Cut the bullshit: Why GenAI systems are neither collaborators nor tutors

Gene Flenady and Robert Sparrow

Working paper only.

A version of this paper appeared as:

Flenady, G., and Sparrow, R. 2025. Cut the Bullshit: Why GenAI Systems are neither Collaborators nor Tutors. *Teaching In Higher Education*. Published online: 17 May 2025. https://doi.org/10.1080/13562517.2025.2497263.

Please cite that version.

Abstract: The rapid development of Generative AI (GenAI) technologies has led to widespread endorsement of GenAI systems serving as a "personal tutor" and learning "collaborator" in higher education. However, because GenAI outputs are prone to "hallucinations," it has been suggested that students take responsibility for the accuracy of GenAI contributions to their learning. We rehabilitate Plato's scepticism regarding writing and draw on Harry Frankfurt's analysis of "bullshit" to demonstrate that GenAI systems are constitutively epistemically irresponsible. We argue that the expectation on tertiary students to assume responsibility for their so-called "tutors" and "collaborators" is pedagogically perverse, amounting to a demand that students take sole responsibility for the accuracy of claims they are not able to properly assess. Moreover, to the extent that GenAI teaching systems replace students' interaction with human teachers, it will be increasingly difficult for students to develop the skills and motivation to hold GenAI outputs to disciplinary standards.

Keywords: Generative AI, responsibility, agency, bullshit, student engagement

Cut the bullshit: Why GenAI systems are neither collaborators nor tutors

Gene Flenady and Robert Sparrow

The recent rapid development of Generative AI (GenAI) technologies has to the possibility of GenAI systems serving as "personal tutor," learning "collaborator," and "dynamic assessor" for tertiary students (Sabzalieva and Valentini 2023). Responses to this possibility might be separated into three camps. At one extreme, the new technology has been championed as the revolutionary future of higher education, in which a hugely expanded worldwide number of tertiary students will learn predominantly through online interactions with GenAI (Khan 2024; Mollick 2024). In the middle, more cautious accounts acknowledge ethical issues with widespread integration of GenAI in education, proposing updated teaching practices and institutional regulations to ethically harness the new technology's potential (Cotton et al 2024; Hartley et al 2024; Fuchs 2023; Kanesci et al 2023; Rudolph, Tan and Tan 2023). A component of this centrist position is that universities should move to accept work that is the result of student-AI "collaboration," but, given the tendency for GenAI to "hallucinate" claims and sources, students should in some way take responsibility for the accuracy of GenAI outputs (Bearman and Ajjawi 2023; Lodge et al 2023; Mollick and Mollick 2023). At the other end of the spectrum, what has been termed a "dystopia is now" strand of academic discourse (Bearman, Ryan and Ajjawi 2022) has expressed fears that GenAI will, for example, deepen rather than relieve the "epidemic" of academic overwork and precarity (Watermeyer et al 2023), and ultimately destroy the reading and writing skills on which the humanities depend (Freiberg 2024).

One common response to scepticism about GenAI's educational use is to raise Plato's famous fears about the emergence of a technology relatively new in his time: writing (Tufekci 2022; Kobb 2024; Plate, Melick and Hutson 2024). It is suggested that, just as Plato's fears about the corrosive effects of writing proved to be unfounded, opposition to the integration of GenAI into higher education will not stand the test of time. The not-so-subtle subtext is that more full-throated criticisms of AI teaching systems come from an irrational fear of change, rather than genuine engagement with and understanding of the new technology and its potential.

In this point of departure, we argue that Plato's routinely derided fears about writing instead illuminate why the expectation that tertiary students take responsibility for the outputs of AI systems is highly problematic. In our view, the implications of the demand for student responsibility are ethically and pedagogically *perverse*. We cannot reasonably expect students to take responsibility for their so-called "collaborator" and "tutor" because GenAI is fundamentally unable to assume responsibility for itself or others. Drawing on recent accounts of GenAI outputs as irresponsible "bullshit," we argue that commentators should in turn "cut the bullshit" and not refer to GenAI in higher education as a collaborator and tutor – words that falsely imply that GenAI is a responsible agent.

I. Plato's Fear

Why *exactly* was Plato "afraid" of writing? Plato, discussing this matter in the *Phaedrus* through the figure of Socrates, has two concerns.

