

CULTIVATING IDENTITY, WORKFORCE READINESS, AND HEUTAGOGICAL LIFELONG LEARNING: THE CASE FOR STUDENT-TRAINED AI AGENTS IN POSTSECONDARY EDUCATION

BY

*James Hutson

Lindenwood University, USA

Receive date: 17-08-2-25

Accept date: 20-08-2025

Abstract

The accelerated integration of human-centered artificial intelligence and human–computer intelligent interaction (HCII) into postsecondary education has inaugurated a paradigm shift in how learners develop cognitive, socio-emotional, and professional competencies. This article argues that requiring students to iteratively build and refine custom AI agents—digital “clones” or GPTs trained on their evolving beliefs, knowledge, values, and interests—addresses three pressing challenges facing contemporary higher education. First, the process of self-modeling via agent training compels deep self-reflection, fostering socio-emotional development and identity formation during the critical neurodevelopmental period that extends into the mid-twenties. Second, as industry trends increasingly prioritize the creation of personal skill and knowledge agents to automate routine tasks and augment human productivity, students who graduate with a mature agent clone are directly prepared for workforce demands in knowledge-intensive sectors. Third, the longitudinal, reflective construction of these agents naturally embeds heutagogical and metacognitive practices, equipping learners with cognitive flexibility, resilience, and the adaptive self-regulation necessary for lifelong learning in a volatile labor market. Drawing on empirical findings from collaborative AI literacy and metacognition research, as well as recent advances in HCII and distributed agent design, this article outlines a transformative educational model that aligns identity work, career preparedness, and adaptive expertise. Such an approach not only individualizes learning but also democratizes access to scalable, context-sensitive cognitive support for a diverse student population.

Keywords: *human-centered AI, metacognition, agent-based learning, socio-emotional development, lifelong learning, heutagogy*

Introduction

Over the past decade, artificial intelligence has advanced from a domain of technical novelty to an omnipresent force shaping cognition, communication, and the contours of professional life. This transformation is embodied most distinctly in the rise of generative models—systems that simulate human-like language, generate complex imagery, and orchestrate multimodal exchanges with remarkable fluency. Such advances have ushered in a new epoch of human–computer intelligent interaction, or HCII, characterized by real-time, dialogic collaboration between individuals and agentic systems across domains as varied as education, industry, and public life (Ding et al., 2024). These agents, equipped with conversational and adaptive capabilities, now engage

as partners in iterative meaning-making, collaborative decision-making, and personalized inquiry. The transition from interface to interlocutor, however, fundamentally alters the landscape of learning. Competence in this environment requires not only operational fluency but also the critical capacity to evaluate, guide, and reciprocally evolve with intelligent agents (Sidra & Mason, 2024). Within this context, the notion of “training an AI to think like oneself” is no longer speculative—it becomes an essential principle for educational praxis.

The imperative to redesign postsecondary education arises from more than the mere proliferation of advanced technologies; it is driven by a profound epistemic crisis in how contemporary societies define, measure, and transmit knowledge, skill, and identity.

Traditional educational models, predicated on the one-time acquisition of content and the episodic demonstration of skill, have become insufficient in a world where professional categories, social roles, and epistemologies themselves are increasingly fluid. Empirical projections underscore that by 2030, the majority of the global workforce will require substantive reskilling—a direct consequence of automation and AI-driven transformation (George, 2024). Contemporary learners now face a landscape that demands navigation across multiple identities and contexts, each marked by distinct expectations and emergent norms. Within this terrain, educational institutions must cultivate cognitive flexibility, robust metacognitive agency, and the socio-emotional intelligence required for ongoing adaptation (Maier & Triff, 2024). The synthesis of human-centered AI with the development of personalized, custom-trained agents offers a unique opportunity to reconceive the educational project as one of continuous identity formation, lifelong cognitive growth, and anticipatory professional rehearsal. By making agent construction and refinement a central, iterative practice, institutions may bridge developmental, strategic, and intellectual objectives within a single, scalable pedagogical structure.

This analysis contends that postsecondary education should be anchored in a tripartite model of agentic learning, wherein each student progressively develops a custom-trained digital agent—a “knowledge clone”—over the duration of their academic formation. The model responds to three interlinked challenges. First, the agent-building process externalizes beliefs, epistemic commitments, and socio-emotional dispositions, rendering them available for deliberate reflection and revision. The digital agent becomes not merely a tool but a mirror, catalyzing ethical self-awareness and more sophisticated forms of personal development (Hutson, 2025). Second, the active construction of functional, task-oriented agents mirrors emergent professional norms in sectors where digital counterparts automate repetitive or information-intensive work. Case studies, such as the deployment of GPT-based assistants in firms like Mineral, illustrate how these systems foster productivity, knowledge reuse, and novel forms of collaboration

(Matthews al., 2025). Third, the recursive cycle of agent tuning—setting goals, monitoring outcomes, and evaluating strategies—grounds and extends metacognitive development, as students become adept at orchestrating their own learning processes amid dynamic, AI-rich environments.

The tripartite framework integrates three mutually reinforcing dimensions: identity formation, agentic professionalization, and lifelong metacognition. In the first dimension, students training an agent to reason or answer as they would must systematically articulate and test their own positions, values, and assumptions. For instance, designing an agent to handle ethical dilemmas in consensus environments compels codification and operationalization of moral intuitions. The second dimension involves building agents that emulate domain-specific expertise—such as summarizing academic literature, generating executable code, or producing data visualizations—reflecting the growing industrial norm of knowledge and skill cloning (Mollick et al., 2024). The third dimension, sustained metacognitive practice, is instantiated as students engage in deliberate agent refinement, informed by validated measures such as the Collaborative AI Metacognition Scale, which correlates reflective AI use with superior task performance and motivational persistence (Sidra & Mason, 2024). Collectively, these vectors instantiate a pedagogy responsive to the neurodevelopmental, economic, and technological imperatives of contemporary higher education.

The significance of this agent-centered model is both conceptual and practical. It repositions the learner not as a passive recipient of content, but as an architect of intelligence—one whose growth emerges from the recursive interplay of identity, skill, and strategic regulation within an interactive, technologically enhanced context. This transformation enables educators to design environments in which learners iteratively construct, critique, and evolve their own cognitive tools, with formative feedback and ethical scaffolding guiding the process. HCI-driven environments, especially those leveraging multimodal, affectively attuned agents, have demonstrated heightened engagement and enhanced self-regulation

in both experimental and applied settings (Ding et al., 2024). The scalable, adaptive nature of these models offers the potential for unprecedented equity in learning support, enabling tailored developmental pathways for diverse student populations.

