

DOI: https://doi.org/10.48009/1_iis_2026_128

From tool to teaching partner: The rise of agentic AI in education

Tala Najem, Oglethorpe University, talanaajem@gmail.com

Leila Halawi, Embry-Riddle Aeronautical University, halawil@erau.edu

Abstract

This paper examines the emergence of agentic artificial intelligence in higher education, exploring the shift from AI as a tool to AI as an institutional actor. Using a structured narrative literature review, the study addresses four research questions: how agentic AI is defined, what autonomous behaviors current systems demonstrate, how increasing AI autonomy reshapes teaching roles and student engagement, and what governance frameworks can guide responsible integration. The analysis establishes that agentic AI is architecturally distinct from prior educational technologies in ways the field has not adequately theorized, and that existing governance structures are poorly equipped for the autonomy these systems now exercise. The paper introduces the Agentic AI Integration Continuum, a four-level framework providing institutions with a shared vocabulary for categorizing AI autonomy, calibrating risk, and aligning governance with actual deployment conditions. The paper concludes with a faculty roadmap prioritizing process-verified learning and explicitly scoped autonomy to preserve human agency in an increasingly automated academic environment.

Keywords: agentic AI, higher education, AI governance, instructional autonomy, academic integrity, AI integration continuum

Introduction

Agentic AI is advancing at a pace that higher education institutions were not designed to accommodate. Unlike previous tools, agentic AI does not wait for a prompt; it autonomously adapts tasks and executes instructional moves (Sukharevsky et al., 2025). It extends beyond providing simple recommendations to actively take instructional actions such as adapting learning tasks and initiating feedback. While organizations have thus far achieved modest productivity gains from basic chatbot implementations and limited-scale pilots, agentic AI represents a transition toward autonomous systems that do not merely suggest tasks but execute entire processes to generate measurable outcomes (Sukharevsky et al., 2025). What began as a theoretical construct is now a \$7.55 billion market, on a trajectory toward \$199 billion by 2034 (Precedence Research, 2025). The impact is already visible: industries adopting these tools are beginning to deliver the measurable ROI that pilots have missed (PwC, 2025). Higher education institutions have traditionally adopted technological change at a measured pace; however, the rapid advancement of agentic AI is creating new challenges that require timely institutional responses. The implications of agentic AI for higher education are significant because these systems increasingly influence learning, teaching, and student success while preparing graduates for an AI-enabled workforce. The field is simultaneously subject to agentic AI disruption and responsible for producing the graduates that every other disrupted industry urgently needs. By early 2026, an agentic AI system demonstrated the ability to automate an entire Canvas course from enrollment to grading, revealing that many online assessment and integrity policies were not designed to detect autonomous AI participation (Palmer, 2026). This demonstration established that AI can

now automate the entire student experience, from enrollment to the final grade. While the tool was removed following legal action, what could not be removed was what it revealed: the standard assessment architecture of most online programs lacked a mechanism to detect it, and no policy language was adequate to address it. Kofinas et al. (2025) found that even AI-knowledgeable faculty cannot consistently identify AI-generated work. Their study warns that authentic assessment, once a primary defense, is no longer a reliable safeguard. As AI evolves from support tools to autonomous systems, existing governance frameworks in higher education may be insufficient to address the challenges posed by greater AI autonomy (Pati, 2025). Characteristics of online learning environments may create conditions in which agentic AI systems can operate more effectively and independently. The very architecture of online education provides a favorable environment for these agentic systems to operate. This paper explores how agentic AI is being used in higher education, pulling insights from educational technology and learning sciences to AI and information systems. Adaptive learning platforms and intelligent tutoring systems are well established; the questions this paper addresses concern something more recent and less well theorized. As AI begins making its own instructional calls and adjusting lessons without human input, it assumes a new kind of authority (Sapkota et al., 2026). Higher education institutions are still working to determine what this means for teaching roles, degrees, and basic oversight. Instead of treating these developments as ethical obstacles to be managed, this paper regards them as structural shifts that need a new analytical framework. It adopts a forward-looking perspective and moves beyond simple ethics to explore how AI's agency transforms teaching, professional roles, and institutional governance.

The paper addresses four primary research questions: (1) What distinguishes agentic AI from existing adaptive and analytical educational technologies? (2) What specific agentic behaviors are current AI systems demonstrating in education? (3) How does increased AI autonomy reshape teaching roles, student engagement, and instructional design? (4) What governance or institutional frameworks can guide the responsible integration of agentic AI in higher education? The remainder of this paper is structured as follows. The literature review defines agentic AI, addresses all four research questions using existing scholarship, and draws the agent-versus-agentic AI distinction required by the analysis. The methodology section describes the structured narrative review approach. The results and discussion synthesize findings and introduce the Agentic AI Integration Continuum, the primary conceptual contribution. Finally, the paper concludes with implications for practice, a faculty roadmap for responsible implementation, and directions for future research.

