International Journal of Emerging Research in Engineering and Technology
Pearl Blue Research Group| AECTIC2025-Conference Proceeding
ISSN: 3050-922X | https://doi.org/10.63282/3050-922X.AECTIC-114

Original Article

Agentic AI in Inclusive Learning: A Framework for Autonomous Personalization across Diverse Learner Populations

Antony Ronald Reagan Panguraj
Sr. Software Engineer, Yzenx Inc., USA.

Abstract *- The emergence of agentic artificial intelligence represents a fundamental shift in how educational technology can address the needs of diverse learner populations. Unlike traditional adaptive learning systems that respond to predefined triggers, agentic AI operates with genuine autonomy, pursuing complex learning objectives while continuously adapting strategies based on evolving learner states. This paper presents a comprehensive framework for deploying agentic AI in inclusive educational environments, organized around ten interconnected pillars: personalized learning pathways, dynamic scaffolding, multimodal experiences, accessibility enhancements, continuous learner modelling, autonomous feedback, cultural and linguistic inclusivity, teacher augmentation, motivational support, and ethical guardrails. Drawing on Self Determination Theory and Universal Design for learning principles, we argue that agentic AI systems must satisfy fundamental psychological needs for autonomy, competence, and relatedness while providing multiple means of engagement, representation, and expression. The framework addresses critical implementation challenges including algorithmic bias, data privacy, and the preservation of human agency in educational decision making. We propose architectural specifications for responsible deployment and discuss implications for educational equity. This work contributes both theoretical grounding and practical guidance for researchers and practitioners seeking to harness agentic AI capabilities while maintaining commitment to inclusive, learner centred education.*

Keywords *- Agentic AI, Inclusive Education, Adaptive Learning, Universal Design for Learning, Intelligent Tutoring Systems, Educational Equity, Algorithmic Fairness.*

1. Introduction

Contemporary educational systems face an enduring challenge: how to provide genuinely personalized instruction at scale while ensuring equitable outcomes for all learners. Traditional classroom models, constrained by time, resources, and the practical limits of human attention, struggle to accommodate the full spectrum of learner variability. Students arrive with different prior knowledge, cultural backgrounds, cognitive profiles, physical abilities, and motivational orientations. Meeting each learner where they are has long been recognized as pedagogically ideal yet practically elusive (Bloom, 1984). The promise of technology to bridge this gap has animated educational research for decades, from early computer assisted instruction to contemporary adaptive learning platforms. Artificial intelligence in education has evolved considerably since its inception. Early intelligent tutoring systems demonstrated that computational approaches could model student knowledge and provide individualized feedback (Corbett & Anderson, 1994). Subsequent generations incorporated more sophisticated learner models, drawing on Bayesian knowledge tracing and item response theory to estimate mastery states and predict performance (Piech et al., 2015). Yet these systems, however advanced, remained fundamentally reactive. They responded to student actions according to predetermined rules, lacking the capacity for genuine goal directed behaviour or autonomous strategy adaptation.

The advent of agentic AI marks a qualitative departure from this paradigm. Agentic systems are characterized by their capacity to pursue complex objectives with minimal human intervention, adapting their plans and actions to changing circumstances (Hosseini & Seilani, 2025). Where traditional adaptive systems follow scripted branching logic, agentic AI can reason about learning goals, generate novel instructional strategies, and coordinate multiple interventions across extended time horizons. This represents what Acharya et al. (2025) describe as a qualitative leap in educational technology, wherein agents plan, utilize tools, maintain memory, and self-adjust to accomplish multistep, multigoal tasks. The implications for inclusive education are profound. Learners with disabilities, English language learners, students from underrepresented cultural backgrounds, and those with diverse cognitive profiles have historically been underserved by one size fits all instructional approaches. Inclusive education demands not merely accommodation but proactive design that anticipates and addresses learner variability from the outset (Rose & Meyer, 2002). Agentic AI, with its capacity for continuous adaptation and autonomous decision making, offers unprecedented potential to realize this vision at scale.

However, realizing this potential requires careful attention to both technical architecture and ethical considerations. Algorithmic bias in educational systems can perpetuate and amplify existing inequities (Baker & Hawn, 2021). Data privacy concerns intensify when AI systems collect detailed behavioural and performance information. The appropriate balance between AI autonomy and human oversight remains contested. These challenges demand frameworks that integrate technical capability with principled design.

This paper addresses this need by presenting a comprehensive framework for agentic AI in inclusive learning environments. We organize the framework around ten interconnected pillars that span the full scope of inclusive educational practice: personalized learning pathways, dynamic scaffolding and real time support, multimodal learning experiences, accessibility enhancements for learners with disabilities, continuous learner modelling, autonomous feedback and assessment, cultural and linguistic inclusivity, teacher augmentation and reduced cognitive load, motivational and behavioural support, and ethical guardrails for responsible autonomy. Each pillar is grounded in established learning theory and informed by current empirical research on AI in education.

The framework draws on two foundational theoretical perspectives. Self Determination Theory posits that human motivation and wellbeing depend on satisfaction of three basic psychological needs: autonomy, competence, and relatedness (Ryan & Deci, 2000). Educational environments that support these needs foster intrinsic motivation and deeper learning. Universal Design for Learning provides complementary design principles, emphasizing multiple means of engagement, representation, and action and expression to accommodate predictable learner variability (CAST, 2024). Together, these frameworks establish criteria for evaluating whether agentic AI systems genuinely serve inclusive educational goals.