Foregrounding the co-evolution of human and machine intelligence in educational design, the tripartite framework advanced in this analysis repositions the university not as a static repository of content, but as a dynamic laboratory for self-directed cognitive engineering and ethical maturation. This reconceptualization draws directly from the principles of heutagogy, defined as the study and practice of self-determined learning in which learners take primary responsibility for identifying their needs, selecting strategies, and evaluating outcomes across unpredictable contexts (Hase & Kenyon, 2007). Heutagogical pedagogy is premised on the understanding that, in an era marked by rapid technological disruption and the proliferation of intelligent systems, the most critical capacity graduates must possess is the ability to adapt, unlearn, and relearn autonomously throughout the lifespan. Self-directed learning, in this sense, is not an ancillary skill but a foundational attribute—one that must be intentionally cultivated within the formal curriculum to ensure that learners continue to evolve and thrive well beyond graduation. Embedding heutagogical strategies within agent-based educational models ensures that students acquire the metacognitive habits, reflective practices, and adaptive dispositions necessary to steer their own intellectual, ethical, and professional trajectories. Such a vision compels institutions to reimagine curriculum, assessment, and policy in ways that treat learning as an ongoing, technologically mediated process of self-authorship and meaningful social engagement.

2. Theoretical Foundations and Rationale

A meaningful transformation of postsecondary education in the context of human-machine symbiosis must be grounded in sound theoretical frameworks that anticipate not only how individuals learn, but also how they adapt. The current epoch, marked by the proliferation of generative AI and the normalization of intelligent agents in daily life, demands that

institutions interrogate and expand their pedagogical assumptions. Traditional instructional models based on pedagogical (teacher-directed) and andragogical (self-directed but structured) principles prove increasingly insufficient in preparing learners for the fluid, personalized, and agent-mediated environments characterizing both contemporary knowledge work and civic participation (Blaschke, 2012). Education must now cultivate not only knowledge and skills but also the capacity for continual learning, reflective self-regulation, and adaptive transfer. Theoretical grounding in heutagogy, which foregrounds learner agency, capability, and metacognitive growth, offers a viable scaffolding for this transition. Heutagogy situates learners as designers of their own learning paths, a stance that aligns inherently with training agents tailored to individual evolving identity, values, and strategic goals (Hase & Kenyon, 2007). The integration of intelligent systems into learning environments further amplifies the necessity of such an approach, transforming AI from a passive content delivery mechanism into an active participant in the cognitive ecosystem of the learner (Goyal, 2025). Thus, a heutagogical framework serves not only as a philosophical anchor but also as a practical imperative for education in the AI era.

2.1 Heutagogy in the Age of Generative AI

Heutagogy, derived from the Greek *heuriskō* (to discover), refers to the study of self-determined learning in which learners are not merely recipients of knowledge but active constructors of meaning and adaptive strategists across novel and complex situations (Hase & Kenyon, 2000). In contrast to pedagogy, which centers on teacher control, and andragogy, which acknowledges adult learner need for relevance and self-direction within guided structures, heutagogy privileges full learner autonomy and emergent learning outcomes (**Table 1**). Its conceptual cornerstone is *capability* rather than *competence*: the former implies not only knowing how to perform a task, but also when, why, and under what conditions to modify one's approach (Blaschke, 2012). This emphasis is particularly salient within generative AI environments where learners must navigate ill-defined problems, collaborate with dynamic systems, and

evaluate generated outputs under uncertain conditions. In this sense, capability includes ethical discernment, emotional self-regulation, and adaptive metacognition—traits that machines cannot replicate, but which they can help surface, refine, and scaffold when deployed within a human-centered learning model (Goyal, 2025). The emergence of collaborative AI, including custom agents like ChatGPT and Khanmigo, creates not only the technological affordance but also the developmental necessity for

heutagogical learning environments (Sidra & Mason, 2024). These agents require learners to engage in planning, monitoring, and evaluating—key metacognitive strategies that align with the core mechanisms of heutagogical development. Thus, heutagogy provides both a descriptive and prescriptive model for designing educational systems in which students train, refine, and ultimately co-evolve with intelligent agents that mirror and amplify their evolving cognitive identity.

Table 1. Comparative Overview: Pedagogy, Andragogy, and Heutagogy

Learning Paradigm	Primary Learner Role	Instructor Role	Degree of Autonomy	Outcomes (Competency vs. Capability)	Example Practices
Pedagogy	Passive recipient of content	Knowledge transmitter and authority	Low	Competency: mastery of specific content	Lectures, standardized testing, skill drills
Andragogy	Self-directed within defined structures	Facilitator, coach, content guide	Moderate	Competency: task-based proficiency	Case studies, problem-based learning, group projects
Heutagogy	Self-determined learner and co-designer	Curator, provocateur, scaffolder of inquiry	High	Capability: adaptive, reflective, lifelong learning	Reflective journaling, agent building, self-assessment, ethical scenario design

2.2 Neurodevelopment and Identity Formation in Emerging Adulthood

Contemporary neuroscience confirms that human brain maturation extends well into the mid-twenties, particularly within the prefrontal cortex, the region responsible for executive function, impulse control, long-term planning, and reflective judgment. Functional imaging and longitudinal neuroanatomical studies demonstrate that the maturation of white matter tracts, myelination, and synaptic pruning continues throughout emerging adulthood, facilitating more efficient, integrated, and self-regulated cognition (Sawyer et al., 2018). These findings reinforce that individuals in postsecondary contexts are neurologically primed for advanced metacognitive tasks, yet still developing the cognitive scaffolding necessary for abstract self-reflection and behavioral synthesis. Concurrently, identity formation during this life phase constitutes a central developmental milestone, marked by the integration of personal

values, vocational goals, ethical commitments, and socio-emotional dispositions into a coherent narrative self (Arnett, 2014). Postsecondary education, when intentionally structured, serves as a crucible for this developmental synthesis, offering cognitive, social, and ethical challenges that prompt students to author and revise their evolving sense of self. Within this context, agentic reflection—particularly through the training of agents modeled after individual reasoning, priorities, and moral frameworks—becomes a powerful developmental tool. Teaching an agent to reflect, decide, and express as its human counterpart compels the learner to articulate, interrogate, and refine personal epistemologies. Rather than treating AI as external scaffolding, this framework situates it as an instrument for internal consolidation, aligning technological interaction with developmental science to cultivate intentional self-authorship during the most formative years of adult identity.

