adversarial

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

manipulation. They introduce two measurable metrics, a Component Synergy Score for inter-agent enablement and a Tool Utilization Efficacy score for correct tool use, and situate the framework against the EU AI Act, NIST AI RMF and ISO/IEC 42001. The review is general-purpose and is not instantiated for education or disability data.

Leo et al. [3] present a layered risk-assessment methodology that combines traditional threat modelling with an explicit measure of how a threat degrades system trustworthiness, evaluated through a case study of RoboPMS, a multi-agent portfolio-management system in finance. Their empirical results show that slow-drift and backdoor-style attacks can remain indistinguishable from benign behaviour until a late stage even while accumulating risk earlier, which motivates trajectory-aware rather than step-isolated risk evaluation. The demonstration is financial-sector; it is not applied to education or disability-linked data.

D. Broader AI-in-Education Landscape

Gonzalez-Calatayud et al. [10] systematically review AI for student assessment across Scopus and Web of Science (22 of 454 records after PRISMA screening) and find that most studies emphasise formative assessment and automated grading while rarely stating the pedagogical model behind their design, and that most work concentrates on higher education. Kabudi et al. [11] map 147 studies (2014-2020) on AI-enabled adaptive learning systems and report that systems and frameworks are the most common interventions, Bayesian networks and neural networks the most frequent techniques, and that many designed systems remain experimental rather than deployed in real classrooms. Bozkurt et al. [12] analyse half a century of AI-in-education research (1970-2020) through social network and text-mining methods and identify five recurring themes, including adaptive personalisation and human-AI interaction, while noting that ethics remains a comparatively neglected research area. Munir et al. [13] systematically review AI and ML in digital education and identify six recurring application themes, intelligent tutoring, dropout and performance prediction, adaptive/predictive learning, and analytics, with neural networks and support vector machines the most frequently used algorithms. None of these four reviews examines disability-specific multi-agent teaching.

Table I: Comparative Summary of Reviewed Literature

Ref.	Strand	Approach	Key finding	Limitation
	Multi-agent teaching	Evaluator + optimizer + analyst agents, Skill-Tree model	Outperforms single-LLM baselines on GSM8K/Algebra	Mathematics / datasets only; no disability focus
	Multi-agent teaching	4-agent teacher-facing system (FACET), teacher-in-the-loop	Scaffolded worksheets rated 4.4-4.5/6 by 70 teachers	Grade-7 math only; dyslexia formatting weak

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Ref.	Strand	Approach	Key finding	Limitation
	Multi-agent teaching	Thematic review, agentic ecosystem perspective	Names fragmentation + inclusion gap directly	Perspective paper; no system built
	Disability agentic AI	4-agent daily-living system (diet, reminders, monitoring)	+27% nutrition adequacy, 81% adherence (synthetic)	Daily-living only; synthetic data
	Disability assistive AI	3 gesture-recognition methods compared	MoveNet+CNN most robust for real-time use	Smart-home only; no teaching link
	Disability AI review	Tool survey + personalised learning model (Bloom's)	Identifies 5 key personalisation factors	Model proposed, not evaluated
	Disability AI review	Case analyses of screen readers, voice, NLP tools	Gains in autonomy; bias/privacy/training gaps	Qualitative, case-based
	Disability AAC review	Narrative synthesis of AAC + AI communication aids	Chatbots, robots aid non-verbal students	Synthesis only; no system
	Governance (TRiSM)	5-pillar TRiSM review, risk taxonomy, CSS/TUE metrics	Layered governance + metrics for agent systems	General-purpose; not education-specific
	Governance (security)	Layered threat model, RoboPMS case study	Trajectory-aware risk beats step-isolated checks	Financial sector; not education
	AI-in-ed review	PRISMA review, 22/454 assessment studies	Formative assessment + auto-grading dominate	Pedagogy rarely stated
	AI-in-ed review	Systematic mapping, 147 studies (2014-20)	Systems/frameworks most common; Bayesian nets top ML	Mostly experimental-stage

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Ref.	Strand	Approach	Key finding	Limitation
	AI-in-ed review	50-yr bibliometric + text-mining review	5 themes; ethics under-studied	Broad, not disability-specific
	AI-in-ed review	Systematic review, 6 application themes	Neural nets/SVM most used algorithms	Broad, not disability-specific

IV. RESEARCH GAPS

Reading the fourteen sources together surfaces five gaps that motivate the framework proposed in Section V.

