




OPINION PIECE

Navigating agentic AI: a call for reimagined academic literacies

Octavio Murekian  and Surabil Sudarshan 

University of Greenwich, UK

ABSTRACT

The emergence of agentic artificial intelligence (AI) in higher education raises a critical question: does the move toward autonomous systems capable of planning, decision-making, and action constitute a fundamental shift, or an acceleration of challenges already posed by generative AI? This article argues that these developments require a reimagining of academic literacies, extending beyond traditional emphases on critical thinking and academic writing toward a more comprehensive conception of AI literacy. Drawing on research in AI ethics and digital pedagogy, it explores the implications of learning in partnership with increasingly autonomous systems. It examines how agentic AI disrupts established understandings of authorship, assessment, and intellectual labour, and proposes a framework centred on critical evaluation, prompt literacy, co-authorship, ethical awareness, and recognition of AI limitations. Learning developers are positioned as central actors in this transition, with a key role in shaping pedagogy, assessment, and institutional policy.

KEYWORDS: agentic artificial intelligence, academic literacies, AI literacy, human-AI collaboration.

ARTICLE HISTORY: Received 17 June 2025. Accepted 28 January 2026.

Introduction

The landscape of artificial intelligence (AI) in higher education is undergoing a profound transformation. AI has evolved from a suite of passive tools that assist with discrete tasks into sophisticated platforms for iterative knowledge synthesis, contextual retrieval, personalised assistance, and collaborative research, and is on a path towards agentic systems capable of

autonomous decision-making, adaptive planning, and goal-directed action (Belcic & Stryker, 2025; Moquin, 2025; PwC, 2025). These developments demand that learning developers hasten the shift from a reactive stance focused on misuse and detection, toward fundamental pedagogical inquiry into the nature of knowledge, learning, and human authorship in an era of intelligent automation. While perspectives differ on whether agentic AI crosses a new threshold or intensifies existing pressures, the practical implications for learning developers demand urgent attention. Despite the pedagogical implications of generative AI having been evident since 2022, institutional adoption of literacy frameworks addressing these challenges remains uneven and often reactive rather than strategic.