First, Plato is worried that writing will "atrophy" memory (Plato 2003, 275a). In a sense, this fear has proven to be well-founded. Those trained in oral cultures have, from the perspective of written cultures, seemingly superhuman powers of recall. For better or worse, written cultures have traded off these mnemonic capacities for the archival forms of memory that writing makes possible.¹

Second, and more importantly for our argument here, Plato was concerned with the way that words, once written down, escape their author's control, and can be interpreted in any number of ways, without their author being present to defend their intended meaning. As Socrates puts it:

"Once an account has been written down, you find it all over the place, hobnobbing with completely inappropriate people no less than with those who understand it, and completely failing to know who it should and shouldn't talk to. And faced with rudeness and unfair abuse it always needs its father [its author] to come to its assistance, since it is incapable of defending or helping itself" (Plato 2003, 275e).

In our view, Plato's ultimate concern in this passage is with *responsibility:* once written down, it becomes much harder to identify who is *responsible to and for* the meaning of a claim, and

¹ The relation between written and oral social memory is complex, and we cannot consider it here. For an overview of debates in the literature see Jones and Russell (2012).

– equally important for Plato – who is responsible for the *effects* that such a claim might have on those unable to properly contextualise and interpret its content. Plato's response was to restrict his "real" doctrine to an oral conversation shared between the select few: philosophers present to and thus directly responsible for the meaning of their words, and those "initiates" (citizens, always male) privileged with the time and opportunity to learn to take those words in the right sense (Hadot 2002).

The modern academy is of course very different to Plato's. Even though elite universities demonstrably privilege those from certain race and class backgrounds (Bhopal and Myers 2023), they are at least in principle egalitarian and meritocratic. Equally obviously, the modern academy depends on writing and written archives. Nonetheless, tertiary institutions remain acutely sensitive to Plato's second worry: universities work to redress the "defencelessness" of words by ensuring that written claims are properly connected to those responsible for them via rigorous practices of attribution: we know who the authors are, we know what institutions they belong to and represent, and we know how to contact them. Whether or not we in fact get a response from a given author when we attempt to engage them in written argument, we know who is *ultimately* responsible for defending a work, and to whom our response to that work ought to be directed. At the same time, we work to reduce the "hobnobbing" of written words by ensuring those who engage with academic texts have either been trained, or are in training, to possess discipline-specific skills that allow them to contextualise, interpret, and assess the text in question. Together these two features of the modern academy - named authorship and attribution combined with disciplinary-specific training in the assessment of claims - have worked to neutralise Plato's fears. However, that we work as conspicuously as we do to neutralize Plato's worry suggests that his concerns remain relevant.

In the age of the internet, of course, the impact of "lies" "fake news" and "bullshit" on public opinion is a serious concern.² Here again Plato's second worry seems legitimate. In the case of internet sources, we do not know *who* or *where* the authors reside, who their words are intended for, or even – increasingly – if they are human. Thus, it has become even more important to inculcate university students with the capacity to distinguish between those works that can be

² MacKenzie and Bhatt (2019, 2020) raise epistemological concerns about the impact of "lies, bullshit, and fake news" on tertiary education; the authors draw, as we do here, on Harry Frankfurt's definition of bullshit. Bhatt and Mackenzie (2019) set out epistemological and pedagogical dangers of encouraging students to "just Google it!" without critical awareness of exogenous factors shaping search engine results – dangers we would suggest are also present in encouraging students to rely on the outputs of GenAl systems.

sourced to a responsible author located within a discipline and an institution and those dangerous words for which authors do not take responsibility.

In our view, those who respond to critics of AI in education by pointing to Plato's fears regarding writing unwittingly draw attention to the ethical heart of the matter. The introduction of GenAI teaching systems is *irresponsible* in the sense that it introduces into the academy words without an author – words for which no one is responsible (Bingham 2024).³ This, we suggest, has the perverse outcome of demanding *students* take responsibility for discourse that they are not yet in a position to properly assess. Moreover, to the extent that GenAI teaching systems substantially replace students' interaction with real human teachers in the future – very much a possibility (Selwyn 2019) – it will become increasingly difficult for students to develop the skills to assess GenAI outputs, just as students become increasingly divorced from the human relationships that give their own academic responsibilities meaning.