2.3 The Future of Work and the Rise of Agentic Professionalization


In parallel with neurological and identity development, the sociotechnical landscape of work is undergoing a rapid transformation defined by the rise of personal AI agents, copilots, and autonomous cognitive assistants. Employers increasingly expect workers to operate in hybrid knowledge ecosystems where intelligent agents serve not only as productivity enhancers but as replicable extensions of individual expertise. Tools such as Microsoft 365 Copilot, Salesforce's Einstein GPT, and open platforms like OpenAI's custom GPTs enable users to offload repetitive, analytical, and semi-creative tasks to AI models tuned to their communication style, disciplinary logic, and task preferences. Notably, organizations like Mineral have instituted structured programs where employees co-develop "GPT pods"—personalized digital assistants fine-tuned to mirror their work processes, decision heuristics, and client-facing functions (Kellogg et al., 2024). Similarly, Khan Academy's *Khanmigo* offers an adaptive pedagogical interface that models Socratic questioning to support learning autonomy and iterative improvement in real time (Sidra & Mason, 2024). The emergence of such agentic infrastructures signals a tectonic shift in both employability metrics and knowledge retention. No longer is productivity judged solely by task execution, but by the capacity to externalize cognitive frameworks into flexible, scalable, and reusable models. Within this context, requiring students to construct and refine their own agentic systems during higher education serves dual purposes: it prepares them for a future of collaborative cognition with digital agents and builds the infrastructure for lifelong learning and re-skilling.

Agent creation is not merely a technical endeavor but a formative act of professional simulation—one that anticipates and embeds learners into the emerging dynamics of intelligent work.

3. The Tripartite Framework of AI Agent Development in Postsecondary Education

The following framework proposes a structured model for embedding agent training within postsecondary learning, grounded in developmental theory, cognitive science, and workforce projections. Rather than treating intelligent systems as passive tools, this model repositions them as co-constructive instruments in the learner's identity, professional, and metacognitive growth. It is comprised of three interlocking dimensions: (1) the use of AI agents as mirrors for ethical formation and identity articulation, (2) the development of domain-specific agentic clones for vocational simulation, and (3) the cultivation of lifelong learning capacities through iterative agent interaction and refinement (**Figure 1**). These dimensions, while analytically distinct, reinforce each other in practice: training an agent to emulate individual reasoning processes simultaneously demands ethical introspection, technical precision, and strategic regulation. The framework not only provides a roadmap for curriculum design and institutional policy, but also models the very competencies—self-authorship, adaptability, and digital fluency—that higher education purports to cultivate. By inviting students to design, critique, and evolve their own intelligent extensions, institutions reposition learners as architects of both cognition and capability in a rapidly transforming epistemic landscape.

Figure 1. The Tripartite Framework of AI Agent Development in Postsecondary Education



3.1 Dimension One: AI Agents as Mirrors for Identity and Ethical Formation

Training an AI agent to think, decide, or argue like oneself necessitates the articulation of deeply held beliefs, epistemological frameworks, and ethical orientations. Unlike conventional assignments, which often reward correct answers within known boundaries, the process of configuring an agent for autonomous reasoning compels students to externalize and formalize their value systems. This may involve scripting ethical decision trees, embedding moral reasoning principles, or training the agent to respond to ambiguous social dilemmas. Reflective practices—such as guided journaling, Socratic self-inquiry, and simulated dialogues with the agent—become essential for ensuring coherence between the learner's identity and the agent's behavior (Hutson, 2024). For instance, students in a political science course might construct an "ethics bot" trained to model utilitarian versus deontological responses to civic policy scenarios. The tension between these programmed outcomes and the student's intuitive responses can catalyze dissonance, prompting meta-ethical interrogation and refinement. This process aligns with established theories of moral development (Kohlberg, 1984), which emphasize the role of cognitive disequilibrium in ethical maturation, and with Baxter Magolda's (2004) theory of self-authorship, which defines maturity as the internal coordination of one's values and decisions. Through agent design, learners externalize and interrogate what they know, why they believe it, and how it guides behavior—transforming AI from an ethical problem into an ethical partner in development.

3.2 Dimension Two: AI Agents as Professional Clones for Work Simulation

In parallel with identity exploration, AI agents function as effective proxies for professional skill replication and simulation. Postsecondary programs in STEM, business, design, law, and health sciences already emphasize scenario-based learning, portfolio development, and domain-specific communication—each of which can be enhanced through agentic modeling. Students might build agents capable of summarizing datasets, drafting business memos, providing first-pass legal analysis, or prototyping user interfaces. Such efforts mirror industry practices,

where organizations increasingly employ generative tools for task automation, client support, and cognitive augmentation (Kellogg et al., 2024). For example, students in a marketing analytics course might construct a prompt-chained agent that generates A/B testing strategies for digital campaigns. In a UX design studio, learners may develop a GPT tuned to critique interface mockups according to accessibility and usability heuristics. These implementations serve not only as simulations but as knowledge artifacts that persist beyond the classroom, available for iterative refinement and reuse. When scaffolded with ethical and technical feedback, these agentic clones become vehicles for competency-based assessment: instructors can evaluate not just what students know, but what their agents are able to perform, justify, and improve. In doing so, education aligns more closely with workforce demands for transparency, scalability, and cognitive reusability—hallmarks of the emerging post-automation economy.

3.3 Dimension Three: AI Agents as Engines of Metacognition and Lifelong Learning

Perhaps the most enduring value of training and refining agents lies in its metacognitive affordances. Effective agent interaction requires students to plan their input strategies, monitor the agent's outputs for alignment or error, and evaluate results against intended goals—mirroring the core metacognitive processes of planning, monitoring, and evaluating (Flavell, 1979). These strategies are further operationalized through tools like the Collaborative AI Metacognition Scale (Sidra & Mason, 2024), which measures user ability to reflect on and regulate their cognitive engagement with AI systems. In practice, students learn to tune prompt specificity, debug hallucinations, revise fine-tuning datasets, and implement system instructions to align agent behavior with desired outputs. For instance, a learner might analyze where an agent's summarization deviated from source intent and revise instructions accordingly—an act that requires not only technical skill but epistemic humility and cognitive flexibility. These practices cultivate adaptive learning behaviors transferrable across disciplines and technologies, rendering the student less dependent on specific tools and more capable of steering diverse systems effectively. Unlike

content mastery, which decays over time, metacognitive fluency is cumulative and generative. When institutionalized through agent interaction, it transforms every engagement into an opportunity for diagnostic reasoning, strategic adjustment, and deliberate growth—hallmarks of resilient lifelong learning in a world of accelerating change.

4. Implementation Strategy and Pedagogical Integration

Translating the tripartite framework into actionable educational practice necessitates an implementation strategy that spans curriculum design, institutional infrastructure, and systems of assessment. A pedagogical model centered on agentic development must be embedded not only in the content of instruction but in the structural and ethical scaffolding

of institutional operations (**Table 2**). Because agent training engages learners in identity formation, professional rehearsal, and metacognitive regulation, its integration into higher education should be sequenced across developmental stages, supported by interdisciplinary faculty, and framed by clear ethical parameters. Implementation must prioritize not only technical fluency but also critical reflection, ensuring that students are not simply building performant systems but cultivating judgment, responsibility, and transferability. This section outlines key pathways for integrating agentic learning into curricular sequences, aligning institutional technology policies, and establishing credentialing mechanisms that validate this emerging literate domain.