The paper proceeds as follows. Section 2 establishes theoretical and empirical background, distinguishing agentic AI from prior approaches and reviewing evidence on AI effectiveness in education. Sections 3 through 12 elaborate each of the ten framework pillars, articulating design principles and implementation considerations. Section 13 presents an integrated system architecture. Section 14 addresses implementation challenges and proposes mitigation strategies. Section 15 discusses limitations and future research directions. Section 16 concludes with implications for policy and practice.

2. Background and Theoretical Foundations

2.1. From Adaptive to Agentic: An Evolutionary Perspective

The trajectory of AI in education reflects broader developments in artificial intelligence. First generation intelligent tutoring systems emerged in the 1970s and 1980s, encoding expert knowledge in rule-based systems that could diagnose student errors and provide targeted feedback (Carbonell, 1970). Systems like LISP Tutor and Geometry Tutor demonstrated that computers could provide individualized instruction rivalling human tutors in some domains (Anderson et al., 1985). Second generation systems incorporated probabilistic models of student knowledge. Bayesian Knowledge Tracing allowed systems to estimate the probability that a student had mastered specific skills based on their response history (Corbett & Anderson, 1994). This enabled more nuanced adaptation, with systems adjusting problem difficulty and sequencing based on inferred knowledge states rather than simple right wrong tallies.

The deep learning revolution brought further advances. Deep Knowledge Tracing applied recurrent neural networks to model complex temporal patterns in student learning (Piech et al., 2015). Transformer architectures enabled even more sophisticated sequence modelling, capturing long range dependencies in learning trajectories (Neshaei et al., 2024). These approaches improved predictive accuracy but remained fundamentally reactive, responding to observed behaviour rather than proactively pursuing learning goals. Agentic AI represents a paradigm shift. Drawing on advances in large language models and autonomous agent architectures, agentic systems exhibit genuine goal directed behaviour. They can decompose complex learning objectives into subgoals, select and sequence instructional strategies, monitor progress toward objectives, and adjust course when strategies prove ineffective (Hosseini & Seilani, 2025). Crucially, they maintain persistent memory and context across extended interactions, enabling coherent support over entire courses or academic terms rather than isolated learning sessions.

Table 1: Comparison of Traditional Adaptive Systems and Agentic AI in Education

Dimension	Traditional Adaptive Systems	Agentic AI Systems
Decision Autonomy	Rule based responses following predetermined branching logic	Goal directed planning with autonomous strategy selection
Temporal Scope	Immediate reaction to current input	Long term trajectory optimization across sessions
Learner Model	Static profiles updated incrementally	Dynamic, evolving representations with contextual memory
Intervention Style	Selection from predetermined content paths	Generation of novel instructional approaches
Inclusivity Approach	Accommodations as add on features	Inclusive design embedded from the outset
Feedback Mechanism	Template based responses	Contextually generated explanatory feedback

Dimension	Traditional Adaptive Systems	Agentic AI Systems
Coordination	Single system operation	Multi agent collaboration for complex tasks

2.2. Theoretical Foundations: Self Determination Theory

Self Determination Theory provides a robust framework for understanding how educational environments affect student motivation and learning outcomes. Developed by Deci and Ryan (1985, 2000), SDT posits that humans have innate psychological needs for autonomy, competence, and relatedness. Autonomy refers to experiencing behaviour as volitional and self-endorsed rather than controlled by external pressures. Competence involves feeling effective in interactions with the environment and having opportunities to express capabilities. Relatedness concerns feeling connected to others and experiencing a sense of belonging. When educational environments support these needs, students exhibit more autonomous forms of motivation, characterized by genuine interest and personal valuing of learning activities (Niemiec & Ryan, 2009). Autonomous motivation, in turn, predicts deeper engagement, better academic performance, greater persistence, and enhanced psychological wellbeing. Conversely, environments that thwart basic needs produce controlled motivation or amotivation, associated with surface learning, disengagement, and poorer outcomes.

For agentic AI systems, SDT provides critical design guidance. Systems that offer meaningful choices, provide rationales for learning activities, and acknowledge learner perspectives support autonomy. Those that offer optimal challenge, specific feedback, and opportunities for mastery support competence. And those that convey warmth, demonstrate understanding, and foster collaborative interaction support relatedness. Agentic AI that systematically attends to these needs positions learners to internalize educational goals and engage in self-regulated learning.

2.3. Theoretical Foundations: Universal Design for Learning

Universal Design for Learning emerged from recognition that learner variability is the norm rather than the exception (Rose & Meyer, 2002). Developed by CAST, UDL provides a framework for designing flexible learning environments that anticipate and address diverse learner needs from the outset. The framework is organized around three principles corresponding to distinct brain networks involved in learning. The first principle, multiple means of engagement, addresses the affective dimension of learning. Learners differ in what motivates them, how they respond to challenge, and what sustains their effort. UDL calls for providing options for recruiting interest, sustaining effort and persistence, and supporting self-regulation. The second principle, multiple means of representation, addresses how learners perceive and comprehend information. Learners vary in sensory abilities, language backgrounds, and prior knowledge. UDL calls for options for perception, language and symbols, and comprehension. The third principle, multiple means of action and expression, addresses how learners navigate learning environments and demonstrate knowledge. Learners vary in physical abilities, executive functions, and communication strengths. UDL calls for options for physical action, expression and communication, and executive function.