Literature Review

The idea of agentic AI draws from long-established conversations about agency across education, psychology, and the learning sciences. Agency in education typically refers to the ability to make choices, act, and shape how learning occurs (Abou Ali et al., 2025). Historically, teachers and students have been seen as primary agents in education. At the same time, technology has mainly served as a tool, supporting or structuring their work without making decisions on its own (Baltimore et al., 2025). Educational technologies have operated in a responsive rather than directive role. Recent AI developments are changing this dynamic. Some AI systems now do more than respond to user input; they actively participate in instructional decision-making by initiating actions that directly affect learning (Sapkota et al., 2026). As a result, questions of agency are no longer limited to human actors alone. Classrooms increasingly include systems that influence instructional flow, feedback timing, and learning pathways. Despite growing attention to this shift, there is currently no universally accepted definition of agentic AI in education. Across the literature, the term is used inconsistently. Some authors describe agentic systems in technical terms, emphasizing planning and tool use (Bandi et al., 2025; Brohi et al., 2025; Madrid-Garcia et al., 2025). Others frame agentic AI more broadly as any system that demonstrates initiative or autonomy in

instructional contexts (Alqurni, 2026; Pati, 2025; Sapkota et al., 2026). In many cases, the distinction between advanced adaptive systems and truly agentic systems remains blurred. This conceptual ambiguity makes it difficult to evaluate claims about autonomy, responsibility, and educational impact. Before defining the terms used in this paper, it is necessary to distinguish between an AI agent and agentic behavior. An agent is a specific software system, while agentic AI describes a system's ability to pursue goals, handle multi-step tasks, and self-correct without a human in the loop (Acharya et al., 2025; Hosseini & Seilani, 2025). For this paper, agentic AI in education refers to artificial intelligence systems that independently initiate and carry out instructional actions within pedagogical boundaries, rather than simply offering recommendations for teachers or students to act on (Pati, 2025). This definition rests on three key criteria: the ability to initiate action without explicit prompting, the capacity to adapt decisions over time, and the requirement to operate within structured educational limits (Kostopoulos et al., 2025). By focusing on action rather than assistance, this definition distinguishes agentic AI from earlier adaptive and analytics-based systems. It provides a clear conceptual basis for analyzing its implications in educational settings.

Distinguishing Agentic AI From Existing Educational Technologies

While older adaptive systems follow pre-set rules, agentic AI moves beyond those boundaries to make its own decisions (Xiong, 2025). Analytics platforms leave the final call to teachers, while tutoring systems reach a limit once a student moves off script. Unlike these tools, agentic AI does not just flag patterns; it acts on them (Kostopoulos et al., 2025). Beyond promotional claims, the substantive shift in agentic AI is architectural. Traditional tools assist instruction, but they never start it. Agentic AI inverts this pattern by monitoring the learning environment and responding autonomously, swapping assignments and redirecting students without waiting for a command (Baltimore et al., 2025; Kostopoulos et al., 2025). The instructional decision-making authority has not been eliminated; it has been transferred from educators to machines. The distinction between traditional AI systems and agentic AI becomes clearer through the AgentAI framework developed by Piccialli et al. (2025). Piccialli et al. (2025) distinguish between systems that operate according to predefined rules in controlled environments and those that can adapt autonomously to changing conditions. The authors argue that agentic AI is defined not merely by computational intelligence but by its capacity to perceive environmental conditions, evaluate available options, and initiate actions within a continuous decision-making cycle that requires minimal human intervention. In higher education, this framework provides a useful distinction by differentiating agentic AI from other forms of educational technology rather than grouping all technologies into a single category. Tools such as analytics dashboards or traditional adaptive platforms can support teaching and personalizing content. However, they still work within fixed rules and only do what an educator sets them up to do. Agentic AI is different. It can monitor what is happening, decide when to step in, and adjust its behavior over time, all within certain educational boundaries. That distinction matters because what higher education is facing is not just a question of doing things faster or more efficiently. It is a deeper shift in how these systems are built and, more importantly, in who or what holds instructional authority.

Degrees of Agency in Educational AI Systems

Agentic AI differs because it can act independently within certain pedagogical constraints. Instead of waiting for a teacher's command or a student's request, these systems initiate responses, track progress over time, and adjust the learning experience as situations evolve (Wang & Chou, 2025). This introduces new questions about how authority, responsibility, and control function in AI-supported classrooms. Educational AI agency is not all-or-nothing; it is a spectrum. A system's place on this scale dictates how it must be governed and integrated into teaching (Sapkota et al., 2026). Some systems operate with constrained autonomy, remaining readily subject to educator override. Others function with substantial independence, making instructional determinations according to pre-established parameters with limited educator involvement (Madrid-Garcia et al., 2025). Wang and Chou (2025) frame this as a foundational question of instructional leadership, a design challenge that most institutions have not yet systematically addressed. Adopting a spectrum model rather than a binary classification provides a more accurate representation of

how these tools function in practice. MATHia, for example, can dynamically adjust its tutoring approach while still operating within the parameters established by the teacher (Carnegie Learning, 2026). Khanmigo directly intervenes in student work to provide contextual feedback as students progress through coursework (Khan Academy, 2023). ChatGPT extends this further, as iterative exchanges progressively simulate a collaborative learning relationship (Latif et al., 2023). These distinctions lay the groundwork for understanding how agentic behaviors appear in educational settings and why they matter for teaching, student engagement, and course design.