Table 2. Modular Curriculum Components for Agentic Literacy

Academic Year	Module Name	Learning Outcomes	Core Activities	Assessment/Portfolio Output
Year 1	Identity and Onboarding	Articulate initial learning goals, identify cognitive preferences, explore ethical values	Self-assessment survey, values inventory, reflective journaling prompts	Foundational AI agent prototype; reflection on learning and identity
Year 2	Academic Integration and Agent Tuning	Apply agent to support learning strategies, integrate academic content, revise self-model	Domain-specific prompt writing, agent-based study planning, skill tracking logs	Agent use-case demos in course assignments; revised agent with academic layer
Year 3	Professional Simulation and Networking	Use agent to simulate workplace tasks, career preparation, and communication patterns	Résumé and cover letter generation, mock interview scripting, simulated professional messaging	Agent as résumé coach; workplace prompt portfolio; interview prep dialogue archive
Year 4	Capstone Reflection and Lifelong Learning Clone	Evaluate agent growth, plan for future upskilling, synthesize personal/professional identity	Ethical testing scenarios, portfolio artifact curation, post-grad goal alignment	Final agent prototype with version history; reflective essay; lifelong learning roadmap

4.1 Curriculum Design and Scaffolding Agentic Literacy

A successful pedagogical integration of agentic development begins with modular curriculum design, strategically distributed across foundational and culminating educational experiences. In the first year,

introductory seminars can incorporate foundational components such as reflective journaling, basic prompt engineering, and agent profiling exercises where students sketch their cognitive patterns, ethical priorities, and decision-making heuristics as inputs for eventual agent modeling. These modules establish a

developmental baseline that is refined over time. In the capstone phase, students engage in full-cycle agent construction, including iterative testing, ethical evaluation, and real-world deployment simulations. Assignments may include creating a domain-specific agent (e.g., a tutoring bot, legal aide, or lab assistant), conducting hallucination analysis, and participating in peer-review critiques of prompt transparency and ethical risk. Faculty support is essential and should include three defined roles: an AI literacy coach who ensures technical fluency and introduces evolving tools; a reflective mentor who scaffolds the metacognitive and identity-based components of agent construction; and a technical liaison who interfaces with IT staff to align platforms, data protocols, and compliance standards. Together, these roles cultivate a holistic agentic fluency that transcends mere tool use and fosters an integrated learning ecology.

4.2 Institutional Infrastructure and Ethical Guardrails

Agentic learning requires robust yet flexible technological infrastructure, grounded in open-access principles and aligned with evolving ethical frameworks. Institutions should adopt interoperable platforms such as OpenAI's ChatGPT (via custom GPT builders), open-source LLMs (e.g., Hugging Face models), and learning management system (LMS) plugins that enable modular agent deployment and privacy-aware feedback loops. These platforms must be accompanied by institutional guidelines for responsible agent construction—covering algorithmic bias, output explainability, hallucination detection, and cultural sensitivity. For example, students should be trained to prompt-test agents across socio-political scenarios to ensure pluralistic reasoning and to audit outputs for factual integrity and epistemic humility. Legal and ethical compliance requires attention to existing regulations: in the United States, FERPA mandates student control over educational records, necessitating secure environments for storing agent logs; GDPR in the EU imposes strict conditions on data processing and transparency; and intellectual property policies must clarify ownership of student-trained agents, particularly if reused or published. Institutional review boards and AI ethics committees should be consulted to establish default policies on agent-generated content, data handling, and permissible use. These infrastructure decisions, while

often invisible to students, are foundational to creating a climate of trust, innovation, and ethical rigor in agentic education.

4.3 Assessment and Credentialing Pathways

Evaluation frameworks must evolve to recognize the complexity and longitudinal nature of AI agent development. Institutions should implement digital badges or microcredentials in areas such as metacognitive agent design, prompt refinement and alignment, and ethical AI implementation, signaling to future employers the learner's fluency in contemporary AI-human collaboration. Formative assessments might include scaffolded rubrics evaluating agent behavior for accuracy, coherence, adaptability, and explainability. For example, students could receive iterative feedback on whether their agent responds with epistemic transparency (e.g., stating confidence or limitations), integrates sources ethically, or improves over time based on user feedback. Summative assessments may take the form of capstone reviews where agents are evaluated alongside a reflective narrative and diagnostic commentary from the student, analyzing what was learned about themselves through the process. Portfolios should document the evolution of the agent across its lifecycle, including design rationales, ethical dilemmas encountered, performance metrics, and future iteration goals. Such portfolios become not only artifacts of learning but also springboards for professional engagement—demonstrating that graduates are not only capable of using AI, but of shaping and governing it with critical, adaptive, and ethical insight.

5. Case Studies and Pilot Programs

A theoretical commitment to agentic learning is most compelling when evidenced by concrete, context-specific practice. The following case studies and pilots illustrate how postsecondary institutions can instantiate the tripartite framework at distinct stages and across varied disciplines. Each example is intentionally detailed, demonstrating both the technological and pedagogical affordances of agentic training, as well as the formative and summative learning outcomes that arise from such work. Through situating these practices in diverse curricular settings—a humanities seminar, a graduate analytics program, and an interdisciplinary capstone—these cases model not only feasibility but also the adaptability and scalability of

agentic education. They also foreground how agent-based assignments create space for ethical inquiry, professional rehearsal, and metacognitive reflection, producing graduates with tangible artifacts of self-authorship and AI literacy.

5.1 Example: Undergraduate Humanities Seminar with Reflective GPT Design

In an upper-level humanities seminar on Ethics, Identity, and Technology, students are tasked with designing a custom GPT-powered conversational agent that simulates their own reasoning about contemporary civic dilemmas (**Figure 2**). The assignment unfolds over a six-week arc: students begin with a series of structured reflective journals, mapping their evolving beliefs about privacy, justice, and free expression in digital culture. These insights are translated into prompt templates and decision trees, which form the foundation for the agent's personality and value logic. Students program the agent to respond to hypothetical scenarios—such as whether to flag misinformation in a

public forum or how to balance transparency and confidentiality in whistle blowing cases—using iterative prompt engineering and scenario testing. Class workshops include peer reviews, where students role-play as both user and agent, surfacing points of cognitive dissonance between the agent's responses and their own intuitions. The final project portfolio includes the agent itself, a critical reflection on moments of surprise or disagreement, and a meta-analysis of how training the agent deepened self-understanding and ethical clarity. Faculty assessment focuses on agentic transparency, consistency, and student capacity to articulate and revise their own frameworks through this human–AI dialogue. In this context, the agent serves as a scaffold for ethical self-authorship and a platform for negotiating value pluralism in contemporary society.