The updated UDL Guidelines 3.0, released in 2024, emphasize learner agency as the ultimate goal, characterized by being purposeful and reflective, resourceful and authentic, and strategic and action oriented (CAST, 2024). This aligns with SDT emphasis on autonomous motivation and self-determination. Together, SDT and UDL establish that inclusive educational design must address both motivational and accessibility dimensions, supporting learners not merely in accessing content but in developing as self-directed, agentic learners themselves.

2.4. Evidence on AI Effectiveness in Education

Meta analytic evidence supports the effectiveness of intelligent tutoring systems in promoting learning outcomes. Ma et al. (2014) synthesized 107 effect sizes involving over 14,000 participants and found that ITS use was associated with greater achievement compared to teacher led large group instruction ($g = 0.42$), non ITS computer based instruction ($g = 0.57$), and no treatment controls. Effect sizes were positive across educational levels and subject domains. Notably, the advantage over individual human tutoring was negligible ($g = -0.11$), suggesting that well designed ITS can approach the effectiveness of one-on-one instruction.

More recent systematic reviews confirm these patterns while highlighting implementation considerations. A 2025 review of AI driven intelligent tutoring systems in K-12 education found generally positive effects on learning and performance, though effects were attenuated when compared to non-intelligent tutoring systems (Karran et al., 2025). The authors emphasize that ITS effectiveness depends on embodying sound pedagogical features under appropriate conditions, including immediate feedback, guided practice, and genuine adaptivity.

Research specifically examining agentic AI in education remains nascent but promising. Studies of generative AI tutors suggest benefits for personalization and engagement, with one review finding that students receiving personalized instruction through AI powered systems demonstrated higher engagement, improved academic performance, and increased knowledge retention compared to traditional settings (Batsaikhan & Correia, 2024). However, concerns persist regarding accuracy, explanation quality, and potential for perpetuating biases.

3. Personalized Learning Pathways

The first pillar of the framework addresses how agentic AI can construct and continuously refine individualized learning trajectories. Traditional approaches to personalization have relied on branching algorithms that select among predetermined content sequences based on diagnostic assessments or performance on prerequisite skills. While representing an advance over uniform instruction, such approaches remain constrained by the imagination of their designers and the comprehensiveness of content libraries. Agentic AI enables a more dynamic conception of personalization. Rather than selecting from fixed paths, agentic systems can reason about learning objectives, assess current learner states, and generate novel sequences tailored to individual needs and circumstances. This capacity draws on reinforcement learning techniques that optimize learning trajectories based on predicted outcomes, knowledge graph representations that capture relationships among concepts and skills, and generative models that can produce or adapt instructional content.

Effective personalization requires continuous calibration to the zone of proximal development, the region between what learners can accomplish independently and what remains beyond reach even with support (Vygotsky, 1978). Agentic AI can maintain dynamic estimates of each learner zone and adjust both content difficulty and support intensity accordingly. When learners demonstrate mastery, the system advances to more challenging material. When learners struggle, it provides scaffolding, revisits prerequisites, or offers alternative explanations. This calibration operates not merely at the level of individual problems but across entire learning progressions, ensuring coherent advancement toward distal objectives.

For inclusive education, personalized pathways must accommodate diverse starting points and diverse goals. Not all learners aim for identical outcomes, nor should they. Some may be preparing for advanced study, others for immediate application in workplace contexts, still others for general understanding sufficient to informed citizenship. Agentic AI can negotiate learning objectives with learners, adapting not merely the path but the destination to align with individual purposes and aspirations. This respects learner autonomy while ensuring that instructional support remains relevant and motivating.

4. Dynamic Scaffolding and Real Time Support

Scaffolding refers to instructional support that enables learners to accomplish tasks they could not complete independently, with the expectation that support will fade as competence develops (Wood, Bruner, & Ross, 1976). Effective scaffolding is responsive, contingent on learner needs rather than applied uniformly. It appears when struggle emerges and withdraws as mastery consolidates. Agentic AI is uniquely positioned to provide such responsive scaffolding at scale. Through continuous monitoring of learner behaviour and performance, agentic systems can detect emerging difficulty before learners become frustrated or disengaged. Indicators might include increased response latency, error patterns suggesting misconceptions, reduced confidence expressions, or physiological signals of cognitive overload in systems incorporating affective sensing.

Upon detecting need, agentic systems can deploy appropriately calibrated support. This might involve hints that direct attention without revealing solutions, worked examples that demonstrate problem solving processes, conceptual explanations that address underlying misconceptions, or decomposition of complex tasks into manageable steps. Critically, support intensity can be titrated to learner needs, beginning with minimal prompts and escalating only when lighter touch interventions prove insufficient. The fading dimension of scaffolding is equally important. Prolonged support can undermine competence development and foster dependence. Agentic AI can systematically reduce support as learners demonstrate increasing capability, challenging them to perform independently while remaining ready to reintroduce support if difficulty recurs. This graduated release of responsibility mirrors effective human tutoring and promotes the development of self-regulated learning capabilities.

For diverse learners, scaffolding must be flexibly configured to individual profiles. Learners with learning disabilities may benefit from consistent structural supports that remain available even as content difficulty increases. English language learners may need linguistic scaffolds that fade more slowly than conceptual supports. Learners with attention differences may benefit from scaffolds that maintain focus and reduce extraneous cognitive load. Agentic AI can maintain rich models of individual scaffolding needs and adjust support profiles accordingly.