Emerging Forms of Agentic AI in Educational Systems

While fully autonomous AI systems are still rare in formal education, recent research shows a growing number of educational technologies with agentic characteristics (Baltimore et al., 2025; Kostopoulos et al., 2025; Pati, 2025). These systems can initiate instructional actions, adjust learning experiences over time, and operate with varying degrees of independence (Sapkota et al., 2026). This shift represents a gradual movement toward greater autonomy in learning environments, though these systems still work within pedagogical goals and institutional boundaries (Kostopoulos et al., 2025). The literature identifies three main areas in which agentic behavior most often occurs. The first involves systems that deliver instructional support without waiting for a prompt, including instructional interventions, autonomous adjustments to learning pathways, and AI tools that support instructors by acting on their behalf (Kostopoulos et al., 2025; Madrid-Garcia et al., 2025; Sapkota et al., 2026). The second area involves dynamic pathway adjustment. Some platforms now make autonomous determinations about instructional content and timing, a shift Pati (2025) argues the field has not fully reckoned with. These platforms do not simply adjust difficulty levels; they dynamically restructure instructional content. By examining weeks of data rather than a single performance, they can redirect students to foundational content or introduce alternative learning resources entirely (Tadimalla, 2025). The third area involves systems that act in the instructor's role directly: grading assignments, managing course resources, tracking progress, and reducing administrative overhead that consumes faculty time that might otherwise go to mentoring and substantive feedback (Kostopoulos et al., 2025). Together, these developments suggest that educational AI is taking on a more active role in shaping how students learn. Table 1 provides representative examples of agentic AI system categories along with educational tools and platforms currently used in higher education and digital learning environments.

Table 1. Agentic AI Systems in Education

Agentic System Type	Core Agentic Behavior	Educational Function	Example Tools / Platforms
Autonomous Tutoring Agents	Adjust lesson difficulty in real time; track engagement using behavioral and cognitive data.	Modify instructional approach without teacher prompting (Pati, 2025).	MATHia, Khanmigo, Squirrel AI Tutor (Carnegie Learning, 2026; Khan Academy, 2023; Squirrel AI Learning, 2026).
AI Grading Agents	Independently grade assignments and student responses.	Provide immediate, consistent feedback; reduce grading time (Kostopoulos et al., 2025).	Gradescope, Turnitin AI grading tools, ChatGPT-assisted rubric grading (Whippy, 2026).
Learning Management Agents	Manage resources, coordinate assessments, and track course progress.	Reduce administrative and cognitive load on teachers (Patel et al., 2025).	Canvas LMS analytics agents, Blackboard AI assistant tools, Coursera adaptive learning systems (Instructure, 2026).

Student Support and Coaching Agents	Provide emotional support, monitor well-being, and manage scheduling.	Support mental health initiatives and student wellbeing (Mendoza-Pinto, 2025).	Mainstay, Ivy.ai campus assistant, Replika-style coaching agents for wellbeing (Ivy.ai, 2026; Mainstay, 2024; Mendoza-Pinto, 2025).
-------------------------------------	---	--	---

Several of the tools listed in Table 1 are already present in educational settings and offer concrete examples of how agentic capabilities work in practice. MATHia, for instance, provides adaptive mathematics instruction by analyzing how students approach problems and adjusting feedback and learning pathways based on what it observes (Carnegie Learning, 2026). Platforms such as Khanmigo and Squirrel AI serve as tutoring assistants that guide students through learning tasks and provide contextual explanations as they progress through coursework (Khan Academy, 2023; Squirrel AI Learning, 2026). In assessment, tools such as Gradescope support automated grouping and evaluation of student responses, reducing the time educators spend on routine grading tasks (Whippy, 2026). One of the most recent real-world examples of agentic behavior is Georgia State University's chatbot, Pounce. Pounce does not wait for a request; it uses behavioral triggers to intervene and offer support before a student may recognize they are at risk (Meyer et al., 2024). A randomized trial indicated that this proactive approach significantly improved grades and reduced course withdrawals. This case is important because it shows that semi-autonomous systems can be scaled responsibly across an entire institution. It provides one of the strongest empirical examples currently available for examining governance approaches to agentic AI, as these systems assume increasingly active roles in supporting student success. Learning management systems increasingly include analytics dashboards that track student engagement and course activity, giving instructors a way to identify students who may be falling behind and intervene before small difficulties become larger ones (Instructure, 2026). Beyond the classroom, conversational agents such as Mainstay and Ivy.ai handle advising and student services by responding to common questions and managing routine communications across institutional systems (Ivy.ai, 2026; Mainstay, 2024). These examples show that agentic AI is already shaping how instruction is delivered, how assessment is managed, and how students are supported within higher education. They also make clear that not all educational AI systems operate with the same degree of independence. Agentic capabilities vary considerably across tools and contexts, underscoring the importance of understanding them as a spectrum rather than neatly fitting into a single category.

Implications of AI Autonomy for Teaching Roles and Instructional Design

As agentic AI becomes more prevalent in educational settings, researchers are paying closer attention to changes in teaching, student engagement, and course design (Alqurni, 2026; Sapkota et al., 2026). Unlike earlier educational technologies that largely supported human decision-making, agentic systems can autonomously initiate instructional actions, delivering feedback, modifying learning tasks, or reorganizing a student's learning pathway based on what the system observes in real time (Kostopoulos et al., 2025; Sapkota et al., 2026). This represents a genuine shift in how instructional authority works. Agency is no longer held exclusively by educators; it is distributed across people and systems, and AI moves from a supporting role to one in which it actively shapes when and how instructional decisions are made. This redistribution of agency has real consequences for what teaching involves. Rather than initiating every instructional intervention directly, educators increasingly find themselves designing, supervising, and evaluating the AI-supported environments in which their students learn (Mah, 2025). Faculty responsibilities shift toward defining the pedagogical boundaries within which these systems operate, determining the appropriate level of autonomy in a given context, and ensuring that what AI does aligns with course goals and academic standards. This requires a form of professional knowledge that goes beyond subject matter expertise, one that involves understanding how agentic systems work, making sense of their outputs, and knowing when to step in and override them. Teaching, in this sense, becomes less about direct

content delivery and more about orchestrating and overseeing a learning environment that includes AI as an active participant (Baltimore et al., 2025).