G1. Fragmented single-purpose tools. Assistive systems recognise gesture, speech or text well but are not designed to hand off to, or coordinate with, other components [7], [9], [14].

G2. Disability is largely absent from multi-agent teaching design. Multi-agent teaching platforms adapt content and coordinate roles, but the reviewed systems do not build disability accommodation into that architecture [1], [4]; broader reviews confirm inclusion for special-needs learners is under-addressed [6], [8].

G3. Agentic disability support stops at daily living. The one disability-focused multi-agent system reviewed here targets diet, reminders and routines rather than classroom teaching or lesson adaptation [5].

G4. Governance is general-purpose, not education-specific. TRiSM and layered risk-assessment frameworks for agentic AI exist [2], [3], but neither is instantiated for learner data or classroom deployment, and education-side sources name privacy and bias as concerns without a concrete control architecture [9].

G5. Limited empirical validation. Where systems are tested, evaluation is confined to one subject, uses synthetic data, or is dataset-based rather than classroom-based [1], [4], [5]; much of the disability-focused literature remains review-level [6], [8], [9], [12], [13].

Main gap: the components needed for accessible, secure, real-time AI-assisted teaching, coordinated agent roles, disability-aware personalisation and agentic-AI governance, exist separately in the literature but have not been combined into one system aimed at learners with disabilities.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

V. PROPOSED FOUR-AGENT FRAMEWORK

A. Design Principles

The framework responds directly to the gaps in Section IV: coordinated roles (G1), disability-first design (G2), a teaching-oriented (not only daily-living) adaptation loop (G3), and education-specific governance (G4). Four principles guide the design.

- Separation of concerns: understanding, teaching, adapting and protecting are handled by four distinct agents, following the role-specialised pattern used for lesson design in [1] and for teacher-facing differentiation in [4].
- Disability-first, not accommodation-added: learner needs are an input to lesson generation itself, not a patch applied afterwards, addressing the gap named in [6] and [8].
- Governance in the message path: every inter-agent message and data access is checked by the Privacy & Security Agent, adapting the TRiSM pillars and risk taxonomy of [2] and the layered, trajectory-aware risk view of [3].
- Human oversight: educators and guardians can review, explain and override agent decisions, consistent with the teacher-in-the-loop principle validated with 30 school principals in [4] and the caregiver-dashboard pattern in [5].

B. Agent Architecture

Fig. 1 Proposed four-agent framework with governance layer

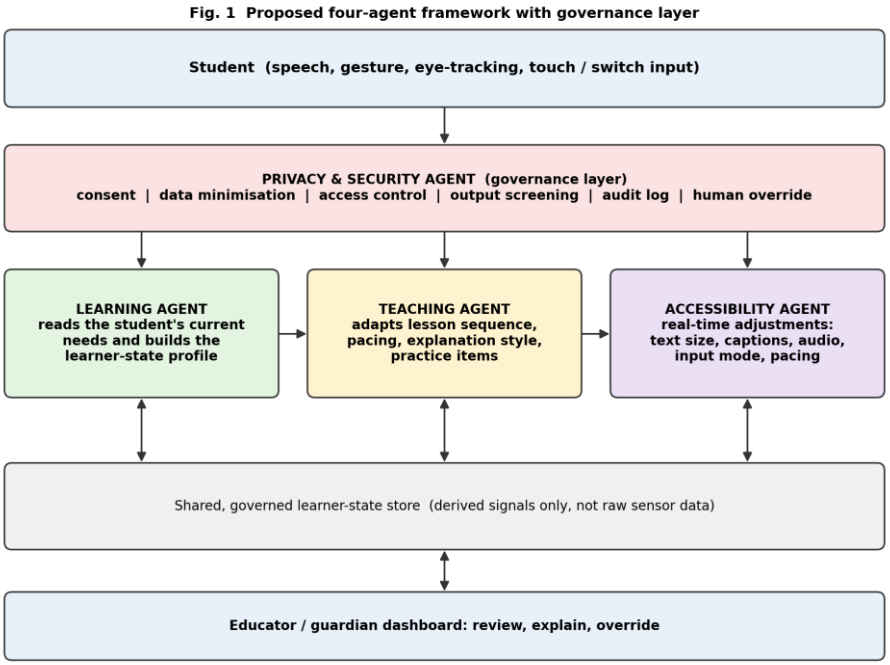


Fig. 1. *Proposed four-agent framework with a governance layer mediating all learner data.*