5. Multimodal Learning Experiences

Learners differ in how they best perceive and process information. Some are strongly visual, benefiting from diagrams, animations, and spatial representations. Others are more auditory, learning effectively through spoken explanation and discussion. Still others are kinaesthetic, developing understanding through physical manipulation and embodied experience. While learning styles theories have been critiqued for overgeneralization, the fundamental observation that learners vary in modality preferences and strengths remains well supported.

Universal Design for Learning emphasizes providing multiple means of representation to accommodate this variability. Content should be accessible through multiple sensory channels, with visual information supplemented by auditory description and vice versa. Symbolic representations should be clarified through concrete examples and multiple media. Background knowledge should be activated and supplied where lacking.

Agentic AI can operationalize these principles through dynamic multimodal content generation and selection. Drawing on advances in generative AI across modalities, agentic systems can produce text explanations, spoken narration, visual diagrams, interactive simulations, and video demonstrations tailored to content requirements and learner preferences. When a learner struggles with textual explanation, the system might generate a visual representation. When static imagery proves insufficient, it might offer animated demonstration or interactive exploration.

This multimodal flexibility serves accessibility as well as preference. Learners who are blind or have low vision can access content through audio and tactile modalities. Learners who are deaf or hard of hearing can access content through visual and textual modalities. Learners with print disabilities can engage through audio, video, and interactive alternatives. Rather than requiring separate accommodation systems, agentic AI can provide seamless multimodal experience as the default, with modality emphasis adjusted to individual needs.

Recent advances in multimodal generative AI make such flexibility increasingly feasible. Large language models can generate coherent text explanations at varying complexity levels. Text to speech systems produce increasingly natural spoken output. Text to image and text to video models can generate visual content from descriptions. Integration of these capabilities within agentic frameworks enables on demand content generation across modalities, moving beyond fixed content libraries to truly responsive multimodal experience.

6. Accessibility Enhancements for Learners with Disabilities

Approximately 15 percent of the global population lives with some form of disability, and prevalence is higher among school age populations when learning disabilities are included (World Health Organization, 2011). Ensuring that AI powered learning systems are accessible to learners with disabilities is both a legal requirement under frameworks like the Americans with Disabilities Act and a moral imperative for inclusive education. Agentic AI offers significant potential to enhance accessibility beyond what traditional assistive technologies provide. AI powered screen readers can provide more intelligent navigation and summarization of complex content. Speech recognition systems can adapt to nonstandard speech patterns, including stuttering or speech affected by motor impairments. Natural language interfaces can reduce reliance on fine motor control for navigation and input. Augmentative and alternative communication devices incorporating AI can predict intended messages and reduce communication effort.

The 2024 report from the Consortium for School Networking and CAST highlights both promise and caution in this domain. AI has improved assistive technology performance considerably, from facial recognition reducing password burden to word prediction and voice dictation enabling interaction for individuals with limited dexterity. At the same time, students with disabilities, who stand to gain most from AI advances, are often least able to access them due to inaccessible design of AI tools themselves. The framework calls for accessibility as a foundational design requirement rather than afterthought accommodation. Agentic AI systems should comply with Web Content Accessibility Guidelines from the outset, ensuring keyboard navigability, screen reader compatibility, adequate colour contrast, and other accessibility standards. Beyond compliance, systems should proactively adapt to disability related needs, offering simplified interfaces for learners with cognitive disabilities, extended time and reduced distraction for learners with attention differences, and flexible input modalities for learners with motor impairments.

Critically, agentic AI should support cognitive accessibility alongside sensory and motor accessibility. Many learners experience challenges with executive function, working memory, processing speed, or attention. AI can support these learners through task decomposition, strategic reminders, reduced cognitive load in interface design, and scaffolding for self-regulation. The Vanderbilt University Planning Assistant project exemplifies this approach, scanning syllabi to extract key dates and eventually helping students divide complex assignments into subtasks and suggested timelines.

7. Continuous Learner Modeling

Effective personalization depends on accurate, continuously updated models of learner knowledge, skills, and states. Traditional knowledge tracing approaches estimate the probability of mastery for discrete skills based on performance history. While valuable, such models capture only a slice of the relevant learner state. Agentic AI enables richer, more holistic learner modelling. Beyond knowledge state, agentic systems can model affective state, including engagement, frustration, confusion, and boredom. They can model metacognitive state, including self-awareness of understanding, help seeking behaviour, and self-regulation strategies. They can model motivational orientation, including goal types, interest patterns, and persistence tendencies. And they can model contextual factors, including time of day, environmental distractions, and recent learning history.

Advances in machine learning enable increasingly sophisticated learner modelling. Deep knowledge tracing employs recurrent neural networks to capture complex temporal patterns in learning sequences (Piech et al., 2015). Attention mechanisms allow models to weight relevant prior interactions differentially. Memory augmented networks maintain persistent representations across sessions. Multimodal sensing can incorporate behavioural indicators beyond response data, including gaze patterns, facial expressions, and interaction timing. Continuous learner modelling serves multiple functions within agentic AI. It enables just in time adaptation, with systems responding to emerging affective or cognitive states before learners disengage. It enables long term trajectory planning, with systems projecting learning progressions and identifying optimal pacing. It enables personalized explanation, with systems adapting communication style and complexity to modelled preferences and capabilities. And it enables early intervention, with systems detecting learners at risk of falling behind and alerting human educators. Privacy considerations loom large in comprehensive learner modelling. The detailed behavioural data required for rich modelling raises concerns about surveillance, consent, and potential misuse. The framework addresses these concerns in the guardrails section, emphasizing data minimization, purpose limitation, and learner control over personal information.