Agentic AI also changes how courses need to be designed. Traditional course structures tend to be linear and fixed, with the instructional sequence largely determined before the course begins. Agentic systems introduce the possibility of learning environments that continuously adapt based on real-time data about student performance and engagement (Tadimalla, 2025). For instructional designers, this means moving beyond static planning and building courses that can accommodate variability and system responsiveness. Learning activities need to be structured in ways that allow AI to intervene usefully without undermining the instructor's intended goals, which emphasizes modular design, clearly articulated learning objectives, and transparent boundaries on what the system is and is not permitted to do (Alqurni, 2026). From the student perspective, agentic AI offers several potential benefits that extend beyond general personalization. First, these systems enable timely, continuous feedback, reducing the delay between a student making an error and receiving guidance, which has been shown to support more effective learning, correction, and retention (Pop et al., 2025). Second, agentic AI supports adaptive learning pathways, allowing students to progress at different paces based on demonstrated understanding rather than fixed course sequences (Tadimalla, 2025). Third, AI-driven systems can provide targeted scaffolding and just-in-time support, particularly for students who may otherwise struggle to seek help independently, improving persistence and engagement (Brohi et al., 2025). Finally, these systems can enhance learning accessibility by offering alternative explanations, formats, and learning routes tailored to individual needs, which is particularly valuable in diverse classroom settings (Kostopoulos et al., 2025). At the same time, researchers have identified important concerns regarding the increasing use of agentic AI in educational settings. When AI systems consistently anticipate errors and provide immediate guidance, students may have fewer opportunities to work through difficulty on their own, reflect on their mistakes, or develop independent problem-solving skills (Alqurni, 2026; Patel et al., 2025). Over time, consistent reliance on automated support could gradually reduce the very agency that education is meant to develop. These developments suggest that the impact of agentic AI on education is shaped less by what technology can do and more by how autonomy is designed and constrained within specific instructional contexts. Getting this right requires balancing automation with meaningful opportunities for human judgment, ensuring that AI supports rather than replaces the cognitive and pedagogical processes that sit at the heart of learning. Integrating agentic AI responsibly is not just about adopting new tools. It requires a genuine rethinking of teaching roles, how courses are built, and what conditions make learning possible (Kostopoulos et al., 2025; D. Patel, 2025).

Institutional Governance and Frameworks for Agentic AI

Higher education governance structures were not built with autonomous AI systems in mind. Most existing policies assume that technology supports human decision-making rather than acting independently on behalf of educators or institutions. Agentic AI challenges this assumption directly by introducing systems that can initiate and carry out instructional or administrative actions within defined parameters (Pati, 2025). As a result, traditional governance approaches are struggling to keep pace, and a gap is opening between what institutional policy assumes and what these technologies can do. This gap is not just a technical problem. It raises fundamental questions about who is responsible when an automated system influences an educational outcome, and how accountability is maintained when human oversight is no longer constant. As AI systems take on more active roles in adjusting learning pathways, generating feedback, and participating in evaluation, institutions need to think carefully about how responsibility is assigned when automated decisions shape student outcomes (D. Patel, 2025). Transparency is becoming especially important. When students and faculty do not fully understand how or why an AI system has acted in a particular way, trust erodes and accountability becomes difficult to establish. Without clear governance structures, institutions risk creating environments where consequential instructional decisions are made without adequate visibility (Alqurni, 2026). In higher education, where academic integrity and institutional trust are foundational, this is a serious concern. Several governance frameworks have emerged in response

to these challenges. The ETHICAL Principles Framework developed at California State University, Fullerton is one prominent example, offering guiding principles organized around transparency, human oversight, accessibility, integrity, and legal compliance (Wynants et al., 2025). Since 2024, institutions such as Harvard University and the Massachusetts Institute of Technology have moved away from blanket AI bans toward tiered systems. This shift allows individual instructors to decide how students use the technology in their specific classroom (Thesify, 2025). By prioritizing student disclosure over total prohibition, these institutions acknowledge that banning a tool students already use can foster a culture of distrust. Frameworks of this kind provide a useful foundation for keeping AI systems aligned with institutional values and ensuring that humans retain meaningful control over high-stakes decisions. That said, principles alone are not enough. Institutions also need operational models that translate those principles into concrete practice (Baltimore et al., 2025). Design frameworks such as the Agentic Product Canvas offer a more applied complement to governance principles. Rather than treating agentic AI as a general-purpose tool, this kind of framework helps institutions define what a system is meant to do, how much autonomy it should have, and under what conditions it should defer to a human decision maker (Wynants et al., 2025). By making these boundaries explicit, institutions are better positioned to manage risk while still leveraging automation. An integrated governance strategy brings both elements together. This means defining different levels of autonomy for different AI applications, maintaining human oversight where the stakes are highest, ensuring that system behavior remains aligned with pedagogical goals, and building ongoing evaluation processes to monitor how systems perform over time (Gao, 2025). This approach allows institutions to move from reacting to problems after they arise to managing agentic AI in a more deliberate and structured way. Governing these systems well requires treating them not simply as tools, but as active participants in educational processes that must be carefully managed, consistently monitored, and aligned with the broader mission of higher education.