Figure 2. Reflective GPT Design: Articulating Ethical Identity through AI

Undergraduate Humanities Seminar

Assignment Title: Reflective GPT Design: Articulating Ethical Identity through AI

Description: This assignment invites students to develop a custom GPT-powered conversational agent that models their personal ethical reasoning when confronted with civic and technological dilemmas. By translating reflective insights and ethical stances into agentic responses, students externalize and interrogate their own frameworks, encountering cognitive dissonance, revising value commitments, and engaging with pluralistic perspectives. The process integrates reflective journaling, prompt engineering, scenario-testing, and critical analysis, culminating in both a functional agent and a metacognitive narrative.

Instructions:

1. Complete three weekly reflective journal entries mapping your beliefs and priorities regarding ethics in digital society.
2. Using your reflections, design a value logic (decision tree or ethical rules) for your GPT agent.
3. Program your agent to respond to at least five hypothetical civic scenarios (e.g., misinformation, privacy, whistle blowing) using prompt engineering.
4. Participate in peer review workshops, role-playing as both user and agent, and record where your agent's responses diverge from your own.
5. Submit a portfolio that includes:
 - a. The agent's prompt templates and design rationale
 - b. A transcript of peer-reviewed scenario dialogues
 - c. A 1000-word reflection analyzing moments of ethical surprise, value conflict, and personal growth through the agent-building process

Assessment Criteria:

- Depth and clarity of ethical reflection in journals and rationale (25%)
- Coherence, transparency, and pluralism in agent responses (25%)
- Technical quality and creativity of prompt engineering and scenario-testing (20%)
- Critical self-analysis in reflection narrative, addressing cognitive dissonance and ethical growth (20%)

Engagement and feedback provided during peer-review sessions (10%)

5.2 Example: Graduate Business Analytics Program Using Agent Workflows

In a professional Master's program in business analytics, students participate in a Digital Workflow and AI Copilots module (**Figure 3**) designed to simulate enterprise-level deployment of agentic tools. Teams are tasked with building domain-specific GPT agents that automate multi-step analytics tasks: ingesting raw sales data, performing statistical analysis, generating dashboard-ready visualizations, and crafting executive summaries tailored for different stakeholders. Each student assumes rotating roles—as prompt architect, test user, or model tuner—mirroring real-world agile team practices. The agent is stress-tested with real historical data sets, and its outputs are audited for accuracy, explainability, and bias. A running log documents every design decision, error

diagnosis, and prompt revision, producing a transparent record of group metacognition and project management. The final deliverable is not only a working agentic workflow but also a set of best-practice guidelines for onboarding future team members and a “post-mortem” on emergent risks, ethical considerations (e.g., avoiding algorithmic discrimination), and opportunities for knowledge retention within the organization. Employers and alumni are invited to provide external feedback, linking the project to real market demands and reinforcing the relevance of AI literacy in high-stakes business environments. In this setting, agent workflows function both as operational prototypes and as instruments for organizational learning and resilience.

Figure 3. AI Copilot Workflow: Building and Auditing an Analytics Agent

Graduate Business Analytics Program

Assignment Title: AI Copilot Workflow: Building and Auditing an Analytics Agent

Description: Students, in teams, will construct and document a domain-specific GPT agent designed to automate and streamline a multi-step business analytics workflow. The project emphasizes prompt architecture, team-based troubleshooting, explainability, and ethical risk management. Students will document the evolution of their agent, track design decisions, analyze agent performance, and synthesize lessons into guidelines for future professional adoption.

Instructions:

1. Form teams and select a business analytics scenario (e.g., sales forecasting, dashboard visualization, executive summary generation).
2. Architect your agent workflow, defining input-output requirements, data security considerations, and user profiles.
3. Iteratively build, test, and refine the agent, recording all prompt chains, error diagnoses, and model revisions in a shared project log.
4. Audit your agent's outputs for accuracy, explainability, and bias using real or simulated business data sets.
5. Prepare a project portfolio including:
 - a. The final agent workflow and prompt documentation
 - b. Annotated logs detailing troubleshooting and collaborative problem-solving
 - c. A set of “best practices” for agent onboarding and handover
 - d. A group post-mortem report addressing ethical risks, stakeholder impact, and recommendations for scaling
6. Present your workflow and findings to a panel of faculty, alumni, and industry advisors.

Assessment Criteria:

- Completeness, transparency, and explain ability of the agent workflow and documentation (25%)
- Rigor and integrity of data auditing, ethical risk management, and stakeholder analysis (25%)
- Depth and clarity of collaborative problem-solving as demonstrated in logs and group reports (20%)
- Applicability and professionalism of best-practice guidelines and recommendations (20%)

Presentation effectiveness and engagement with external feedback (10%)

5.3 Example: Interdisciplinary Capstone with Lifelong Learning Clone

The culminating case features an interdisciplinary senior capstone, open to students from computing, education, health sciences, and design. Each participant is tasked with constructing a Lifelong Learning Clone (**Figure 4**)—a custom-trained agent designed to serve as their personal learning companion beyond graduation. Students begin by reflecting on their disciplinary journey, mapping strengths, blind spots, motivational triggers, and preferred learning strategies. The agent is iteratively trained on this meta-cognitive and affective data, as well as on curated knowledge bases from the student's field. Its initial use cases include scheduling study sessions, recommending research articles based on emerging

interests, generating practice quizzes, and even providing motivational feedback during periods of self-doubt. Students work in advisory pods, troubleshooting technical and ethical dilemmas such as over fitting the agent to fixed preferences, ensuring data privacy, or setting boundaries for when the agent should “push back” with alternative strategies. The final capstone presentation includes a demonstration of the agent's capacity to adapt to novel learning contexts, a portfolio documenting its developmental arc, and a reflection on the student's evolving theory of self-directed learning. Program alumni are surveyed six months post-graduation to assess agent utility in real-world upskilling and adaptation, providing longitudinal evidence of educational impact. Here, the learning clone becomes not a static artifact but a living extension of the student's heutagogical and metacognitive agency.

Figure 4. Lifelong Learning Clone: Designing a Personal AI Companion for Future Growth

Interdisciplinary Capstone

Assignment Title: Lifelong Learning Clone: Designing a Personal AI Companion for Future Growth

Description: In this capstone, students from diverse fields will create and train a “Lifelong Learning Clone”—a custom AI agent intended to serve as a personal up skilling partner after graduation. The project integrates self-assessment, knowledge curation, motivational profiling, and ethical boundary-setting. Students will iteratively document the agent's development, demonstrate its adaptability, and reflect on the process of preparing for lifelong learning in a volatile world.