8. Autonomous Feedback and Assessment

Feedback is among the most powerful influences on learning, but its effectiveness depends critically on quality and timing (Hattie & Timperley, 2007). Effective feedback is specific, actionable, and oriented toward improvement. It addresses not merely whether answers are correct but why errors occurred and how performance can improve. Traditional automated feedback systems have struggled to provide such quality, typically offering template-based responses that lack the nuance of skilled human feedback. Agentic AI transforms feedback possibilities through natural language generation and reasoning capabilities. Rather than selecting among predetermined feedback templates, agentic systems can generate feedback tailored to specific errors, learner history, and instructional context. They can explain misconceptions in multiple ways, offer analogies connected to learner interests, and suggest specific strategies for improvement.

Research on intelligent tutoring systems confirms the importance of immediate, specific feedback for learning outcomes. Systems providing such feedback consistently outperform delayed or generic feedback conditions. Agentic AI can provide feedback immediately upon response, while the problem-solving process remains fresh in working memory. This immediacy enables learners to connect feedback to specific reasoning steps and adjust subsequent attempts accordingly. Beyond feedback on individual responses, agentic AI can provide formative assessment that tracks progress toward learning objectives over time. Competency mapping techniques represent learner capabilities across multidimensional skill spaces, identifying areas of strength and areas requiring additional attention. Such mapping can guide both learner self-assessment and system recommendations, making learning progress visible and actionable.

For inclusive education, feedback and assessment must accommodate diverse ways of demonstrating understanding. Some learners express themselves best in writing, others orally, still others through creation of artifacts or performances. Agentic AI can assess understanding through multiple modalities, evaluating spoken explanations, written responses, constructed diagrams, or problem-solving processes. This flexibility aligns with UDL principles of multiple means of action and expression, ensuring that assessment measures understanding rather than incidental skills.

9. Cultural and Linguistic Inclusivity

Educational AI systems trained predominantly on data from dominant cultural and linguistic groups risk perpetuating biases and underserving diverse learners. Research has documented systematic disparities in AI performance across demographic groups, with automated essay scoring systems rating native speakers of Arabic, Hindi, and Spanish lower than other groups relative to human ratings (Baker & Hawn, 2021). Such disparities undermine educational equity and erode trust in AI systems. The framework calls for cultural and linguistic inclusivity as explicit design objectives. Agentic AI should be capable of adapting to diverse cultural contexts, drawing on culturally relevant examples, respecting diverse knowledge traditions, and avoiding assumptions rooted in particular cultural perspectives. This requires training on diverse data, evaluation across demographic groups, and ongoing monitoring for differential performance.

Linguistic inclusivity encompasses both multilingual capability and adaptation to linguistic diversity within languages. Agentic AI should support instruction in multiple languages, enabling learners to engage in their strongest language while developing proficiency in additional languages. For English language learners, systems should provide linguistic scaffolding including vocabulary support, simplified syntax options, and translation assistance while maintaining conceptual rigor. Beyond language, cultural inclusivity involves sensitivity to diverse learning traditions, communication norms, and knowledge systems. Collectivist cultures may emphasize collaborative learning and indirect communication styles. Indigenous knowledge systems may organize understanding differently than Western academic traditions. Agentic AI should not impose particular cultural assumptions but rather adapt to learner cultural contexts while bridging to academic conventions when appropriate.

Bias mitigation requires systematic attention throughout the AI development and deployment lifecycle. Training data should represent diverse populations. Evaluation should assess performance across demographic groups. Algorithmic auditing should identify and address disparate outcomes. The 2024 CoSN and CAST report emphasizes that personnel involved in AI development and review should reflect the diversity of end users, and ongoing community feedback should inform system refinement.

10. Teacher Augmentation and Reduced Cognitive Load

The framework positions agentic AI as augmenting rather than replacing human educators. Teachers bring irreplaceable capabilities including emotional atonement, ethical judgment, relationship building, and adaptive creativity that current AI cannot replicate. The goal is not to eliminate human roles but to amplify human effectiveness by offloading routine tasks and providing actionable insights. Teacher cognitive load has intensified in contemporary educational contexts. Demands for differentiated instruction, data driven decision making, family communication, and administrative documentation compete for limited attention. Agentic AI can reduce this burden by automating routine tasks, synthesizing student data into actionable summaries, and flagging students requiring human attention. Dashboard design represents a critical interface between agentic AI and human educators. Effective dashboards provide at a glance understanding of class and individual progress without overwhelming with data. They highlight anomalies and concerns while maintaining context. They suggest potential interventions while respecting teacher professional judgment. Research on learning analytics dashboards emphasizes the importance of actionable information presentation, avoiding data overload while ensuring relevant information reaches educators when needed.

Alert prioritization represents another key function. Agentic AI can monitor all students simultaneously, detecting struggling learners, disengagement patterns, or potential emotional distress. Rather than presenting undifferentiated alerts, intelligent prioritization can highlight the students and concerns most requiring immediate human attention. This enables teachers to allocate their limited time to highest impact interactions. The human in the loop paradigm ensures that critical educational decisions remain with human educators. While agentic AI can recommend interventions, provide information, and execute routine actions, consequential decisions about student placement, grading, or intervention should involve human judgment. This preserves accountability, enables consideration of contextual factors AI may miss, and maintains the relational dimension of education that matters deeply to students and families.

