

Agentic AI for Learning

Sukrit Sharma

Artificial Intelligence (AI) has reshaped the educational landscape with personalized tutors and adaptive platforms where agentic AI is the latest paradigm. Automated grading, and content recommendation engines have already taken their place in classrooms, lecture halls, and online learning platforms across the globe. At the centre of this revolution has been the rise of large language models (LLMs), most prominently GPT systems, which have enabled personalized support, tutoring, and scalable content generation.

Yet, even in their current form, these AI systems are still constrained by the stagnation of their design. They often lack nuanced reasoning, depend heavily on static data, and fail when required to adapt dynamically to real-world educational contexts. This expresses a dire need for a more sustainable and innovative framework, one that builds upon the pre-existing foundations of generative AI but extends beyond its current limitations. A radical shift in pedagogy to accommodate the newer complexities of educational framework.

Here comes a newer paradigm, Agentic AI. While it draws on the power of LLMs, Agentic AI is not just a better chatbot. It is fundamentally different from generally perceived AI. It is a dynamic orchestra of LLMs. In this paradigm, multiple LLM-powered “agents” interact, reason, and collaborate with each other, pursuing goals autonomously rather than waiting for predefined instructions. Each agent can be programmed with a task, assist other agents, and/or adapt dynamically to new conditions. Unlike static prompt-response systems, agentic workflows involve multiple-step reasoning where outputs are reviewed, refined, and contextualized by other agents before being presented to the learner or educator.

This goal-oriented, autonomous decision-making allows Agentic AI to move beyond simple input-output mappings and toward workflows that are dynamic and context-aware. In practice, this means an educational task – say, helping a student struggling with fractions – would not be solved by a single LLM output. Instead, one agent might diagnose the gap in knowledge, another might select teaching strategies, a third might retrieve appropriate materials, and yet another might monitor engagement and comprehension. Together, they produce a richer, more adaptive, and pedagogically aligned intervention.




Figure 1: A simplified workflow of an agentic AI tutoring system

Generative AI can thus be seen as the precursor, the raw material upon which Agentic AI builds. Agentic AI abstracts away direct prompting from the user, this allows agents themselves to generate, refine, and direct inputs. This shift creates workflows where the learner interacts not with a single

model but with a collaborative system of reasoning agents. To conclude, education no longer has just one “AI tutor” – it has a team of AI collaborators.

Despite its potential boons, the adoption of Agentic AI in large-scale educational systems remains at a very early stage, most modern educational systems at scale have not yet adopted an agentic framework even though most employ LLMs and Gen AI such as embibe. Current initiatives – like Agent4EDU or intelligent tutoring systems (ITS) – offer glimpses of what agent-based education might look like. They show that multi-agent systems (MAS) can coordinate learning experiences and augment teachers, but their implementations remain far from industry-grade solutions.

Several shortcomings stand out. Early agentic systems often lack seamless tool integration, which means they cannot draw dynamically from APIs, institutional databases, or real-time learning analytics personalized to user data. They also struggle with personalization at depth. While many claim to offer tailored content, few systems truly analyze past performance, detect gaps, and craft individualized learning trajectories at scale. Moreover, as Yusuf et al. (2025) argue, these efforts often treat personalization as content delivery rather than holistic student support – including emotional well-being, engagement, and pacing. Simply put, tool integration employed at its full implementation is absolute for understanding and improving student performance in all domains, it holds immense power to pinpoint each nuance detail of the user/students past mistakes and responses. Thus, it can change the teaching pattern in real time for a case-to-case basis.

Another flaw in frameworks like *Agent4EDU* is their underutilization of the *human-in-the-loop* principle. These systems sometimes suggest replacing, rather than augmenting, educators. This is a critical mistake. The power of Agentic AI is not in sidelining teachers but in amplifying their reach—helping them deliver personalized attention at scale without being overwhelmed. A system that overshadows the teacher’s role risks alienating the very practitioners it seeks to empower.

Yet another obstacle lies ahead, many early attempts are hampered by issues inherited from older ITS approaches: limited ICT training for teachers, poor alignment between AI-based methods and traditional pedagogy, and lack of interpretability. While agentic AI theoretically solves some of these barriers – because agents can adapt across pedagogical styles and even train other agents – the reality is that *Agent4EDU* and similar projects rarely demonstrate this flexibility in practice.

In short, while these pioneering efforts reveal the need for multi-agent systems in education, they also highlight a wide-open opportunity: to design frameworks that integrate tools seamlessly, personalize deeply, keep educators central, and address ethical-technical gaps head-on.

What makes this moment exciting is that agentic AI is no longer confined to research labs. Thanks to platforms like **CrewAI** and **n8n**, agentic frameworks are accessible to anyone – even those without programming expertise.

CrewAI provides a no-code MAS interface where agents can be configured with simple natural language instructions. Educators can, for example, assign an “explainer” agent to break down concepts, a “checker” agent to evaluate student responses, and a “coach” agent to monitor motivation – all without writing a single line of code (Perumal et al., 2024). These agents can interact across platforms, analyzing a YouTube transcript for lecture notes one moment and updating a personalized study plan through an LMS API the next.