II. Irresponsible Bullshit

Harry Frankfurt famously defined bullshit as "lack of concern with truth, or an indifference to how things really are" (Frankfurt 2002, 340; Frankfurt 2005). Where most people tell the truth, and liars pay attention to the truth to shape their lies, the bullshitter is only interested in provoking a response from their audience and doesn't care whether what they say is true or not.

Sparrow, Koplin and Flenady (2023) and Hicks, Humphries and Slater (2024) have argued that the outputs of Generative AIs like ChatGPT are bullshit in Frankfurt's precise sense. GenAI is neither designed to nor is capable of truthfully representing the world; instead, "they are designed to convey convincing lines of text" (Hicks, Humphries and Slater 2024, 38). That is, GenAI is designed to produce a particular reaction in its users – to *convince* its users of the usefulness of the response – rather than to say something true of the world. They are, then, fundamentally bullshitters.⁴

³ In line with the view presented here, Bingham argues that GenAI outputs are Plato's fears about writing "come to life," insofar as such outputs lack an author who might defend them (2024: 41). Unlike us, however, Bingham follows Jacques Derrida's deconstructive reading of the *Phaedrus* in *Dissemination* (1981) by claiming we ought to "celebrate" the educational and philosophical opportunities of the breakdown of the distinction between properly authored words and the "unnatural" outputs of GenAI. ⁴ One might object that GenAI systems, lacking intention, cannot be characterised as producing "bullshit" in Frankfurt's sense. While a human agent bullshits in order to achieve some goal, GenAI systems merely probabilistically "stitch together" sentences (Bender et al 2021). Hicks, Humphries and

Importantly, to say AI outputs are bullshit is not to claim that they are *always* untrue, nor to say that they can play no role in the justification of belief. A bullshitter might in certain cases achieve their rhetorical goals via the utterance of true claims. Similarly, AI might – perhaps even for the most part – produce outputs that reliably correspond with states of affairs in the world.

That the outputs of GenAI constitute "bullshit" should be enough, we think, to at least dampen the enthusiasm for their widespread integration in tertiary education. However, the problem of the introduction of GenAI in educational contexts is deeper still, for reasons that emerge once we connect the definition of GenAI as "bullshit" with the discussion of responsibility for discourse above.

It is widely believed that machines cannot be moral agents: that is, machines cannot be *responsible* for their acts (Johnson 2006; Hakli & Mäkelä 2019; Sparrow 2021; Véliz 2021). It makes no sense, for example, to send a military drone to jail for wrongly killing civilians; responsibility for such an action must be assigned to some human agent. The ethical issue in such cases is the difficulty in identifying the human beings who ought to "fill in" responsibility for the machine, giving rise to the notorious "responsibility gap" (Sparrow 2007).

Importantly for our discussion of machines in education, the making of an epistemically contentful claim – that is, giving "testimony" about some state of affairs – is itself a kind of action for which the individual is responsible.⁵ Giving testimony implicitly commits the testifier to providing further reasons as justification for their belief in the truth of their report. The testifier is committed to defending their testimony, "standing behind it" to "back it up," when requested, with further reasons. Individuals may fail to meet this obligation, but their act is intelligible as an act of testimony on condition that some responsibility to meet it is assumed.

If giving testimony is an act that entails the assumption of responsibility by a moral agent, and if machines are not moral agents, it follows that machines cannot give testimony.

Now, when a human bullshitter gives testimony that such-and-such is the case, while they do not take seriously or indeed entirely disregard the implicit responsibility to defend their claims,

Slater anticipate this objection, arguing that GenAI systems can be meaningfully viewed as "having intention" insofar as they are "designed to give the impression of concern with truth" (2024: 32). ⁵ In pursuing this line of argument we are directly influenced by Robert Brandom's pragmatist semantics (1994, 2019). See Heinrichs and Knell (2021) and Sparrow and Flenady (2025a) for Brandomian arguments that Generative AI is not a responsible agent and so does not make genuinely contentful knowledge claims.

they nonetheless can be held responsible for them. Part of the motivation, then, for a philosophical account of the concept of bullshit is to enable individuals with a commitment to the truth to "call out" bullshit, that is, to highlight the way in which the bullshitter does not take their own epistemic responsibilities seriously, and, if necessary, to demand that they "cut the bullshit," so to speak.