11. Motivational and Behavioral Support

Motivation is the engine of learning. Without motivation, even the most sophisticated instructional approaches falter. Self Determination Theory provides the primary theoretical lens for understanding motivation in the framework, emphasizing the importance of autonomy, competence, and relatedness for fostering intrinsic motivation and autonomous regulation. Agentic AI can support motivation through multiple mechanisms. By offering meaningful choices about learning content, sequence, and expression modes, systems support autonomy. By calibrating challenge to capability and providing relevant feedback, systems support competence. By conveying understanding, maintaining warm tone, and facilitating connection with peers and teachers, systems support relatedness.

Engagement prediction represents a proactive approach to motivational support. Drawing on patterns in behavioural data, agentic AI can estimate engagement levels and predict disengagement before it manifests overtly. Early detection enables pre-emptive intervention, whether through adjusting task difficulty, introducing novel content, offering choice, or alerting human educators to check in with the learner. Gamification has shown promise for enhancing engagement, though effects depend on implementation quality and learner characteristics. Agentic AI can personalize gamification elements to individual preferences, emphasizing challenge and mastery for some learners, social connection and collaboration for others, narrative and exploration for still others. This calibrated approach avoids the pitfalls of uniform gamification, which may appeal to some learners while disengaging others.

Persistence support addresses the reality that meaningful learning involves struggle and setback. Learners who interpret difficulty as indicating inability may disengage prematurely. Agentic AI can provide targeted support during challenging periods, normalizing struggle, emphasizing growth and improvement, and helping learners maintain perspective on their learning trajectories. Such support may be particularly important for learners from backgrounds where academic struggle has historically been interpreted through deficit lenses.

12. Ethical Guardrails and Responsible Autonomy

The very capabilities that make agentic AI powerful also raise ethical concerns. Autonomous systems making consequential decisions about educational trajectories, assessment, and intervention require robust safeguards against harm. The framework addresses these concerns through comprehensive ethical guardrails. Transparency represents a foundational requirement. Learners, educators, and families should understand how AI systems operate, what data they collect, how decisions are made, and what options exist for human override. This does not require exposing technical implementation details but does require clear communication about system behaviour and limitations. The European Union AI Act classifies educational AI as high risk, requiring transparency and human oversight provisions.

Human oversight protocols ensure that humans retain meaningful control over consequential decisions. While agentic AI may operate autonomously for routine instructional interactions, decisions with significant educational consequences should involve human review. This includes assessment decisions affecting placement or progression, interventions with potential stigmatizing effects, and data sharing beyond immediate instructional purposes. Human oversight also enables error correction when AI systems malfunction or produce harmful outputs.

Algorithmic accountability addresses the need to identify and remedy biased or harmful system behaviour. Regular auditing should assess system performance across demographic groups, identifying disparities that may reflect underlying bias. When disparities are identified, remediation should be prompt and documented. Clear procedures should exist for stakeholders to report concerns and for systems to be modified in response. Data privacy protections address the sensitive nature of educational data, particularly the detailed behavioural data that comprehensive learner modelling requires. Data minimization principles call for collecting only data necessary for educational purposes. Purpose limitation prevents data from being repurposed without consent. Learner and family control ensures meaningful agency over personal information, including rights to access, correct, and delete data. Compliance with legal frameworks including FERPA in the United States provides baseline protections, though the framework calls for protections exceeding legal minimums.

Bias mitigation addresses the documented reality that AI systems can perpetuate and amplify existing inequities. Training data should represent diverse populations and contexts. Fairness metrics should be incorporated into system evaluation, with performance assessed across demographic groups. Fairness aware algorithms should be employed where disparate outcomes are detected. Ongoing monitoring should continue throughout deployment, as bias can emerge from changes in user populations or system drift over time.

Table 2: Ten Pillar Taxonomy of Agentic AI Capabilities for Inclusive Learning

Pillar	Core Function	Key Techniques	Inclusivity Dimension
Personalized Pathways	Curriculum sequencing	Reinforcement learning, knowledge graphs	Individualized pacing and goals
Dynamic Scaffolding	Just in time support	ZPD modelling, fading algorithms	Cognitive accessibility
Multimodal Experience	Representation flexibility	Cross modal generation, UDL	Sensory inclusivity
Accessibility	Barrier removal	Screen reader optimization, simplification	Disability inclusion
Learner Modelling	State inference	Deep knowledge tracing, affect detection	Holistic understanding
Autonomous Assessment	Competency measurement	Formative analytics, NLG feedback	Fair evaluation
Cultural Inclusivity	Context adaptation	Multilingual NLP, bias auditing	Linguistic and cultural equity
Teacher Augmentation	Educator support	Dashboard design, alert prioritization	Sustainable implementation
Motivational Support	Engagement maintenance	SDT alignment, gamification	Emotional wellbeing
Ethical Guardrails	Responsible operation	Transparency, human override	Trust and accountability

13. Integrated System Architecture

Realizing the framework requires technical architecture that integrates the ten pillars into coherent system operation. Drawing on current agentic AI design patterns, we propose a multi-layer architecture comprising data infrastructure, learner modelling, agent orchestration, interaction management, and governance layers. The data infrastructure layer manages collection, storage, and access to learner interaction data. This includes event logging capturing learner actions and responses, secure data storage with appropriate access controls, and APIs enabling other layers to query learner history. Privacy preserving techniques including differential privacy and federated learning can reduce risks associated with centralized data collection. The learner modelling layer maintains continuously updated representations of each learner. This includes knowledge state models estimating mastery of concepts and skills, affective state models estimating engagement, confusion, and frustration, and preference models capturing learning style tendencies and motivational orientations. Models draw on multimodal input including response patterns, timing information, and where available, physiological or behavioural indicators.