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Fig. 1 shows the four agents. The Learning Agent parallels the perception role of the gesture-recognition pipeline in [7] and the Learner Agent in [4], but combines speech, gesture and eye-tracking into one learner-state profile rather than a single modality. The Teaching Agent adapts the lesson plan the way the evaluator-optimizer loop in [1] refines instructional content, but conditions on accessibility needs rather than only test performance. The Accessibility Agent applies moment-to-moment adjustments (text size, captions, audio description, input mode), a role with no direct counterpart in the reviewed multi-agent teaching systems, which is precisely gap G2. The Privacy & Security Agent applies the TRiSM pillars from [2] to every message and data request, rather than governing the system only at the perimeter.

Table II: Agent Responsibilities

Agent	Responsibility	Key inputs	Key outputs
Learning	Build learner-state profile	Speech, gesture, gaze, interaction logs	Engagement, fatigue, preferred modality
Teaching	Adapt lesson content	Learner-state profile, curriculum goals	Sequence, pacing, practice items
Accessibility	Real-time accommodation	Learner-state profile, pending content	Text size, captions, audio, input mode
Privacy & Security	Govern data + agent behaviour	All messages, data requests	Allow/redact/block, audit log

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

C. Data Flow

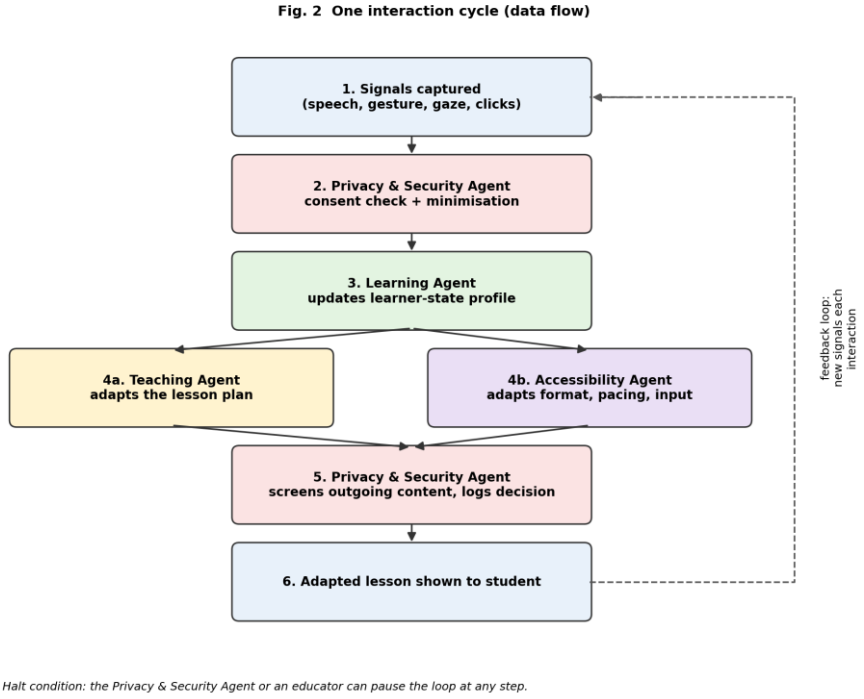


Fig. 2. *One interaction cycle: signal capture, governed processing, parallel adaptation, and delivery.*

A single interaction cycle proceeds as follows: (1) sensors and interface events produce raw signals; (2) the Privacy & Security Agent applies consent checks and data minimisation before any other agent sees the data; (3) the Learning Agent updates the learner-state profile; (4) the Teaching and Accessibility Agents act on that profile in parallel, one adapting the lesson, the other its presentation; (5) the Privacy & Security Agent screens outgoing content and logs the decision; (6) the adapted lesson is shown, and new signals restart the cycle. Passing a structured, derived profile between agents rather than raw sensor data limits how much sensitive information travels through the system, addressing a concern raised generally for agentic pipelines in [2] and specifically for learner data in [9].