In the case of a GenAI system, however, the problem of the irresponsibility of bullshit is much more significant: GenAI is not an agent that can even potentially be held accountable. As in the "responsibility gap," the problem of a "testimony gap" arises: the problem of appropriately identifying those human agents to be held accountable for the activity of machines. A "testimony gap" opens where GenAI systems are taken – despite the limitations of their design – by human agents to be reporting on the world, but cannot be held responsible for mistaken or otherwise epistemically inadequate reports.⁶

III. Perverse Outcomes

The response in the literature, one which tacitly accepts the irresponsibility of GenAI outputs, has been to propose *outsourcing* responsibility from GenAI to students. There are naïve and considered versions of this view.

The naïve version insists that the need to take responsibility for GenAI systems as liable to err or "hallucinate." Students are responsible for error checking their so-called "collaborators" and "tutors" (Mollick and Mollick 2023). There are two problems with this.

Firstly, this view fundamentally misunderstands GenAI outputs as reporting on the world, with the suggestion that AI systems are mostly – and will become more – accurate representations of true states of affairs. As discussed above, that GenAI outputs are bullshit is not an observation about their accuracy: to bullshit is to be unconcerned with the truth, and accurate reports can serve as bullshit. As such, "hallucination" is a misleading term: it gives the impression that GenAI *wants* to report on the world, but suffers some breakdown in that reporting (Hicks, Humphries and Slater 2024).

Secondly, this way of assigning responsibility has the perverse outcome of making individual students *entirely* responsible for the truth of the claims in the work they submit, in place of

⁶ For a detailed account of AI outputs as giving rise to a "testimony gap," see Sparrow and Flenady (2025b).

genuine collaborator and teacher-student relationships that are characterised by *mutual* responsibility to the truth.

Calling GenAI a "collaborator" disregards accepted understandings of that term: to collaborate is to work together with another agent who is responsible for their contributions, who is capable of holding themselves and others to collective standards governing a shared project. There is significant higher education research on collaborative learning as developing students' capacity to constructively provide and respond to feedback and to monitor their own and others' contributions in light of shared goals (Slavin 1983; Hargreaves 2007; Johnson et al 2007). GenAI, however, is unable to take responsibility for itself or for others in this way; it cannot collaborate, and calling it a "collaborator" drains actual student collaboration of its meaning and pedagogical significance. Enthusiasm for GenAI as a "collaborator" is stubbornly unconcerned with both the reality of GenAI and of genuine collaboration. It is, in a word, bullshit.

Similarly, teaching is constituted by a relationship between individuals. In a properly constituted teacher-student relationship, responsibility is bidirectional: university students ought to hold their tutors and lecturers to account, questioning their claims and the justification for them, just as teaching staff ought to give ongoing feedback on students' claims. That is, arguably, what training in the humanities is: learning to contest the knowledge claims of others and to vindicate one's own. There is, again, significant higher education research on effective feedback as a dialogic, co-constructive process, in which teachers and students take responsibility for respectfully but critically engaging with one another's claims to knowledge (Carless 2016; Carless and Winstone 2023). Introducing GenAI as a "tutor" inserts a fundamentally irresponsible agent - an agent who cannot even *potentially* assume responsibility – into the co-constructive teaching and learning space, corroding the mutual responsibility that makes teaching and feedback effective. Yet the perverse response to tacit acknowledgement of the irresponsibility of GenAI tutoring systems is to expect students to assume responsibility for their "tutor." This at the very least reverses the common assumption that pedagogical responsibilities flow, at least initially, from teacher to student: it is the responsibility of the teacher to make possible a co-constructive site of mutual responsibility between teacher and students, by modelling "care respect" (Zhou et al 2021) and foregrounding their own vulnerability to feedback (Carless and Boud 2018; Carless and Winstone 2023). The assignation of responsibility to students for their GenAI "tutors" should indicate that GenAI

systems are not *really* tutors at all; they can only be called as much by disregarding the reality of GenAI and of genuine teaching. It is, again, bullshit.