Instructions:

1. Begin with a self-assessment mapping strengths, knowledge gaps, learning preferences, and motivational triggers.
2. Curate and input relevant disciplinary content, resources, and metacognitive prompts into your agent's knowledge base.
3. Train and iterate your agent to provide learning recommendations, generate study plans, quiz you, and offer motivational feedback for new learning goals.
4. Participate in advisory pod meetings, addressing challenges such as overfitting, privacy, and ethical limits of your agent's autonomy.
5. Prepare a final portfolio including:
 - a. The agent's development log and use-case demonstrations
 - b. A reflective essay (1200 words) on the process, including adaptations for future scenarios
 - c. Peer and self-assessments on collaboration, troubleshooting, and ethical reasoning
 - d. A six-month plan for continuing to update and adapt your clone post-graduation

Assessment Criteria:

- Insight and authenticity of self-assessment and content curation (20%)
- Agent's adaptability, functionality, and ethical safeguards in use-case demonstrations (30%)
- Depth of reflection on learning processes and strategies for future growth (20%)
- Quality and constructiveness of peer collaboration and advisory participation (15%)

Clarity and feasibility of the post-graduation adaptation plan (15%)


5.4 Longitudinal Model Training: The Four-Year Agentic Development Sequence

A truly innovative application of agentic learning leverages the full arc of the undergraduate experience,

embedding longitudinal model training as an evolving scaffold that shapes academic, professional, and personal growth (**Figure 5**). In this extensive implementation, students begin their journey with the introduction of a personalized AI agent—a foundational “identity clone”—in their first-year seminar or orientation course. Rather than treating the agent as a disposable class project, this model is envisioned as a living, adaptive companion, designed to co-evolve with the learner over the course of four years. The resulting artifact, upon graduation, is not

merely a technical demonstration or reflective journal, but a robust, context-aware agent that encodes the student’s best practices for learning, decision-making, and self-regulation. This lifelong cognitive partner provides not only a durable record of identity formation and metacognitive insight but also a competitive, indispensable portfolio artifact as graduates navigate an employment market transformed by automation and the scarcity of traditional entry-level roles.

Figure 5. Four-Year Longitudinal AI Agent Development Pathway



Year 1: Foundational Self-Modeling and Academic Onboarding

In the first semester, students are introduced to the principles of human-centered AI, metacognition, and heutagogy as part of a required “Introduction to University Life” course. Early activities focus on self-assessment: students articulate initial academic goals, preferred learning strategies, ethical values, and personal interests using structured surveys and reflective essays. These inputs form the first training dataset for their initial agent, which is implemented on a campus-supported platform (e.g., a custom GPT instance or an open-source alternative linked to the learning management system). Students engage in weekly exercises, such as feeding the agent with reflections from academic or social challenges (e.g.,

procrastination, navigating group work, or ethical dilemmas in digital life), and receive personalized feedback or reminders from the agent. Faculty and peer mentors monitor progress, helping students refine both their prompts and their self-understanding, and introducing foundational concepts such as bias mitigation and ethical use of AI.

Year 2: Expanding Context and Academic Specialization

In the second year, as students declare majors or enter professional tracks, the agent is updated to reflect evolving disciplinary interests and cognitive styles. Assignments prompt students to feed the agent with domain-specific content—summaries of favorite classes, preferred note-taking or study techniques, and insights from early internships or research projects.

The agent is scaffolded to serve as a meta-cognitive coach: suggesting new strategies for tackling difficult readings, flagging patterns of procrastination, or recommending campus resources (tutoring, wellness, networking events) based on individual needs. Periodic check-ins require students to compare their agent's suggestions with their actual choices, deepening both metacognitive awareness and alignment between internal goals and external action. Throughout, the agent's transparency and ethical reasoning are stress-tested through scenario-based learning modules, preparing students to distinguish when and how to trust AI in complex decision-making.

Year 3: Professionalization, Networking, and Identity Synthesis

The third year centers on the agent's transition from an academic companion to a pre-professional collaborator. Students train their models to help draft résumés, tailor cover letters, generate industry-relevant interview questions, and map networks of alumni or professional associations. Advanced modules enable the agent to simulate workplace scenarios, critique written communication for tone and clarity, or synthesize feedback from internships and field placements. At the same time, the agent records longitudinal data on learning habits, evolving motivations, and adaptive strategies during setbacks or periods of high stress. Faculty advisors use the agent's behavioral logs—always with student consent—to support discussions of career fit, identity consolidation, and next steps. In reflective workshops, students compare their agent's simulated "advice" to that of human mentors, surfacing any divergence as a prompt for deeper self-authorship. By this stage, the agent has become not just a record of academic progress but an active rehearsal space for professional growth.

Year 4: Capstone Integration and Lifelong Learning Launch

In the final year, the agent is integrated into the capstone or portfolio requirement. Students audit the model's performance, identifying where its recommendations or behavior best matched or diverged from their own values and learning outcomes. The agent is tasked with helping students set up

personalized upskilling plans, recommend graduate programs or certifications, and simulate adaptive responses to anticipated disruptions in their field (such as job automation or new AI technologies). Capstone deliverables include a demonstration of the agent's evolution—via logs, visualizations, and use-case narratives—and a reflective synthesis essay that articulates how building and iterating the model shaped the student's identity, learning philosophy, and career aspirations. In a culminating presentation or digital showcase, students present their "AI clone" to faculty, employers, or external reviewers, illustrating how it will serve as a lifelong learning partner in the unpredictable world beyond graduation.

Portfolio and Market Value

Upon completion, each graduate possesses not only a transcript or résumé but a dynamic, portable digital artifact that provides evidence of four years of self-directed learning, ethical inquiry, and adaptive growth. As entry-level positions contract and employers demand demonstration of cognitive flexibility, technical fluency, and personal agency, the presence of a bespoke lifelong learning agent offers a compelling market signal. More importantly, it provides graduates with an ongoing partner for upskilling, reflection, and adaptive reinvention—ensuring that the value of a university education extends far beyond degree conferral and into a future where human-machine co-evolution is both a challenge and a promise. In this model, building an AI agent across four years becomes not a classroom novelty, but a rite of passage: the symbolic and practical foundation for thriving in the AI-mediated century.