Methodology

This conceptual paper does not collect primary data or test empirical hypotheses. Instead, it synthesizes cross-disciplinary literature from the learning sciences, computer science, and management information systems to examine the emerging role of agentic AI in higher education. The study seeks to clarify a rapidly evolving concept, synthesize existing knowledge, and develop a framework for understanding agentic AI governance in higher education. Following Torraco's (2005) integrative review methodology, this study employs a structured narrative review to synthesize literature from multiple disciplines. This approach was selected because research on agentic AI remains conceptually fragmented and spans several academic domains. The four research questions addressed in this study also varied in nature, encompassing taxonomic, descriptive, theoretical, and normative concerns that are not well suited to a traditional systematic review. Literature searches were conducted between January and March 2026 using ERIC, Education Source, PsycINFO, ABI/INFORM, and Google Scholar databases. Search terms included agentic AI, autonomous AI agents, AI agency in education, higher education governance, learning management system automation, and instructional design autonomy. Boolean operators were used to combine primary and contextual search terms to identify relevant literature across educational, technological, and governance perspectives. Sources were included if they addressed agentic AI, autonomous systems, educational applications of AI, governance frameworks, instructional design, student support systems, or higher education policy. Sources were excluded if they focused exclusively on technical system development without educational implications, addressed AI applications unrelated to education, were duplicate records, or lacked sufficient relevance to the research questions. Priority was given to peer-reviewed journal articles, conference proceedings, dissertations, and scholarly publications. The screening process involved an initial review of titles and abstracts, followed by a full-text assessment of potentially relevant sources. Because peer-reviewed research on agentic AI in educational settings remains limited,

selected industry reports, organizational publications, and practitioner sources published between 2020 and 2026 were included to capture emerging developments and documented implementations. These sources were evaluated critically and used to supplement, rather than replace, peer-reviewed literature. The final corpus of literature was selected based on its relevance to the four research questions and its contribution to understanding the educational, organizational, and governance implications of agentic AI.

Results and Discussion

The findings of this study carry implications beyond technology adoption: agentic AI does not merely augment the educator's role; it redistributes instructional authority. The core difference between agentic AI and earlier tools is not speed or power, but the shift from responding to initiating. This functional change transforms these systems from passive support tools into active participants that already shape decisions across tutoring, grading, and advising. In each domain, the pattern is consistent: systems that once required human activation now monitor conditions, interpret data, and act before a faculty member has decided. Together, these findings provide the foundation for the Agentic AI Integration Continuum introduced in this paper.

The Agentic AI Integration Continuum

The Agentic AI Integration Continuum constitutes the primary conceptual contribution of this review. It provides higher education institutions with a shared vocabulary for locating current AI practices, assessing the risks associated with each level of autonomy, and making deliberate rather than reactive governance decisions. The Continuum maps four levels of AI autonomy, each representing a different answer to one foundational question: who holds decision-making authority, the educator or the system?

Level One: AI as an Instructional Supplement

At Level One, the answer is always the educator. AI functions as a responsive tool that supports grammar review, summarization, and quiz generation, activated and directed by a human actor. The learning environment remains under full educator control. Governance at this level appropriately centers on disclosure and academic integrity policy (Kostopoulos et al., 2025; Wynants et al., 2025).

Level Two: AI as an Adaptive Assistant

At Level Two, the system becomes an active participant in the learning process. Unlike earlier tools, the Level Two deployment at the University of Michigan uses Gemini to act as a conversational partner that prioritizes inquiry over information delivery (Michigan Ross, 2025). By deliberately withholding direct answers, the system shifts from a passive resource to an active participant that shapes how students navigate their own learning. The defining characteristic of Level Two systems is that they remain user-initiated. Although the system may personalize responses, adapt content, or guide learning interactions, it acts only after a student, instructor, or administrator initiates engagement. Observable indicators include adaptive feedback, conversational tutoring, personalized recommendations, and interactive support that requires ongoing human participation to trigger action.

Level Three: AI as a Semi-Autonomous Agent

At Level Three, the system stops waiting for commands and starts taking the initiative. Instead of just reacting to a user, it constantly tracks data behind the scenes, spots trends, and steps in entirely on its own. This looks like automatic check-ins, instant alerts, and shifting learning paths based on how a student is actually behaving and performing. Essentially, the system catches and fixes a fading student's trajectory before a professor even realizes there's an issue (Carnegie Learning, 2026; Khan Academy, 2023). Georgia State University's "Pounce" advising system is the perfect real-world example here. For nearly a decade, Pounce has used academic and behavioral triggers to text students out of the blue, offering help before

anyone asks for it. The data prove it works, too: a randomized controlled trial found that students receiving these proactive nudges were far more likely to earn a B or higher and much less likely to drop out (Meyer et al., 2024).