Considered accounts of how students might accept responsibility for the work they produce using GenAI argue for the need to go beyond error checking and assess GenAI outputs according to disciplinary standards, which commendably seeks to distribute responsibility to students *and* teachers (Bearman and Ajjawi 2023). Teachers have a responsibility to inculcate students with the discipline-specific skills to properly assess the relevance and value of discourse, including that generated by, or with, AI. The limitation of this approach, however, appears once we consider the nature of pedagogical responsibility as holding mutually between teachers and students.

As students, one important way we come to a sense of responsibility to disciplinary standards is *through* relationships with real representatives of those standards, that is, by addressing ourselves in speech and writing to teachers who serve as models of those standards for us. We come to identify with and invest in a discipline because representatives of that discipline have invested in us in some meaningful way. In other words, we accept responsibility to a discipline and its standards in the context of mutual responsibility, being held by another to normative standards, and then coming to hold others to that standard in turn. In a future where GenAI tutors and grading systems substantially replace or mediate pedagogical relationships this felt sense of responsibility to another is highly likely to be corroded. More concretely, if students come to feel they are writing with and for a machine rather than with and for human beings if their education emphasises prompting and revising GenAI outputs, with submissions in turn to be marked by GenAI – it is plausible to suggest that they will not feel responsible to representatives of disciplinary standards, and thus not to disciplinary standards at all. Asking students to critically interrogate GenAI outputs in line with disciplinary standards thus faces a significant conceptual problem: we may come to ask students to exercise responsibility to a standard outside of the human contexts in which those standards have meaningful grip.

This might be put in terms of student engagement. Although the higher education literature on engagement is complex,⁷ there is evidence that student engagement is significantly influenced by teacher behaviour (Umbach and Wawrzynski 2005), and that students are more likely to engage when teachers are prepared, approachable, set high standards and challenge students

⁷ For an overview of the history of the engagement literature, see the introduction to Bowden, Tickle and Naumann (2021).

(Mearns, Meyer and Bharadwaj 2007, Bryson and Hand 2007). It may be hard to imagine a GenAI tutoring system promoting engagement in this way, but proponents of GenAI have predicted the new technology will in fact increase student engagement (Fuchs 2023). However, the evidence available thus far is ambivalent, with GenAI in some cases promoting *dis*engagement (Lo, Hew and Jong 2024). Indeed, there is already evidence that students are more likely to rely uncritically on ChatGPT – effectively using it to cheat – in cases where they feel "degree apathy," that is, where they do not feel invested or engaged in their education (Playfoot, Quigley and Thomas 2024). It is plausible then, on current evidence, to predict a vicious circle: substantial introduction of GenAI systems into higher education will make it less likely for students to critically engage with GenAI outputs, which is just what students are being asked to do to justify the widespread introduction of GenAI into tertiary education in the first place.

Conclusion

The "responsibilization" of students for the outputs of the AI they use fits neatly with the neoliberal model of the student as consumer (Molesworth, Nixon and Scullion 2009). It fits much less neatly with the University student understood as a member of a scholarly community. "University," we should remember, is drawn from the Latin for "community of teachers and scholars" [*Universitas Magistrorum et Scholarium*]. The University, on this conception, is made up of individuals responsible to one another for their knowledge claims. We believe that GenAI threatens the University community insofar as it represents an extreme form of the epistemic irresponsibility that motivated Plato's famous criticisms of writing. GenAI is a constitutively irresponsible source of knowledge claims insofar as GenAI outputs have no author to answer for them. GenAI threatens to undo the work that community does to ensure that a claim can be properly traced to an agent that can answer for and defend it.

Part of the solution here is to cut the bullshit: GenAI is constitutively unable to serve as a collaborator or a tutor and should not be said to do so, if the meaning of those roles – and the responsibilities those roles entail – is not to be corroded. To be clear, our goal here has not been to argue for a total ban on the use of GenAI systems in higher education – such a ban is neither possible nor desirable – but rather to incite discussion about our language for and conceptualisation of GenAI teaching systems in higher education, in the interests of robust regulatory oversight and pedagogically sound integration.

References

Bearman, M. and R. Ajjawi. 2023. Learning to work with the black box. Pedagogy for a world with artificial intelligence. *British Journal of Educational Technology* 54: 1160-1173.