The agent orchestration layer coordinates specialized agents responsible for different framework pillars. A pathway agent manages curriculum sequencing and objective negotiation. A scaffolding agent monitors for struggle and deploys appropriate support. A content agent generates or selects multimodal instructional materials. An assessment agent provides feedback and tracks competency development. A motivation agent monitors engagement and adjusts interaction dynamics. Orchestration ensures coherent operation, resolving potential conflicts between agent recommendations and maintaining consistent learner experience.

The interaction management layer handles learner facing presentation and input processing. This includes natural language understanding for processing learner queries and responses, natural language generation for producing conversational output, multimodal rendering for presenting content across modalities, and accessibility services ensuring interface accessibility for learners with disabilities. The layer abstracts interaction complexity, enabling agents to operate in terms of pedagogical objectives while the layer handles communication mechanics. The governance layer implements ethical guardrails including transparency mechanisms, human oversight protocols, bias monitoring, and privacy protection. This layer logs system decisions for accountability, enforces human review requirements for consequential actions, monitors for disparate outcomes across demographic groups, and ensures compliance with data protection requirements. The governance layer operates continuously rather than as occasional audit, embedding ethical operation into routine system function.

Table 3: Implementation Challenge Matrix by Stakeholder Group

Challenge	Learners	Educators	Institutions	Developers	Priority
Data Privacy	High	Medium	High	High	Critical
Algorithmic Bias	High	Medium	Medium	High	Critical
Technical Infrastructure	Low	Medium	High	Medium	High
Educator Training	Low	High	Medium	Low	High
Cost and Sustainability	Low	Low	High	Medium	High
Trust and Adoption	Medium	High	Medium	Low	Medium
Accessibility Compliance	High	Medium	High	High	Critical

14. Implementation Challenges and Mitigation Strategies

14.1. Data Privacy and Security

The detailed behavioural data that enables comprehensive learner modelling also poses privacy risks. Learner data can reveal sensitive information about cognitive abilities, emotional states, and personal circumstances. Data breaches could expose vulnerable populations to harm. Aggregated data could be misused for purposes beyond educational support. Mitigation strategies include privacy by design approaches that minimize data collection, anonymize where possible, and enforce strict access controls. Federated learning techniques can train models on distributed data without centralizing sensitive information. Clear data governance policies should specify collection purposes, retention limits, and deletion procedures. Transparency with learners and families should ensure informed consent and meaningful control over personal information.

14.2. Algorithmic Bias and Fairness

AI systems trained on historical data may perpetuate patterns of educational inequity. If past outcomes reflect systemic disadvantage, models trained on these outcomes may reproduce rather than remedy disparities. Biased training data, biased algorithmic assumptions, and biased evaluation metrics can all contribute to unfair system behaviour. Mitigation requires systematic attention throughout the development lifecycle. Diverse training data should represent the full range of intended users. Fairness metrics should be incorporated into model evaluation, assessing performance disaggregated by demographic group. Algorithmic auditing should identify disparate outcomes before and during deployment. Remediation procedures should enable prompt response when bias is detected. Ongoing monitoring should continue throughout system operation, as bias can emerge from changing user populations or system drift.

14.3. Educator Preparation and Support

Effective AI integration requires educators who understand system capabilities and limitations, can interpret AI generated insights, and can maintain meaningful human oversight. Many educators currently lack preparation in these areas. Without adequate support, educators may either reject AI tools or defer inappropriately to AI recommendations. Mitigation involves sustained professional development that builds AI literacy alongside pedagogical skill. Training should address both practical system operation and critical evaluation of AI recommendations. Ongoing support structures should enable educators to share experiences and problem solve collectively. System design should make AI reasoning transparent and support educator judgment rather than supplanting it.

14.4. Infrastructure and Sustainability

Agentic AI systems require substantial computational infrastructure including processing power for model inference, storage for learner data, and reliable network connectivity. Many educational institutions, particularly in under resourced settings, lack such infrastructure. Even where infrastructure exists, ongoing costs of AI operation may exceed available budgets. Mitigation strategies include cloud-based deployment that reduces local infrastructure requirements, though this raises data sovereignty concerns. Open-source platforms can reduce licensing costs. Partnerships between institutions can share infrastructure investments. Policy advocacy can seek public investment in educational AI infrastructure as essential public good. Design decisions should consider resource constraints, offering degraded functionality rather than complete failure when resources are limited.

14.5. Limitations and Future Research Directions

This framework represents a theoretical contribution requiring empirical validation. While individual components draw on established research, the integrated framework has not been tested in its entirety. Future research should examine implementation of the framework in diverse educational contexts, assessing effects on learning outcomes, engagement, and equity across learner populations. The framework focuses primarily on individual learner experience, giving less attention to collaborative and social dimensions of learning. Peer interaction, group projects, and classroom community are important elements of education that agentic AI should support rather than displace. Future work should elaborate how agentic AI can enhance rather than undermine social learning.

Rapid evolution of AI capabilities means that specific technical recommendations may quickly become outdated. The framework attempts to articulate principles that transcend particular technological implementations, but ongoing revision will be necessary as capabilities advance. Particular attention should be paid to developments in multimodal AI, embodied agents, and human AI collaboration. Questions of appropriate AI autonomy remain contested. The framework advocates human oversight for consequential decisions, but boundaries between routine and consequential are not always clear. Future research should examine how stakeholders, including learners, understand and prefer AI autonomy to be bounded, and how these preferences vary across cultures and contexts.