Level Four: AI as an Autonomous Institutional Actor

At Level Four, faculty oversight is no longer present in real time. Level Four systems are already managing course flows and administrative tasks with almost no human intervention (Palmer, 2026). The introduction of an autonomous agent in 2026 demonstrated this capability, showing that a system could navigate Canvas, participate in forums, and submit work entirely on its own (Palmer, 2026). This highlights a major structural weakness: online education is built on the assumption that users are human students. As autonomous systems become capable of independently completing coursework and navigating online learning environments, many existing assessment models and academic integrity policies may be insufficient to address these emerging challenges. Most platforms assume users are always human students, leaving current grading and integrity policies unprepared for cases where that is not true. Institutions operating at Level Four without an established governance framework are not effectively managing the risks associated with autonomous AI deployment. At this stage, clear accountability chains and limits on autonomy must be integrated into the operational design from the outset. Governance at this level must be built into the operational architecture from the ground up, with explicit accountability chains, defined autonomy parameters, and continuous monitoring established before deployment. As institutions consider where their current practices fall on this continuum, the central question is not whether AI should be adopted but whether governance structures will be in place before autonomous systems begin making consequential decisions without them.

Instructional Agency, Teaching Roles, and Course Design

The redistribution of instructional authority that the Continuum makes visible produces consequences on two fronts simultaneously, for educators and for students, and literature establishes these not as separate concerns but as two sides of the same structural transformation. On the faculty side, the shift is well documented. Kostopoulos et al. (2025) and Sapkota et al. (2026) establish that when AI takes the lead in giving feedback and mapping out learning paths, the teacher no longer makes every instructional call alone. Baltimore et al. (2025) and Mah (2025) build on this, showing that a professor's role is shifting away from simple content delivery. Instead, faculty are becoming the architects of the system, setting its operational boundaries, ensuring it aligns with pedagogical goals, and retaining the final say to override the AI whenever it falls short on student development or academic standards. Wang and Chou (2025) frame this as a foundational question of instructional leadership that most institutions have not yet resolved. Tadimalla (2025) observes that agentic systems restructure lessons in real time, pushing course design away from linear paths toward modular frameworks. To manage this, instructors must establish clear learning goals and transparent boundaries for these automated adaptations (Alqurni, 2026). What the Continuum adds to these findings is precision: the nature and extent of this role shift is not uniform. It is determined by which level of autonomy an institution has deployed. A faculty member working alongside a Level Two system faces a fundamentally different professional reality than one whose institution has deployed a Level Three or Level Four system without defined oversight protocols.

On the student side, literature identifies genuine benefits and genuine risks, and both deserve equal analytical weight. Pop et al. (2025) and Brohi et al. (2025) show that consistent feedback and targeted support help prevent students from dropping out. However, Alqurni (2026) and Patel et al. (2025) caution that if an AI intervenes too early, students lose the opportunity to develop their own problem-solving and self-regulation skills. This creates a tension in which maximizing support can weaken deep learning. The challenge for designers is not simply using AI but ensuring it does not remove the productive struggle students need to grow.

Governance as an Institutional Obligation

The governance literature reviewed in this paper reflects a field that has diagnosed the problem clearly but has not yet produced tools adequate to manage it in practice. California State University (CSU), Fullerton's ETHICAL Principles framework (Wynants et al., 2025) provides a solid baseline for core values such as transparency and compliance. But top universities are already moving past theoretical rules. Instead of flat-out banning AI, schools like Harvard and the Massachusetts Institute of Technology (MIT) are rolling out tiered policies to encourage people to use the tech responsibly rather than just avoid it (Thesify, 2025). Vague ethical guidelines just aren't enough when an AI can rewrite learning tracks or act without checking in first (Alqurni, 2026; Gao, 2025). Without clear boundaries, these systems risk overstepping before anyone catches the mistake. You need practical design guardrails, like the Agentic Product Canvas (Wynants et al., 2025), to actually force the AI to stand down and let a human take over. Ethics tell you what to do; design frameworks actually make the system do it. Together, they establish institutional values while defining operational limits for AI autonomy. What has been missing from the literature is a shared vocabulary for determining which systems require which levels of governance oversight. The Agentic AI Integration Continuum is intended to address that gap. By mapping AI into four distinct levels of autonomy, the Continuum lets colleges match their rules and oversight to the actual risk of each tool, instead of treating all AI the same. Governing these systems isn't a one-and-done task. It takes ethical boundaries to set the rules, design controls to enforce them, and constant tracking to watch how the tech evolves. Ultimately, the Continuum shows that these guardrails have to be active *before* deployment, not thrown together after something breaks.

Conclusion

This paper addressed four primary research questions concerning the emergence of agentic AI in higher education. First, it established that agentic AI differs from earlier adaptive and analytics-based educational technologies by independently initiating instructional actions and making decisions with limited human intervention (Alqurni, 2026; Pati, 2025; Sapkota et al., 2026). Second, it demonstrated that current educational AI systems already exhibit agentic behaviors through adaptive tutoring, personalized advising, automated interventions, and proactive student support (Carnegie Learning, 2026; Khan Academy, 2023; Meyer et al., 2024). Third, it showed that increasing AI autonomy is reshaping teaching roles, student engagement, and instructional design by shifting faculty responsibilities toward oversight, governance, assessment redesign, and the supervision of AI-supported learning environments (Baltimore et al., 2025; Kostopoulos et al., 2025; Mah, 2025). Fourth, it identified the need for governance frameworks that align institutional oversight with varying levels of AI autonomy. The Agentic AI Integration Continuum provides a conceptual framework for guiding governance, accountability, and implementation decisions (Gao, 2025; Wynants et al., 2025). As demonstrated by the Georgia State University case, the successful integration of agentic AI requires intentional governance and human oversight before autonomous systems are deployed, ensuring that increasing AI autonomy enhances rather than undermines educational practice (Meyer et al., 2024).