D. Governance Controls

Table III maps the TRiSM pillars identified in [2] and the layered threat view of [3] onto concrete controls for the Privacy & Security Agent.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Table III: Governance Controls Mapped to TRiSM

Control point	Mechanism	Basis in reviewed work
Data intake	Consent check, minimisation, encryption	Model privacy pillar [2]
Inter-agent messages	Schema validation, least-privilege access	Coordination-failure risk [2], [3]
Model output	Screening for unsafe or biased content	Explainability pillar [2]
Oversight	Audit trail, human override	Lifecycle governance [2]

VI. COMPARATIVE ANALYSIS

Fig. 3 scores the reviewed literature clusters and the proposed framework against the five dimensions that define the gaps in Section IV, on a 0 (absent) to 3 (strong) scale assigned from what each cluster reports. The proposed framework's scores are design intentions, not measured results; the chart makes explicit that only real-world validation with learners (the rightmost dimension) remains unaddressed by any reviewed source, including the proposed design.

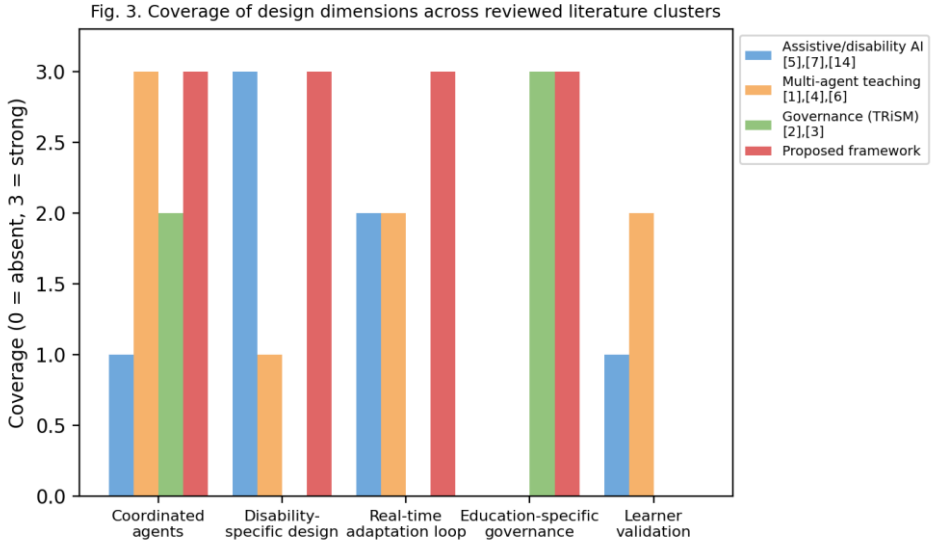


Fig. 3. *Coverage of five design dimensions across literature clusters and the proposed framework.*

Fig. 4 places the reviewed sources against an illustrative motif of how agentic-AI and AI-in-education publication activity has grown, drawn from the trend descriptions in the governance review [2], the ecosystem perspective [6], the systematic mapping [11] and the 50-year bibliometric review [12]. The curves are illustrative of the direction and relative pace those sources describe; they are not extracted

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

publication counts and should not be read as precise bibliometric data.