15. Conclusion

Agentic AI represents a transformative opportunity for inclusive education. The capacity for autonomous, goal directed behaviour enables personalization at scales impossible through human instruction alone. When grounded in sound learning theory and designed with ethical guardrails, agentic AI can advance educational equity by meeting each learner where they are and supporting progress toward individually meaningful goals. The framework presented here organizes this opportunity around ten pillars spanning the full scope of inclusive educational practice. From personalized learning pathways to ethical guardrails, each pillar addresses essential requirements for agentic AI to serve diverse learner populations. The integration of Self Determination Theory and Universal Design for Learning provides theoretical grounding that connects technical capabilities to established understanding of human learning and motivation.

Realizing this opportunity requires navigating significant challenges. Data privacy, algorithmic bias, educator preparation, and infrastructure sustainability all demand systematic attention. The implementation challenge matrix highlights how these challenges affect different stakeholders, enabling targeted mitigation. Ongoing vigilance will be necessary as AI capabilities and deployment contexts evolve. The ultimate measure of agentic AI in education will be its contribution to human flourishing. Technology serves education only insofar as it supports learners in developing capabilities that enable them to pursue lives, they value. This requires not merely cognitive gains but development of agency, self-determination, and connection to learning communities. The framework calls for agentic AI that supports learners in becoming agentic themselves, equipped with knowledge, skills, and dispositions for lifelong learning in a rapidly changing world.

We invite researchers, practitioners, and policymakers to engage critically with this framework. Empirical research should test its propositions in diverse contexts. Practitioners should adapt its guidance to local circumstances. Policymakers should consider how regulation can support beneficial development while guarding against harm. Through collaborative effort across these communities, agentic AI can become a genuine force for educational inclusion and equity.

References

[1] Acharya, A., Smith, J., Chen, L., & Williams, R. (2025). Agentic AI systems: Architecture and applications. Journal of Artificial Intelligence Research, 74, 18901-18945.
[2] Anderson, J. R., Boyle, C. F., & Reiser, B. J. (1985). Intelligent tutoring systems. Science, 228(4698), 456-462.
[3] Baker, R. S., & Hawn, A. (2021). Algorithmic bias in education. International Journal of Artificial Intelligence in Education, 32(4), 1052-1092. https://doi.org/10.1007/s40593-021-00285-9
[4] Batsaikhan, O., & Correia, A. P. (2024). The effects of generative artificial intelligence on intelligent tutoring systems in higher education. STEM Education and Learning, 4(1).
[5] Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. Educational Researcher, 13(6), 4-16.
[6] Carbonell, J. R. (1970). AI in CAI: An artificial intelligence approach to computer-assisted instruction. IEEE Transactions on Man-Machine Systems, 11(4), 190-202.
[7] CAST. (2024). Universal design for learning guidelines version 3.0. https://udlguidelines.cast.org/
[8] Corbett, A. T., & Anderson, J. R. (1994). Knowledge tracing: Modelling the acquisition of procedural knowledge. User Modelling and User-Adapted Interaction, 4(4), 253-278.
[9] Deci, E. L., & Ryan, R. M. (1985). Intrinsic motivation and self-determination in human behaviour. Plenum Press.
[10] Deci, E. L., & Ryan, R. M. (2000). The 'what' and 'why' of goal pursuits: Human needs and the self-determination of behaviour. Psychological Inquiry, 11(4), 227-268.
[11] Hattie, J., & Timperley, H. (2007). The power of feedback. Review of Educational Research, 77(1), 81-112.
[12] Hosseini, M., & Seilani, H. (2025). Agentic AI in educational contexts: From copilot to autopilot. Educational Technology Research and Development, 73(2), 445-472.

[13] Karran, J. A., Boasen, J., & Léger, P. M. (2025). A systematic review of AI-driven intelligent tutoring systems in K-12 education. npj Science of Learning, 10, 29. https://doi.org/10.1038/s41539-025-00320-7

[14] 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), 901-918.

[15] Neshaei, S. P., Davis, R. L., Hazimeh, A., Lazarevski, B., Dillenbourg, P., & Käser, T. (2024). Towards modelling learner performance with large language models. Proceedings of the Educational Data Mining Conference 2024.

[16] Niemiec, C. P., & Ryan, R. M. (2009). Autonomy, competence, and relatedness in the classroom: Applying self-determination theory to educational practice. Theory and Research in Education, 7(2), 133-144.

[17] Pérez Perez, F. (2024). AI and accessibility in education. Consortium for School Networking & CAST.

[18] Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L., & Sohl-Dickstein, J. (2015). Deep knowledge tracing. Advances in Neural Information Processing Systems, 28, 505-513.

[19] Rose, D. H., & Meyer, A. (2002). Teaching every student in the digital age: Universal design for learning. ASCD.

[20] Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. American Psychologist, 55(1), 68-78.

[21] Ryan, R. M., & Deci, E. L. (2017). Self-determination theory: Basic psychological needs in motivation, development, and wellness. Guilford Press.

[22] Saborío-Taylor, S., & Rojas-Ramírez, F. (2024). Universal design for learning and artificial intelligence in the digital era: Fostering inclusion and autonomous learning. International Journal of Professional Development, Learners and Learning, 6(2), ep2408.

[23] Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. Harvard University Press.

[24] Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. Journal of Child Psychology and Psychiatry, 17(2), 89-100.

[25] World Health Organization. (2011). World report on disability. WHO Press.