Bearman, M., Ryan, J., and R. Ajjawi. 2022. Discourses of artificial intelligence in higher education: A critical literature review. *Higher Education* 86: 369–385.

Bender, E., Gebru, T., McMillan-Major, A., and S. Schmitchell. 2021. On the dangers of stochastic parrots: can language models be too big? *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 610-623.

Bowden, J., Tickle, L. and K. Naumann. 2021. The four pillars of tertiary student engagement and success: a holistic measurement approach. *Studies in Higher Education* 46 (6): 1207-1224,

Brandom, R. 1994. *Making it explicit: reasoning, representing, and discursive commitment.* Cambridge, MA: Harvard University Press.

Brandom, R. 2019. *A spirit of trust: a reading of Hegel's Phenomenology*. Cambridge, MA: Belknap.

Bryson, C. and L. Hand. 2007. The role of engagement in inspiring teaching and learning. *Innovations in Education and Teaching International* 44 (4): 349-362.

Carless, D. 2016. Feedback as dialogue. *Encyclopedia of educational philosophy and theory*. Edited by M. A. Peters. Singapore: Springer.

Carless, D. and D. Boud. 2018. The development of student feedback literacy: enabling uptake of feedback. *Assessment & Evaluation in Higher Education* 43(8): 1315-1325.

Carless, D., and N. Winstone. 2023. Teacher Feedback Literacy and Its Interplay with Student Feedback Literacy. *Teaching in Higher Education* 28(1): 150-163.

Cotton, D. R. E., Cotton, P. A., and J. R. Shipway. 2024. Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International* 61 (2): 228–239.

Bingham, C. 2024. Education and Artificial Intelligence at the Scene of Writing: A Derridean Consideration. *Futurity Philosophy* 3 (4): 34-46.

Bhopal, K. and M. Myers. 2023. *Elite Universities and the Makin of Privilege: Exploring Race and Class in Global Educational Economies*. London: Routledge.

Bhatt, I. and A. MacKenzie. (2019). Just Google it! Digital Literacy and the epistemology of Ignorance. *Teaching in Higher Education* 24 (3): 302-317.

Derrida, J. 1981. *Dissemination*. Translated by B. Johnson. Chicago: University of Chicago Press.

Frankfurt, H. 2002. "Reply to Cohen." In *The Contours of Agency: Essays on themes from Harry Frankfurt*, edited by S. Buss and L. Overton. Cambridge, MA: MIT Press.

Frankfurt, H. 2005. On Bullshit. Princeton: Princeton University Press.

Freiberg, C. 2024. Generative AI and the necessity of an existential crisis for the liberal arts. *Educational Philosophy and Theory*, October: 1-11.

Hadot, P. 2002. *What is ancient philosophy?* Translated by M. Chase. Cambridge, MA: Harvard University Press.

Hakli, R. and P. Mäkelä. 2019. Moral responsibility of robots and hybrid agents. *The Monist* 102 (7): 259–275.

Hargreaves, E. 2007. The validity of collaborative assessment for learning. Assessment in Education: Principles, Policy & Practice, 14 (2): 185–199.

Hartley, P., Beckingham, S., Lawrence, J., and S. Powell. 2024. Using generative AI effectively in higher education. In *Using Generative AI Effectively in Higher Education: Sustainable and Ethical Practices for Learning, Teaching and Assessment,* edited by Beckingham, S., Lawrence, J., Powell, S. & Hartley, P. London: Routledge.

Heinrichs, B. and S. Knell. 2021. Aliens in the Space of Reasons? On the interaction between humans and artificial intelligence agents. *Philosophy and Technology* 34 (4): 1569-1580.

Hicks, M. T., Humphries, J. and J. Slater. 2024. ChatGPT is bullshit. *Ethics and Information Technology* 26, 38.

Johnson, D. G. 2006. Computer systems: Moral entities but not moral agents. *Ethics and Information Technology* 8 (4): 195-204.

Johnson, D.W., Johnson, R.T. and K. Smith. 2007. The state of cooperative learning in postsecondary and professional settings. *Educational psychology review* 19: 15-29

Jones, S. and L. Russell. 2012. Archaeology, Memory and Oral Tradition: An Introduction. *International Journal of Historical Archaeology* 16: 267-283.