6. Challenges, Limitations, and Future Research

While the integration of custom-trained AI agents into postsecondary education offers immense potential for identity development, professional preparation, and lifelong learning, it also surfaces critical challenges that must be addressed with theoretical rigor and institutional care. The use of generative AI in learning environments invites questions about over-reliance, digital divides, cognitive offloading, and ethical uncertainty. As with any emergent paradigm, enthusiasm must be tempered by continuous critical inquiry and iterative empirical validation. This section outlines three key limitations—risks of over-

automation, equity in access and design, and the current paucity of longitudinal data on educational and cognitive outcomes. Addressing these constraints is essential to developing AI-augmented learning ecosystems that are not only technically functional but pedagogically sound, socially equitable, and developmentally durable.

6.1 Risks of Over-Automation and Metacognitive Laziness

One of the primary concerns associated with the widespread adoption of AI agents in education is the phenomenon of *metacognitive laziness*—a tendency among learners to offload cognitive processes onto intelligent systems rather than actively engaging in reflective and strategic thinking (Fan et al., 2025). As generative AI systems improve in fluency and utility, there is a growing risk that students may become passive recipients of synthesized content, rather than critical evaluators or co-constructors of knowledge. When AI-generated responses appear authoritative, they may suppress epistemic curiosity or reinforce superficial understanding. This is particularly problematic in domains where ambiguity, complexity, and moral judgment are integral to mastery. Moreover, students accustomed to AI-mediated support may struggle with unstructured problem-solving or fail to develop the endurance required for sustained intellectual labor. To mitigate these effects, educational designers must embed deliberate friction points into the learning process—scenarios where students are required to interrogate, revise, or justify agent outputs using transparent reasoning. Just as calculators did not replace numeracy but changed how it is taught, so too must AI prompt the redesign of pedagogies that prioritize *when not to use AI* as much as how to use it effectively.

6.2 Equity and Access in AI Agent Training

The benefits of agentic learning are not distributed evenly. Students from under-resourced institutions or socioeconomically marginalized backgrounds often face compounded barriers to participation in AI-augmented education, including limited access to hardware, bandwidth, and advanced digital literacy instruction. Furthermore, linguistic, cultural, and epistemological biases embedded in foundational language models may alienate learners who do not conform to normative user profiles (Bender et al.,

2021). Without careful scaffolding and inclusive design, agentic learning environments risk reinforcing existing educational inequities rather than redressing them. There is also a risk of uneven instructional preparation: while some faculty may be fluent in AI systems and pedagogy, others may struggle to meaningfully engage with or assess agent-based work. Institutions must thus invest in faculty development, multilingual platforms, and student-centered AI onboarding systems that democratize participation in agentic design. Accessibility audits, bias testing protocols, and optional low-bandwidth agent tools are necessary infrastructural elements to ensure inclusion. In short, equity in agent training is not a technical add-on but a pedagogical precondition for ethical implementation.

6.3 Future Research Questions: Cognitive Impact, Transferability, Longitudinal Efficacy

The theoretical case for agentic education is strong, but empirical research remains in early stages. Critical questions remain unanswered about the *cognitive impact* of AI agent co-development: How does iterative interaction with a self-trained agent shape learners' self-regulation, moral reasoning, or epistemic vigilance? Are certain cognitive profiles more likely to benefit from agentic reflection than others? *Transferability* is another open domain—can metacognitive strategies practiced through agent design be generalized to unfamiliar technologies, disciplinary problems, or non-digital contexts? Moreover, *longitudinal efficacy* must be established through sustained post-graduate studies. Do students who train lifelong learning clones actually return to and refine them over time? Are their career trajectories measurably more adaptive, ethical, or reflective as a result? Methodologically, this calls for mixed-methods designs combining experimental comparisons, reflective artifacts, agent behavior logs, and delayed interviews. It also demands interdisciplinary collaboration across cognitive psychology, AI ethics, learning analytics, and instructional design. As institutions scale these practices, research must follow not only to validate their effectiveness, but to ensure that this educational transformation is aligned with the enduring goals of intellectual autonomy, social equity, and human flourishing.

7. Conclusion

The rapid emergence of generative AI and custom agentic systems compels higher education to fundamentally reorient its mission, moving beyond content delivery toward the cultivation of cognitive adaptability, agentic development, and the dynamic co-construction of identity. In this environment, the capacity to self-direct learning, critically engage with technology, and reflect upon one's beliefs and strategies becomes not merely advantageous but indispensable. The university's role, therefore, is not simply to transmit established knowledge but to act as a crucible in which the next generation learns to thrive in continuous partnership with intelligent systems. By making the design, training, and refinement of AI agents—personalized cognitive collaborators—central to the curriculum, institutions acknowledge that the ability to shape, critique, and adapt digital extensions of oneself is as vital as any technical or disciplinary skill. This is not a pedagogical gimmick, but a profound rite of passage for graduates entering an era defined by fluid boundaries between human and machine cognition. Building an AI clone, in this light, represents the convergence of intellectual autonomy, ethical growth, and technological fluency, equipping learners to navigate and lead within the complexities of an AI-mediated century. Higher education, reimagined through this agentic lens, becomes both a sanctuary for reflective self-authorship and a launchpad for collaborative human-machine innovation.

References

- Albus, P., & Seufert, T. (2024). Artificial Intelligence in Education: the Importance of Metacognitive. *The Impact of Artificial Intelligence on Societies*, 109.
- Aluko, H. A., Aluko, A., Offiah, G. A., Ogunjimi, F., Aluko, A. O., Alalade, F. M., & Nwani, C. H. (2025). Exploring the effectiveness of AI-generated learning materials in facilitating active learning strategies and knowledge retention in higher education. *International Journal of Organizational Analysis*.
- Arnett, J. J. (2023). *Emerging adulthood: The winding road from the late teens through the twenties*. Oxford University Press.
- Bashynska, I., Garachkovska, O., Kichuk, Y., Podashevskaya, T., & Bigus, O. (2021). Smart education 4.0: Balancing dual-distance and reskilling revolution. *Studies of Applied Economics*, 39(6).
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021, March). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 610-623).
- Blaschke, L. M. (2012). Heutagogy and lifelong learning: A review of heutagogical practice and self-determined learning. *The International Review of Research in Open and Distributed Learning*, 13(1), 56-71.
- Borer, C., & Weidinger, W. (2024, June). From Life Skills to Future Skills: Competencies that Nowadays' Children Will Need in Ten Years' Time. In *Conference Proceedings. The Future of Education 2024*.
- Brusilovsky, P., Parra, D., Rahdari, B., Raj, S., & Torkamaan, H. (2025, March). HealthIUI: Workshop on Intelligent and Interactive Health User Interfaces. In *Companion Proceedings of the 30th International Conference on Intelligent User Interfaces* (pp. 183-186).
- Ding, Z., Ji, Y., Gan, Y., Wang, Y., & Xia, Y. (2024). Current status and trends of technology, methods, and applications of Human-Computer Intelligent Interaction (HCII): A bibliometric research. *Multimedia Tools and Applications*, 83(27), 69111-69144.
- Ellikkal, A., & Rajamohan, S. (2024). AI-enabled personalized learning: empowering management students for improving engagement and academic performance. *Vilakshan-XIMB Journal of Management*.
- European Parliament. (2016). *General Data Protection Regulation (GDPR)*. Regulation (EU) 2016/679. <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- Ezzaim, A., Dahbi, A., Aqqal, A., & Haidine, A. (2024). AI-based learning style detection in adaptive learning systems: a systematic