Limitations

It is important to acknowledge the following limitations. Because this review is conceptual rather than empirical, it cannot establish causal relationships between the deployment of agentic AI and educational outcomes. The Agentic AI Integration Continuum still requires empirical validation across institutional contexts, and its four levels represent analytical distinctions rather than defined operational boundaries.

Directions for Future Research

Future research should test the Agentic AI Integration Continuum across different types of institutions to examine how governance plays out in practice. Future research should also include long-term studies to determine whether sustained reliance on autonomous systems diminishes students' capacity for independent problem-solving. Examining this question through information systems theory could help clarify who is truly accountable when autonomous decisions directly affect student success.

Recommendations

The recommendations below translate the Agentic AI Integration Continuum into practical considerations for faculty, instructional designers, and higher education institutions. Effective implementation requires deliberate decisions regarding the degree of autonomy granted to AI systems and the governance structures used to oversee their operation.

- **Establish clear AI use policies:** Faculty and institutions should clearly communicate the appropriate purposes and boundaries of AI use in coursework. Policies should distinguish between AI as a learning support tool and AI used as a substitute for student learning and intellectual engagement.
- **Redesign assessments for authentic learning:** Assessment strategies should be reviewed to ensure that student performance reflects genuine learning. Assignments that can be completed primarily through autonomous AI systems should be redesigned to emphasize process documentation, iterative development, reflection, and demonstrated human participation.
- **Maintain human oversight:** Faculty should continue to review AI-generated feedback, recommendations, and instructional interventions to ensure alignment with course objectives, academic standards, and student needs. Human judgment remains essential for addressing contextual and developmental factors that automated systems may not fully capture.
- **Promote critical AI literacy:** Educational programs should provide students with opportunities to critically evaluate AI-generated outputs, recognize system limitations, and develop independent judgment when interacting with automated systems.
- **Support continuous course evaluation:** Institutions and instructional teams should regularly evaluate the impact of AI-supported learning environments to ensure that efficiency gains do not diminish meaningful engagement, critical thinking, or other core educational outcomes.

The recommendations should be implemented through a phased roadmap beginning with assessment redesign, followed by AI governance integration and ongoing institutional evaluation. Documenting implementation efforts will strengthen both institutional practice and the broader evidence base for responsible adoption.

Acknowledgements

The authors confirm they reviewed and followed the journal's AI policy. The editor tool within Word was used only sparingly for sentence-level edits and minor proofreading; all substantive content, analysis, and conclusions are the authors' original work.

References

- Abou Ali, M., Dornaika, F., & Charafeddine, J. (2025). Agentic AI: A comprehensive survey of architectures, applications, and future directions. *Artificial Intelligence Review*, 59(1), Article 11.
- Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous intelligence for complex goals: A comprehensive survey. *IEEE Access*, 13, 18912–18936. <https://doi.org/10.1109/ACCESS.2025.3532853>
- Alqurni, J. (2026). Exploring the role of agentic AI in fostering self-efficacy, autonomy support, and self-learning motivation in higher education. *Frontiers in Artificial Intelligence*, 9, Article 1738774. <https://doi.org/10.3389/frai.2026.1738774>
- Baltimore, A., Bush, K., & Uppinghouse, J. (2025). *A study on faculty and student perceptions of the use of AI on a college campus* [Doctoral dissertation, Lipscomb University].
- Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The rise of agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges. *Future Internet*, 17(9), 404.
- Brohi, S., Mastoi, Q.-a., Jhanjhi, N. Z., & Pillai, T. R. (2025). A research landscape of agentic AI and large language models: Applications, challenges, and future directions. *Algorithms*, 18(8), 499.
- Carnegie Learning. (2026). *Meet MATHia*. Retrieved April 17, 2026, from <https://www.carnegielearning.com/>
- Gao, Y. (2025). *Agentic AI in higher education* [Doctoral dissertation, The Ohio State University].
- Hosseini, S., & Seilani, H. (2025). The role of agentic AI in shaping a smart future: A systematic review. *Array (New York)*, 26, Article 100399. <https://doi.org/10.1016/j.array.2025.100399>
- Instructure. (2026). *Data insights and analytics for Canvas LMS*. Retrieved March 5, 2026, from <https://www.instructure.com/solutions/data-insights>
- Ivy.ai. (2026). *AI chatbots and virtual assistants for higher education*. Retrieved March 5, 2026, from <https://ivy.ai/>
- Khan Academy. (2023). *Khanmigo: AI-powered tutor and teaching assistant*. Retrieved March 1, 2026, from <https://www.khanmigo.ai/>
- Kofinas, A. K., Tsay, C. H., & Pike, D. (2025). The impact of generative AI on academic integrity of authentic assessments within a higher education context. *British Journal of Educational Technology*, 56(6), 2522–2549. <https://doi.org/10.1111/bjet.13585>
- Latif, E., Mai, G., Nyaaba, M., Wu, X., Liu, N., Lu, G., Li, S., Liu, T., & Zhai, X. (2023). *AGI: Artificial general intelligence for education* (arXiv:2304.12479). <https://arxiv.org/abs/2304.12479>
- Madrid-García, A., Benavent, D., & Merino-Barbancho, B. (2025). From chat to act: Large language model agents and agentic AI as the next frontier of AI in rheumatology. *EULAR Rheumatology Open*, 1, 147–156.