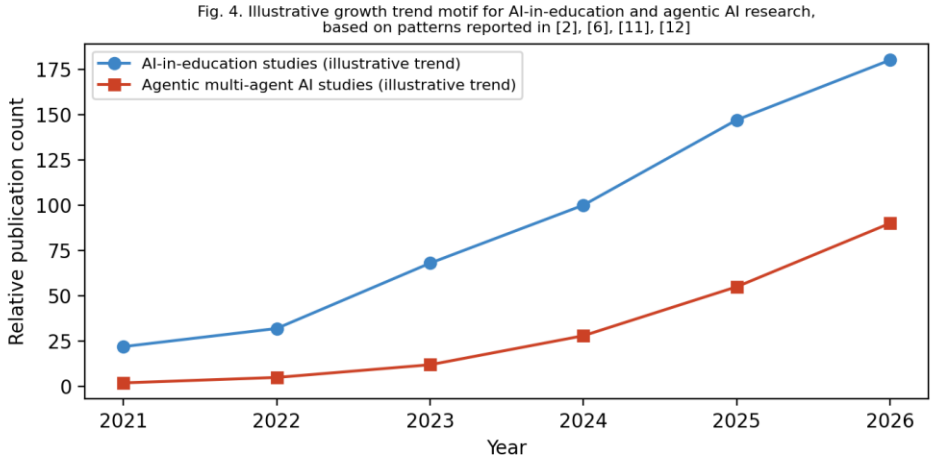


Fig. 4. *Illustrative growth motif for AI-in-education and agentic-AI research, based on trends reported in [2], [6], [11], [12].*

VII. ETHICAL AND PRACTICAL CONSIDERATIONS

Disability-linked data (gaze, speech, gesture, health-adjacent signals) is highly sensitive; the literature review above shows this concern recurring across strands rather than being specific to education. Kooli and Chakraoui [9] document algorithmic bias and privacy risk in existing assistive tools; Raza et al. [2] catalogue data leakage and adversarial manipulation as agentic-AI risks generally; and Jan et al. [5] build an explainable-AI layer specifically so caregivers can question system decisions involving disability-related routines. The proposed framework follows these precedents by minimising collected data, passing derived state rather than raw signals between agents, and keeping a human reviewer, an educator here rather than a caregiver, in the decision loop as in [4] and [5].

Bias is a second concern. Recognition models can perform unevenly on atypical speech, gesture or gaze patterns, and Panagiotou et al. [7] report that no single recognition method dominates across all scenarios, which argues for per-group evaluation rather than a single aggregate accuracy figure. Sudarshan et al. [6] and Ahmed et al. [8] both note that learners with special educational needs are under-represented in the design and evaluation of current systems, a participation gap the framework cannot close by architecture alone; co-design with disabled learners and educators is a precondition for responsible deployment, not an optional extension.

VIII. EVALUATION PLAN

Because the framework has not been implemented, Table IV sets out the questions and measures a future study would need, adapting metrics already proposed in the reviewed literature rather than inventing new ones where a suitable metric exists.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Table IV: Proposed Evaluation Plan

Question	Measure	Source of metric
Does adaptation help learning?	Comprehension vs. non-adaptive baseline	Worksheet ratings method in [4]
Do agents coordinate well?	Component Synergy Score, Tool Util. Efficacy	Metrics proposed in [2]
Is recognition fair across users?	Per-group error rate (speech/gesture/gaze)	Multi-method comparison approach in [7]
Is data protected?	Audit against layered risk taxonomy	Threat-modelling method in [3]
Is it usable for teachers?	Teacher rating survey (Likert, open feedback)	Survey design in [4]

Studies involving learners would require ethical approval and informed consent from participants and guardians, following the consent and voluntary-participation statement adopted in [4].

IX. CONCLUSION AND FUTURE WORK

This review examined fourteen studies across multi-agent AI for teaching, disability-focused assistive and agentic AI, and TRiSM governance for agentic systems, and found that each strand solves part of the problem of AI-assisted teaching for disabled learners without solving all of it: coordinated agents exist for mainstream teaching [1], [4]; coordinated agents exist for disability support outside the classroom [5]; and governance frameworks exist for agentic AI generally [2], [3]. Building on this gap, we proposed a conceptual four-agent framework, Learning, Teaching, Accessibility and Privacy & Security, that combines these threads for disability-first classroom teaching. The contribution is a literature-grounded design and evaluation agenda; it has not been implemented or tested with learners.