From tools to agents: examining the nature of change

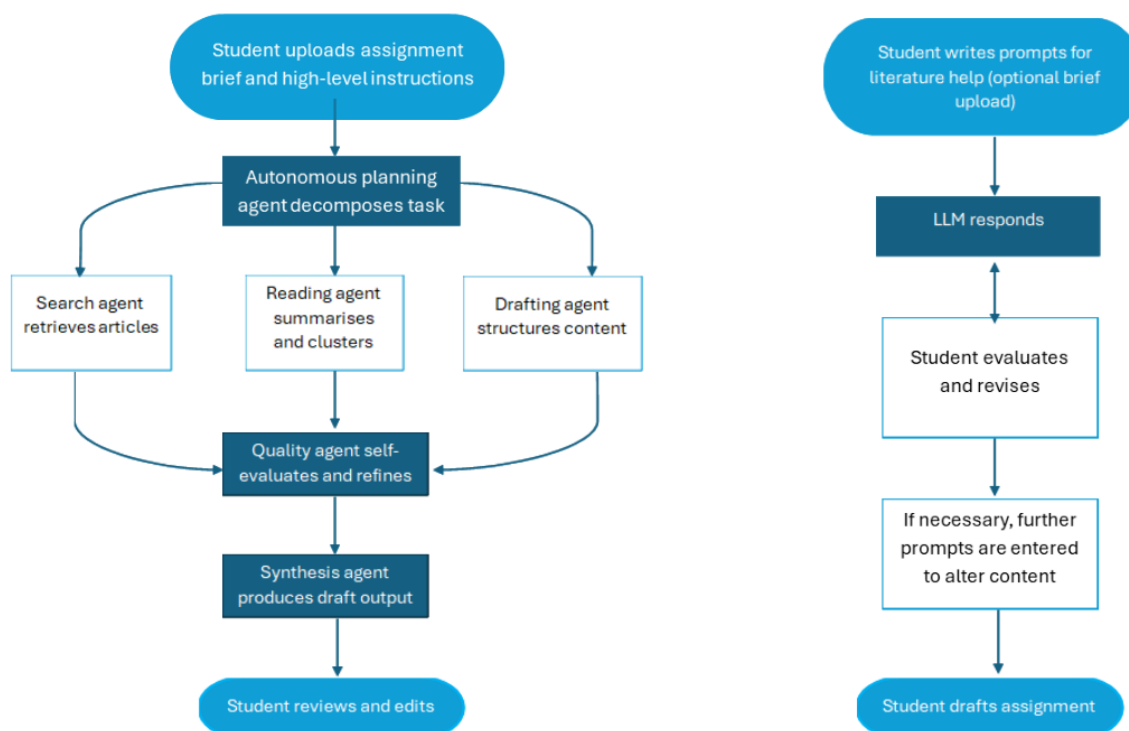
For years, educational AI took the form of supportive technologies (grammatical correctors, plagiarism detectors, or recommendation systems) that augmented, but did not displace, human cognition. With the emergence of generative AI, particularly large language models (LLMs), this relationship evolved. Now, agentic AI systems may represent a significant acceleration in the evolution of educational ontology. To clarify terminology: while generative AI tools like ChatGPT, Claude, and Gemini respond to individual prompts, agentic AI systems can plan, execute, and iterate across multiple steps with minimal human direction. For example, Microsoft's Copilot agents and application-specific agents in Word or PowerPoint represent early implementations of agentic capabilities: they can retrieve information, format documents, and coordinate tasks across applications based on high-level instructions. These systems are designed to initiate actions, make decisions, self-correct, and execute workflows, often without human oversight (Russell & Norvig, 2021; Somasegar, 2025). They can retrieve and synthesise information, interact across digital platforms, and engage in multi-step reasoning tasks. Current implementations of agentic AI in education demonstrate varying degrees of autonomy. Sakana AI's 'AI Scientist' (2024) claims to generate complete research papers with minimal human input, though independent reviews identify significant limitations (Verspoor, 2024). Commercial platforms like those described by Balaceanu (2023) deploy advising agents that handle student enrolment queries autonomously, while Microsoft's Copilot agents operate 'semi-autonomously', requiring human oversight of outputs. This variation underscores an important point: autonomy exists on a spectrum. Full autonomy



(where AI systems make decisions without human intervention) remains aspirational in most educational applications.

However, even semi-autonomous systems that aggregate information, draft responses, and suggest actions represent a qualitative shift from prompt-and-respond tools, as they reduce opportunities for student engagement with source material and intermediate reasoning steps. To illustrate the difference in educational terms: a student using ChatGPT for a dissertation might prompt it to suggest research questions, then separately request help drafting literature review sections, then seek data analysis guidance, with each step requiring student direction and integration. An agentic system could be given the assignment brief and, with minimal direction, propose topics, conduct literature searches, draft methodology chapters, and structure findings, potentially completing what should be months of learning in hours, with minimal student cognitive engagement (see Figure 1). The implications are not only operational but epistemic: what does it mean for learning when an AI system, rather than the student, can create original work and submit it with minimal oversight?

Figure 1. Single LLM tool versus agentic AI workflow in an academic assignment.



This further blurs foundational distinctions between assistance and substitution. When students upload an assignment brief to an agentic AI system and the system autonomously constructs a coherent, well-argued submission, we face a critical epistemological challenge:



whose learning is being assessed? Recent studies confirm that even experts can struggle to distinguish AI-generated academic texts from human-written ones (Floridi, 2023; Waltzer et al., 2024; Williams, 2024). The reliability of AI detectors, often proposed as a safeguard, is also in serious doubt. OpenAI, the creators of ChatGPT, have explicitly acknowledged that AI-generated text detectors are generally unreliable and prone to both false positives and false negatives (OpenAI, 2023) which still remains an issue (Vertu, 2025). It is worth noting that major AI providers are consciously maintaining 'human-in-the-loop' frameworks, positioning agents as collaborative rather than replacement tools. This design choice (whether driven by liability concerns, technical limitations, or pedagogical values) creates space for educational intervention but does not eliminate the challenge of distinguishing student from machine cognition in submitted work.