- Mah, C. (2025). *Co-constructing AI literacy with writing teachers at the dawn of generative AI* [Doctoral dissertation, Stanford University].
- Mainstay. (2024). *AI student success platform for higher education*. Retrieved April 5, 2026, from <https://mainstay.com/>
- Mendoza-Pinto, R. (2025). Exploring the potential of ChatGPT in developing a virtual assistant for prevention and support in cases of school bullying. *CLEI Electronic Journal*, 28(4), Paper 6.
- Meyer, K., Page, L. C., Mata, C., Smith, E. N., Walsh, B. T., Fifield, C. L., Tyson, M., Eremionkhale, A., Evans, M., Frost, S., & Jung, E. E. (2024). *Let's chat: Leveraging chatbot outreach for improved course performance* (EdWorkingPaper No. 22-564). Annenberg Institute at Brown University. <https://doi.org/10.26300/es6b-sm82>
- Michigan Ross. (2025). *U-M's Ross School of Business and Google Public Sector launch virtual teaching assistant pilot program* [Press release]. Retrieved March 15, 2026, from <https://news.umich.edu/u-ms-ross-school-of-business-google-public-sector-launch-virtual-teaching-assistant-pilot-program/>
- Palmer, K. (2026). Agentic AI can complete whole courses. Now what? *Inside Higher Ed*. Retrieved April 20, 2026, from <https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2026/02/26/agentic-ai-can-complete-whole-courses-now>
- Pati, A. K. (2025). Agentic AI: A comprehensive survey of technologies, applications, and societal implications. *IEEE Access*, 13, 151824–151837.
- Patel, D. (2025). *Investigating the artificial intelligence policy landscape in higher education* [Doctoral dissertation, Frostburg State University].
- Patel, K., Shah, M., Qureshi, K. M., & Qureshi, M. R. N. (2025). A systematic review of generative AI: Importance of industry and startup-centered perspectives, agentic AI, ethical considerations and challenges, and future directions. *Artificial Intelligence Review*, 59(1), Article 7.
- Piccialli, F., Chiaro, D., Sarwar, S., Cerciello, D., Qi, P., & Mele, V. (2025). AgentAI: A comprehensive survey on autonomous agents in distributed AI for Industry 4.0. *Expert Systems With Applications*, 291, 128404.
- Pop, M. V., Tont, G., Flonta, F. V., & Flore, M. (2025). Agentic AI in STEM education: Enhancing cognitive flexibility and workforce readiness. *BRAIN: Broad Research in Artificial Intelligence and Neuroscience*, 16(Suppl. 1), 239–249.
- Precedence Research. (2025). *Agentic AI market size, share, and trends 2025 to 2034*. Retrieved March 12, 2026, from <https://www.precedenceresearch.com/agentic-ai-market>
- PwC. (2025). *PwC's AI agent survey: Key trends about AI agents in the enterprise*. Retrieved March 15, 2026, from <https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html>
- Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2026). AI agents vs. agentic AI: A conceptual taxonomy, applications, and challenges. *Information Fusion*, 126, Article 103599.

- Squirrel AI Learning. (2026). *Squirrel AI adaptive learning platform*. Retrieved April 20, 2026, from <https://squirrelai.com/>
- Sukharevsky, A., Kerr, D., Hjartar, K., Hamalainen, L., Bout, S., Di Leo, V., & Dagorret, G. (2025). *Seizing the agentic AI advantage*. McKinsey & Company. Retrieved April 1, 2026, from <https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage>
- Tadimalla, S. Y. (2025). *Re-framing AI identity and education through socio-technical models for human-AI ecosystems* [Doctoral dissertation, The University of North Carolina at Charlotte].
- Thesify. (2025). *Generative AI policies at the world's top universities: October 2025 update*. Retrieved April 1, 2026, from <https://www.thesify.ai/blog/gen-ai-policies-update-2025>
- Torraco, R. J. (2005). Writing integrative literature reviews: Guidelines and examples. *Human Resource Development Review*, 4(3), 356–367. <https://doi.org/10.1177/1534484305278283>
- Wang, N.-C., & Chou, Y.-C. (2025). Using agentic AI to enhance the quality of academic libraries' responses for student queries. In *Proceedings of the Association for Information Science and Technology: 88th Annual Meeting* (Poster session).
- Whippy. (2026). *Best AI software for educational institutions*. Retrieved March 15, 2026, from <https://www.whippy.ai/ai-software/best-ai-software-for-educational-institutions>
- Wynants, S., Childers, G., De La Torre Roman, Y., Budar-Turner, D., & Vasquez, P. (2025). *ETHICAL principles AI framework for higher education*. California State University, Fullerton.
- Xiong, F. (2025). *Developing artificial intelligent agents for accounting and auditing with large language models* [Doctoral dissertation, Rutgers, The State University of New Jersey].