Kobb, K. 2024. From Plato to AI: Are we losing our minds? *Resilience*. https://www.resilience.org/stories/2024-09-08/from-plato-to-ai-are-we-losing-our-minds/

Lo, C. K., Hew, K. F., and M. S. Jong 2024. The influence of ChatGPT on student engagement: A systematic review and future research agenda. *Computers and Education* 219, no. C.

Lodge, J. M., Suijing Yang, Furze, L. and P. Dawson 2023. It's not like a calculator, so what is the relationship between learners and generative artificial intelligence? *Learning: Research and Practice* 9 (2): 117-124.

Mearns, K., Meyer, J. and A. Bharadwaj. 2007. "Student engagement in human biology practical sessions." Paper presented at the 16th Annual Teaching and Learning Forum, Curtin University of Technology, Perth, Australia, January 2007.

Molesworth, M., Nixon, E. and R. Scullion. 2009. Having, Being and Higher Education: The Marketisation of the University and the Transformation of the Student into Consumer. *Teaching in Higher Education* 14 (3): 277–87.

Mollick, E. 2024. Co-Intelligence: Living and Working with AI. London: Penguin.

Mollick, E., & Mollick, L. 2023. Assigning AI: Seven approaches for students, with prompts. *SSRN*. https://ssrn.com/abstract=4475995

Plate, D. Melick, E. and J. Hutson. 2024. Pedagogical Foundations of AI Integration. *Faculty Scholarship* 671.

Plato. 2003. Phaedrus. Translated by R. Waterfield. Oxford: Oxford University Press.

Playfoot, D., Quigley, M. and A. Thomas. 2024. Hey ChatGPT, give me a title for a paper about degree apathy and student use of AI for assignment writing. *The Internet and Higher Education*. 62, 100950.

Rudolph, J., Tan, S. and S. Tan. 2023. ChatGPT: Bullshit spewer or end of traditional assessments in higher education? *Journal of Applied Teaching and Learning* 6 (1): 342-363.

Sabzalieva, E. and A. Valentini. 2023. ChatGPT and artificial intelligence in high-er education: Quick start guide. *United Nations Educational, Scientific and Cultural Organization*. https://unesdoc.unesco.org/ark:/48223/pf0000385146

Slavin, R.E. 1983. Cooperative learning. New York: Longman.

Sparrow, R. 2021. Why machines cannot be moral. *AI & Society: Journal of Knowledge, Culture and Communication* 36 (3): 685-693.

Sparrow, R. 2007. Killer robots. Journal of Applied Philosophy 24 (1): 62-77.

Sparrow, R., and G. Flenady. 2025a. Bullshit universities: The future of automated education. *AI & Society: Journal of Knowledge, Culture and Communication*. Published online: 2 April 2025. https://doi.org/10.1007/s00146-025-02340-8.

Sparrow, R., and G. Flenady. 2025b. The testimony gap: Machines and reasons. *Minds and Machines*. Published online: 24 February 2025. https://doi.org/10.1007/s11023-025-09712-5.

Sparrow, R., Koplin, J., and G. Flenady. Generative AI is dangerous – but not for the reason you might think. *ABC*. <u>https://www.abc.net.au/religion/why-generative-ai-like-chatgpt-is-</u>bullshit/102010238

Tufekci, Z. 2022. What Would Plato Say About ChatGPT? *The New York Times*. https://www.nytimes.com/2022/12/15/opinion/chatgpt-education-ai-technology.html

Umbach, P.D. and M.R. Wawrzynski. 2005. Faculty do matter: The role of college faculty in student learning and engagement. *Research in Higher Education* 46 (2): 153-184.

Véliz, C. 2021. Moral zombies: Why algorithms are not moral agents. AI & Society: Journal of Knowledge, Culture and Communication, 36: 487–497.

Watermeyer, R., Phipps, L., Lanclos, D., and C. Knight. 2024. Generative AI and the Automating of Academia. *Postdigital Science and Education* 6 (2): 446-466.

Zhou, J., Dawson, P., Tai, J. H. M. and M. Bearman. 2021. How conceptualising respect can inform feedback pedagogies. *Assessment & Evaluation in Higher Education* 46 (1): 68-79.