Whether this constitutes a fundamental shift or an intensification of existing challenges may be debated; what is clear is that the autonomous nature of these systems compresses cognitive processes in ways that demand pedagogical responses. Consequently, existing assessment models, rooted in assumptions of individual authorship and intellectual ownership, are being destabilised. The educational emphasis on original thought and academic voice now coexists uneasily with systems capable of convincingly simulating both.

The result is a profound tension: institutional frameworks that emphasise original thought and academic voice are confronting systems that can convincingly simulate both. This is not merely a threat to assessment integrity, the autonomy and sophistication of agentic AI systems hasten a reappraisal of how we define academic contribution, cognitive development, and intellectual labour. This presents a profound and urgent challenge to the structures and values underpinning higher education as well the critical skills students should develop.

Redefining academic literacies

Learning development in higher education has long centred on equipping students with a set of foundational competencies: the ability to read and write academically, to reason critically, to locate and evaluate scholarly sources, and to learn independently within disciplinary norms (Lea & Street, 1998; Wingate, 2006). However, the advent of agentic AI compels a reassessment of these literacies. Traditional approaches assumed a human author, whose



thinking could be evaluated through writing or other academic outputs. That assumption is increasingly incompatible with systems capable of simulating reasoning, analysis, and creativity.

Our framework diverges from existing AI literacy models by tailoring AI literacy to higher-education learning development. Unlike competency-based models such as Long and Magerko's (2020) or the technically or ethically focused frameworks reviewed by Panke (2025), it integrates AI literacy with academic literacies and foregrounds epistemic dimensions, positioning AI as a knowledge actor reshaping authorship and cognition. Developed through analysis of current pedagogical challenges and research in AI ethics, digital pedagogy, and academic literacies, it supports Learning Development practitioners in fostering students' epistemic agency, co-authorship ethics, metacognitive documentation, and disciplinary understanding of AI's limits. AI literacy must therefore be foregrounded as an essential educational goal. This literacy is not reducible to operational competence with AI systems. Rather, it constitutes a complex epistemic disposition: a capacity to engage with AI as both a partner in cognition through the co-creation of work and a subject of critical scrutiny. It entails several interrelated dimensions, explained below.

Critical evaluation remains foundational to academic literacy, but agentic AI fundamentally alters what students must evaluate. Traditional evaluation assumed human authorship with traceable expertise and institutional affiliations. Agentic systems generate outputs lacking these markers: there might be no author to assess, no peer review to trust, only probabilistic models synthesising patterns from vast datasets (AI-produced citations may be hallucinated, inaccurately linked, or stripped of interpretive context). Crucially, in the context of agentic AI, what matters is not only the presence of failures such as hallucinations presenting fabricated information confidently, invented yet correctly formatted citations, and structurally coherent arguments lacking genuine disciplinary understanding, but the way these can be orchestrated and scaled across an automated workflow. Agentic systems can propagate such errors across multiple linked tasks with minimal human oversight, reducing the points at which students would ordinarily exercise evaluative judgement.

Drawing on traditions in media literacy (Buckingham, 2015) and critical data studies (Mittelstadt & Floridi, 2016), learning developers should situate AI as a discursive actor whose outputs must be interrogated, not accepted at face value. To support this, students must also



learn to triangulate AI outputs with scholarly and domain-specific sources. This involves checking for fabricated citations, distinguishing between surface plausibility and evidence-based claims, and recognising disciplinary standards of evidence that AI may not fully capture. As AI systems increasingly simulate expert language, discernment becomes foundational to academic integrity.