Future work should prioritise a prototype for one disability category, built and evaluated with the participation of disabled learners and educators, following the co-design gap identified in [6] and [8]; a shared, versioned learner-state schema; and an empirical study measuring the coordination metrics proposed in [2] alongside the fairness and usability measures in Table IV.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

REFERENCES

[1] X. Zhang, C. Zhang, J. Sun, J. Xiao, Y. Yang, and Y. Luo, "EduPlanner: LLM-Based Multiagent Systems for Customized and Intelligent Instructional Design," *IEEE Trans. Learn. Technol.*, vol. 18, pp. 416-427, 2025, doi: 10.1109/TLT.2025.3561332.

[2] S. Raza, R. Sapkota, M. Karkee, and C. Emmanouilidis, "TRiSM for Agentic AI: A Review of Trust, Risk, and Security Management in LLM-based Agentic Multi-Agent Systems," arXiv:2506.04133, 2025.

[3] M. Leo, F. Tan, T. Miao, and G. Anand, "From threat to trust: assessing security risks of agentic AI systems," *Int. J. Inf. Secur.*, vol. 25, no. 1, art. 23, 2026, doi: 10.1007/s10207-025-01185-y.

[4] J. Gonnermann-Müller, J. Haase, N. Leins, M. Igel, K. Fackeldey, and S. Pokutta, "FACET: Multi-Agent AI Supporting Teachers in Scaling Differentiated Learning for Diverse Students," arXiv:2601.22788, 2026.

[5] S. Jan, T. A. Syed, A. Akarma, A. Ali, and M. R. Belgaum, "Agentic AI Framework for Individuals with Disabilities and Neurodivergence: A Multi-Agent System for Healthy Eating, Daily Routines, and Inclusive Well-Being," in *Proc. Int. Conf. Business and Digital Technology*, Bahrain, Springer Nature, 2025.

[6] V. K. Sudarshan, A. Sisodia, R. A. Ramachandra, B. Sia, and J. C. L. Leng, "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," Nanyang Technological Univ., 2025.

[7] C. Panagiotou, E. Faliagka, C. P. Antonopoulos, and N. Voros, "Multidisciplinary ML Techniques on Gesture Recognition for People with Disabilities in a Smart Home Environment," *AI*, vol. 6, no. 1, art. 17, 2025, doi: 10.3390/ai6010017.

[8] S. Ahmed, M. S. Rahman, M. S. Kaiser, and A. S. M. S. Hosen, "Advancing Personalized and Inclusive Education for Students with Disability Through Artificial Intelligence: Perspectives, Challenges, and Opportunities," *Digital*, vol. 5, no. 2, art. 11, 2025, doi: 10.3390/digital5020011.

[9] C. Kooli and R. Chakraoui, "AI-driven assistive technologies in inclusive education: benefits, challenges, and policy recommendations," *Sustain. Futures*, vol. 10, art. 101042, 2025, doi: 10.1016/j.sftr.2025.101042.

[10] V. Gonzalez-Calatayud, P. Prendes-Espinosa, and R. Roig-Vila, "Artificial Intelligence for Student Assessment: A Systematic Review," *Appl. Sci.*, vol. 11, no. 12, art. 5467, 2021, doi: 10.3390/app11125467.

[11] T. Kabudi, I. Pappas, and D. H. Olsen, "AI-enabled adaptive learning systems: A systematic mapping

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

of the literature," *Comput. Educ.: Artif. Intell.*, vol. 2, art. 100017, 2021, doi: 10.1016/j.caeai.2021.100017.

[12] A. Bozkurt, A. Karadeniz, D. Baneres, A. E. Guerrero-Roldan, and M. E. Rodriguez, "Artificial Intelligence and Reflections from Educational Landscape: A Review of AI Studies in Half a Century," *Sustainability*, vol. 13, no. 2, art. 800, 2021, doi: 10.3390/su13020800.

[13] H. Munir, B. Vogel, and A. Jacobsson, "Artificial Intelligence and Machine Learning Approaches in Digital Education: A Systematic Revision," *Information*, vol. 13, no. 4, art. 203, 2022, doi: 10.3390/info13040203.

[14] R. Malviya and S. Rajput, "AI to Empower People with Disabilities Communication and Autonomy," in *Advances and Insights into AI-Created Disability Supports*, Springer, 2025, pp. 97-113, doi: 10.1007/978-981-96-6069-8_6.