- literature review. *Journal of Computers in Education*, 1-39.
- Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., ... & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, 56(2), 489-530.
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American psychologist*, 34(10), 906.
- Frank, E. (2024). The Influence of artificial intelligence on education: Enhancing personalized learning experiences. *EasyChair Preprint*, 14675.
- George, A. S. (2024). Workforce Impacts By 2030: Job Losses and Opportunities in an Era of Automation and AI. *Partners Universal International Research Journal*, 3(4), 41-47.
- Goyal, A. (2025). AI as a Cognitive Partner: A Systematic Review of the Influence of AI on Metacognition and Self-Reflection in Critical Thinking. *International Journal of Innovative Science and Research Technology*, 10(3), 1231-1238.
- Hase, S., & Kenyon, C. (2000). From andragogy to heutagogy. *Ulti-BASE In-Site*.
- Hase, S., & Kenyon, C. (2007). Heutagogy: A child of complexity theory. *Complicity: An international journal of complexity and education*, 4(1).
- Hutson, J. (2025). Enhancing Metacognitive Competencies through Human-Centered AI: The Role of Custom-Trained Intelligent Agents in Workforce Upskilling. *KOS Journal of AIML, Data Science, and Robotics*, 1(1).
- Kaswan, K. S., Dhattewal, J. S., & Ojha, R. P. (2024). AI in personalized learning. In *Advances in technological innovations in higher education* (pp. 103-117). CRC Press.
- Kellogg, K., Bach, N., Ciccirelli, L. J., & Mollick, E. (2024). No shortcuts for generative AI upskilling. *MIT Sloan School of Management and Boston Consulting Group*. Retrieved: <https://mitsloan.mit.edu/press/no-shortcuts-generative-ai-upskilling>
- Kohlberg, L. (1987). The psychology of moral development. *Ethics*, 97(2).
- Kovalenko, I., & Baranivska, N. (2024). Integrating Artificial Intelligence in English Language Teaching: Exploring the potential and challenges of AI tools in enhancing language learning outcomes and personalized education. *Європейські соціо-правові та гуманітарні студії*, (1), 86-95.
- Li, L. (2024). Reskilling and upskilling the future-ready workforce for industry 4.0 and beyond. *Information Systems Frontiers*, 26(5), 1697-1712.
- Lozano-Osorio, I., Ruiz-Olmedilla, S., Pérez-Marín, D., & Paredes-Velasco, M. (2024). Learning by Teaching: Creation of Tutorials in the Field of Vocational Training. *IEEE Transactions on Education*.
- Maier, M., & Triff, Z. (2024). Artificial intelligence and the development of metacognitive structures involved in problem solving. In *INTED2024 Proceedings* (pp. 2574-2582). IATED.
- Magolda, M. B. B. (2023). Self-authorship as the common goal of 21st-century education. In *Learning partnerships* (pp. 1-35). Routledge.
- Matthews, M. J., Su, R., Yonish, L., McClean, S., Koopman, J., & Yam, K. C. (2025). A Review of Artificial Intelligence, Algorithms, and Robots Through the Lens of Stakeholder Theory. *Journal of Management*, 01492063241311855.
- Mazari, N. (2025). Building Metacognitive Skills Using AI Tools to Help Higher Education Students Reflect on Their Learning Process. *RHS: Revista Humanismo y Sociedad*, 13(1), 2.
- Mijakovic, I., & Rahimi, S. (2024). Learning by teaching efficiently enhances learning outcomes in molecular biology of the cell course. *Biochemistry and Molecular Biology Education*, 52(1), 15-24.
- Mollick, E., Mollick, L., Bach, N., Ciccirelli, L. J., Przystanski, B., & Ravipinto, D. (2024). AI agents and education: Simulated practice at scale. *arXiv preprint arXiv:2407.12796*.

- Ojeda-Ramirez, S., Rismanchian, S., & Doroudi, S. (2023). Learning about AI to learn about learning: Artificial intelligence as a tool for metacognitive reflection.
- Recommendations, E. C. (2018). Key competences for lifelong learning. *Official Journal of the European Union*. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF>.
- Rekha, K., Gopal, K., Satheeskumar, D., Anand, U. A., Doss, D. S. S., & Elayaperumal, S. (2024, May). Ai-Powered Personalized Learning System Design: Student Engagement and Performance Tracking System. In *2024 4th International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE)* (pp. 1125-1130). IEEE.
- Ribeiro, B. A., Coelho, H., Ferreira, A. E., & Branquinho, J. (2024). Metacognition, Accountability and Legal Personhood of AI. *Multidisciplinary perspectives on artificial intelligence and the law*, 169.
- Sajja, R., Sermet, Y., Cikmaz, M., Cwiertny, D., & Demir, I. (2024). Artificial intelligence-enabled intelligent assistant for personalized and adaptive learning in higher education. *Information*, 15(10), 596.
- Sawyer, S. M., Azzopardi, P. S., Wickremarathne, D., & Patton, G. C. (2018). The age of adolescence. *The lancet child & adolescent health*, 2(3), 223-228.
- Sidra, S., & Mason, C. (2024, June). Reconceptualizing AI literacy: The importance of metacognitive thinking in an artificial intelligence (AI)-Enabled workforce. In *2024 IEEE Conference on Artificial Intelligence (CAI)* (pp. 1181-1186). IEEE.
- Syafii, A., Purwati, O., Munir, A., & Suhartono, S. (2024). AI-Powered Pedagogy: Elevating Student Engagement through Integrated Feedback. *Asian Journal of Management, Entrepreneurship and Social Science*, 4(04), 721-739.
- Taguma, M., & Barrera, M. (2019). OECD future of education and skills 2030: Curriculum analysis. *Disponibile su: https://www.oecd.org/education/2030-project/teaching-and-learning/learning/skills/Skills_for_2030.Pdf*.
- Turchi, T., Malizia, A., Paternò, F., Borsci, S., Chamberlain, A., & Fish, A. (2025, March). Adaptive XAI: Advancing Intelligent Interfaces for Tailored AI Explanations. In *Companion Proceedings of the 30th International Conference on Intelligent User Interfaces* (pp. 172-174).
- U.S. Department of Education. (2023). *Family Educational Rights and Privacy Act (FERPA)*. <https://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html>