Prompt literacy refers to students' ability to formulate effective, intentional queries that elicit meaningful responses from AI systems. Crafting a well-structured prompt is increasingly akin to formulating a research question or designing an experimental procedure. It requires an understanding of how LLMs interpret linguistic cues and how to anticipate and refine output based on iterative interaction. As Mollick and Mollick (2023) argue, prompting is not a neutral act but a form of intellectual shaping. Furthermore, prompt literacy must be situated within disciplinary conventions: prompting for a policy brief, lab report, or critical essay each require tailored strategies. Students should develop genre-aware prompting practices in 'dialogue' with AI that align with the expectations and epistemologies of their fields.

Collaborative co-authorship entails working with AI as a partner in the production of knowledge. This includes understanding when to rely on AI for drafting or ideation, how to integrate AI-generated material into human-authored work, and where to draw ethical and intellectual boundaries. Students must develop an awareness of the ways in which AI co-constructs meaning, while also affirming their own epistemic agency and accountability (Bozkurt, 2024). To support this, students should be encouraged to document the evolution of AI-human co-authored texts, including reflections on how AI was used, where human input shaped the outcome, and how boundaries between co-creation and authorship were maintained. Practical strategies include using documentation tables to record tool names, versions, tasks, and evaluative comments; appending sample statements or integrity declarations; and following discipline-specific or publisher guidelines for AI disclosure (see Weaver, 2024 for an example).

Understanding AI limits is essential in resisting the illusion of machine intelligence. Although systems like GPT-5 can generate human-like prose, they lack consciousness, intentionality, and embodied experience. These models offer increasingly personalised outputs, adjusting tone, style, and disciplinary register to user prompts, which can make their limitations and hallucinations harder to detect in ostensibly well-tailored academic text. Students must be



able to identify the limitations of AI in relation to nuance, ethics, and context. This literacy draws on cognitive science and critical AI studies (Bender et al., 2021; Chelli et al., 2024; Marcus & Davis, 2020), challenging the myth of AI infallibility. Literacy education should also include an understanding of phenomena such as hallucination, where AI systems fabricate information, citations, or factual content (Xu et al., 2024). Awareness of these risks equips students to verify and cross-reference outputs, ensuring that generated material does not mislead or misrepresent. Recognising the absence of lived experience in AI outputs also helps students critically assess the limitations of empathy, affect, and ethical reasoning in synthetic language.

Ethical AI use expands the scope of AI literacy beyond academic settings. It includes an awareness of the broader social consequences of AI systems, such as data surveillance, algorithmic bias, epistemic injustice and environmental impact (Birhane, 2021; Eubanks, 2018; van Uffelen et al., 2024). Students must learn not just how to use AI responsibly, but how to interrogate the power structures embedded in its design and deployment.

Together, these dimensions can form a robust framework for academic AI literacy. This is not a supplemental skill set; it is a redefinition of what it means to be academically literate in the 21st century. AI is now part of the apparatus through which knowledge is constructed (Ayyoub et al., 2025). Students must therefore be equipped not only to use AI but to critique its epistemological status and participate in shaping its educational integration.

Practical implications for learning developers

The redefinition of academic literacies in response to agentic AI is not a theoretical concern alone; it has immediate implications for practice. First, learning developers can lead efforts to support students in using AI transparently and reflectively. Rather than framing AI use as inherently dishonest or threatening, they can provide guidance on how to document and acknowledge AI contributions within academic work. This includes educating students on citation practices for AI-generated content, clarifying institutional policies, and challenging reductive understandings of plagiarism. Emerging research suggests that many students are uncertain about whether AI use constitutes misconduct, and misconceptions can lead to both misuse and underuse (Gonsalves, 2024). Learning developers can intervene to demystify these grey areas, by, for example, introducing a 'decision tree' for students to follow outlining



ethical AI use (Staufer & Gold, 2024), as well as providing case-study examples as a guide (Newcastle University, n.d.).

Second, there is a pressing need to reimagine assessment design. Traditional summative essays are increasingly vulnerable to AI replication. Learning developers can collaborate with faculty to co-design assessments that are more resilient to automation: for example, tasks that require students to critique AI outputs, document their reasoning process, or integrate experiential learning. Process-oriented and metacognitive assessments, including portfolios and reflective commentaries, offer productive alternatives. Such designs foreground student thinking over polished output and reward critical engagement with AI as a tool, rather than punishing its use.

Third, learning developers can take the lead in developing AI literacy initiatives that move beyond technical training. Workshops, online modules, and co-curricular resources should be developed to cultivate critical AI literacy across the curriculum. These should not only explain how to prompt or use AI tools effectively, but also engage students in reflecting on the limits, implications, and ethical dimensions of AI use.

Fourth, there is a vital role for learning developers in academic staff development. Many faculty are unsure how to incorporate AI into teaching and assessment, or remain focused on detection rather than design. Learning developers can facilitate structured dialogue about the pedagogical affordances of AI, offering models and case studies for integrating AI use into learning outcomes. This shift in mindset, from adversarial to exploratory, requires institutional support, collaboration, and sustained professional learning.

Fifth, a critical, yet often overlooked, dimension is equipping students and staff to recognise and manage AI hallucinations—false or fabricated outputs that LLMs produce (Xu et al., 2024). Learning developers can promote awareness of hallucination risks and foster skills to verify AI-generated content rigorously. This includes embedding strategies for cross-referencing outputs, using emerging detection tools critically, and maintaining an epistemic vigilance essential for academic rigour in an AI-mediated environment. Strategies to mitigate hallucinations include the use of tools that detect hallucinations (Saxena, n.d.) as well as introducing ‘hallucination hunts’ as a practice to embed vigilance in students (York University, 2025).



Finally, learning developers should have a voice in institutional AI policymaking. Learning developers can advocate for flexible, inclusive, and pedagogically informed AI policies that balance innovation with integrity. They can also ensure that policies do not inadvertently widen equity gaps by privileging students with greater technological fluency or access.

These practical implications affirm that learning developers are not peripheral to the AI conversation, they are central. Their pedagogical expertise, institutional insight, and proximity to both students and staff position them as critical actors in shaping a post-agentic AI higher education.

Acknowledgements

The authors would like to thank an anonymous reviewer for drawing attention to emerging security concerns surrounding the potential misuse of agentic AI systems, including the risk of malicious actors exploiting such technologies to hijack or manipulate digital systems. While these issues are increasingly being raised at the time of writing, their full implications are not yet well understood. A detailed exploration of the security dimensions of agentic AI lies beyond the scope of this paper and is more appropriately addressed through dedicated technical, institutional, and IT governance frameworks/departments.

Disclosure statement

The authors used the following AI tool in the preparation of this manuscript: ChatGPT-4. The manuscript was originally written by the authors in full. ChatGPT-4 was subsequently used to refine phrasing, improve clarity, and enhance the overall readability of the text. No content, argumentation, or structure was generated by the AI, and all ideas remain the intellectual work of the authors. The authors have complied with the JLDHE's principles of AI use.

Funding

The authors received no funding for this research and declare no competing interests.

References

Ayyoub, A. M., Khlaif, Z. N., Shamali, M., Eideh, B. A., Assali, A., Hattab, M., Barham, K. A., & Bsharat, T. R. K. (2025). Advancing higher education with GenAI: factors influencing



educator AI literacy. *Frontiers in Education*, 10, Article 1530721.

<https://doi.org/10.3389/feduc.2025.1530721>

Balaceanu, D. (2023, February 23). Agentic AI paves the way for a new dawn in higher education. *Druid AI*. <https://www.druidai.com/blog/agentic-ai-paves-the-way-for-a-new-dawn-in-higher-education>

Belcic, I., & Stryker, C. (2025). AI agents in 2025: expectations vs. reality. *IBM Think*. <https://www.ibm.com/think/insights/ai-agents-2025-expectations-vs-reality#:~:text=%E2%80%9CIBM%20and%20Morning%20Consult%20did,declaration%20is%20not%20without%20nuance>

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big?. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, Canada*, 610–623. <https://doi.org/10.1145/3442188.3445922>

Birhane, A. (2021). Algorithmic injustice: a relational ethics approach. *Patterns*, 2(2), Article 100205. <https://doi.org/10.1016/j.patter.2021.100205>

Bozkurt, A. (2024). GenAI et al.: cocreation, authorship, ownership, academic ethics and integrity in a time of generative AI. *Open Praxis*, 16(1), Article 654. <https://openpraxis.org/articles/10.55982/openpraxis.16.1.654>

Buckingham, D. (2015, May 21). The blanding of media literacy. *David Buckingham Blog*. <https://davidbuckingham.net/2015/05/21/the-blanding-of-media-literacy/>

Chelli, M., Descamps, J., Lavoué, V., Trojani, C., Azar, M., Deckert, M., Raynier, J-L., Clowez, G., Boileau, P., & Ruetsch-Chelli, C. (2024). Hallucination rates and reference accuracy of ChatGPT and Bard for systematic reviews: comparative analysis. *Journal of Medical Internet Research*, 26, Article e53164. <https://doi.org/10.2196/53164>

Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.

Floridi, L. (2023). Norms for academic writing in the era of advanced artificial intelligence. *Digital Society*, 2, Article 48. <https://link.springer.com/article/10.1007/s44206-023-00079-7>

Gonsalves, C. (2024). Addressing student non-compliance in AI use declarations: Implications for academic integrity and assessment in higher education. *Assessment & Evaluation in Higher Education*, 50(4), 592–606. <https://doi.org/10.1080/02602938.2024.2415654>



- Lea, M. R., & Street, B. V. (1998). Student writing in higher education: an academic literacies approach. *Studies in Higher Education*, 23(2), 157–172.
<https://doi.org/10.1080/03075079812331380364>
- Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, USA*, 1–16.
<https://doi.org/10.1145/3313831.3376727>
- Marcus, G., & Davis, E. (2020). *Rebooting AI: Building artificial intelligence we can trust*. Penguin.
- Mollick, E. R., & Mollick, L. (2023). Using AI to implement effective teaching strategies in classrooms: Five strategies, including prompts. *The Wharton School Research Paper*, 1–26
<http://dx.doi.org/10.2139/ssrn.4391243>
- Moquin, S. (2025, May 8). Five real use cases of agentic AI in higher ed that aren't just fancy buzzwords. *Enrollify*. <https://www.enrollify.org/blog/5-real-use-cases-of-agentic-ai-in-higher-ed-that-arent-just-fancy-buzzwords>
- Newcastle University. (n.d.). *Acknowledging use of AI*. <https://www.ncl.ac.uk/academic-skills-kit/good-academic-practice/artificial-intelligence/acknowledging/>
- OpenAI. (2023, January 31). *New AI classifier for indicating AI-written text*.
<https://openai.com/blog/new-ai-classifier-for-indicating-ai-written-text>
- Panke, S. B. (2025). *GenAI literacy: What is it, and how should we teach it? Frameworks, reviews, and approaches*. AACE Review. <https://aace.org/review/genai-literacy-what-is-it-and-how-should-we-teach-it-frameworks-reviews-approaches/>
- Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- Sakana AI. (2024, August 13). *The AI scientist: Towards fully automated open-ended scientific discovery*. <https://sakana.ai/ai-scientist/>
- Saxena, C. (n.d.). Top tools and plugins to detect AI hallucinations in real-time. *ISHIR Blog*.
<https://www.ishir.com/blog/183214/top-tools-and-plugins-to-detect-ai-hallucinations-in-real-time.htm>
- Somasegar, S. (Host). (2025, 24 April). Reinventing recruiting: SeekOut's Anoop Gupta on the rise of agentic AI. [Audio podcast episode]. In *Founded & Funded podcast*. Madrona Venture Group. <https://www.madrona.com/rise-of-agentic-ai-in-recruiting-seekout-anoop-gupta/>
- Stauffer, J., & Gold, J. (2024, July 10). A decision tree to guide student AI use. *Edutopia*.
<https://www.edutopia.org/article/student-use-ai-helpful-framework/>



PwC. (2025, June 3). *The fearless future: 2025 global AI jobs barometer*.

<https://www.pwc.com/gx/en/services/ai/ai-jobs-barometer.html>

van Uffelen, N., Lauwaert, L., Coeckelbergh, M., & Kudina, O. (2024). Towards an environmental ethics of artificial intelligence. *arXiv*, Article 2501.10390.

<https://doi.org/10.48550/arXiv.2501.10390>

Vertu. (2025, June 22). *Is OpenAI's AI detector accurate enough for 2025*. <https://vertu.com/ai-tools/openai-ai-detector-accuracy-features-user-experiences-2025/>

Vespoor, K. (2024, August 20). A new AI scientist can write science papers without any human input—here's why that's a problem. *The Conversation*. <https://theconversation.com/a-new-ai-scientist-can-write-science-papers-without-any-human-input-heres-why-thats-a-problem-237029>

Weaver, K. D. (2024). The artificial intelligence disclosure (AID) framework: An introduction. *College & Research Libraries News*, 85(10), 407–411. <https://doi.org/10.5860/crln.85.10.407>

Waltzer, T., Pilegard, C., & Heyman, G. D. (2024). Can you spot the bot? Identifying AI-generated writing in college essays. *International Journal for Educational Integrity*, 20, Article 11. <https://doi.org/10.1007/s40979-024-00158-3>

Williams, R. T. (2024). The ethical implications of using generative chatbots in higher education. *Frontiers in Education*, 8, Article 1331607.

<https://doi.org/10.3389/feduc.2023.1331607>

Wingate, U. (2006). Doing away with 'study skills'. *Teaching in Higher Education*, 11(4), 457–469.

<https://doi.org/10.1080/13562510600874268>

Xu, Z., Jain, S., & Kankanhalli, M. (2024). Hallucination is inevitable: An innate limitation of large language models. *arXiv*, Article 2401.11817. <https://doi.org/10.48550/arXiv.2401.11817>

York University. (2025). *Artificial Intelligence (AI) pedagogies resources*.

<https://www.yorku.ca/teachingcommons/resources-by-topic/artificial-intelligence-ai-pedagogical-resources/>

Author details

Octavio Murekian is a Lecturer in Marketing at the University of Greenwich and Programme Leader for the MA in Marketing Management. His research focuses on consumer culture and behaviour, with particular emphasis on online communities, status, technology, and security-related consumption. His teaching interests include consumer behaviour, marketing research,



digital marketing, and the use of technology to enhance learning and teaching in higher education.

Surabil Sudarshan is a Senior Teaching Fellow at the University of Greenwich Business School with over 20 years of industry experience spanning technology, business, and consulting. He designs and leads industry-focused modules and has contributed extensively to undergraduate and postgraduate programme development. His expertise includes artificial intelligence, particularly AI agents and agentic AI, data analytics, digital marketing, and business technology consulting. His research examines responsible AI, consumer trust, content visibility, and human–AI collaboration in educational and applied contexts. His teaching focuses on AI, data-driven digital marketing, analytics, innovation, and technology-enhanced learning.

ORCID

Octavio Murekian  <https://orcid.org/0009-0001-8437-2682>

Surabil Sudarshan  <https://orcid.org/0009-0001-7480-8630>

