AI & SOCIETY (2025) 40:5285–5296
https://doi.org/10.1007/s00146-025-02340-8

RESEARCH

Bullshit universities: the future of automated education

Robert Sparrow[1,2] · **Gene Flenady**[1]

Received: 13 November 2024 / Accepted: 3 February 2025 / Published online: 2 April 2025
© The Author(s) 2025

Abstract
The advent of ChatGPT, and the subsequent rapid improvement in the performance of what has become known as Generative AI, has led to many pundits declaring that AI will revolutionize education, as well as work, in the future. In this paper, we argue that enthusiasm for the use of AI in tertiary education is misplaced. A proper understanding of the nature of the outputs of AI suggests that it would be profoundly misguided to replace human teachers with AI, while the history of automation in other settings suggests that it is naïve to think that AI can be developed to assist human teachers without replacing them. The dream that AI could teach students effectively neglects the importance of 'learning how' in order to 'learn that', that teachers are also role models, and the social nature of education. To the extent that students need to learn how to use AI, they should do so in specialized study skills units. Rather than creating a market for dodgy educational AI by lowering their ambitions about what they can offer, universities should invest in smaller class sizes and teachers who are passionate about their disciplines. To flourish in the future, just as much as they do today, societies will need people who have learned to think and not—or not just—intelligent machines.

Keywords Higher education · Generative AI · ChatGPT · Artificial intelligence · Pedagogy · Knowledge

1 Introduction

The advent of ChatGPT, and the subsequent rapid improvement in the performance of what has become known as Generative AI (GenAI), has led to many pundits declaring that AI will revolutionize education, as well as work, in the future (Trumbore 2023; Sidorkin 2024; Khan 2024. On work see Mollick 2024a; Taulli 2023). In this paper, we argue that enthusiasm for the use of AI in tertiary education is misplaced.[1] A proper understanding of the nature of the outputs of AI suggests that it would be profoundly misguided to replace human teachers with AI, while the history of automation in other settings suggests that it is naïve to think that AI can be developed to assist human teachers without replacing them. The dream that AI could teach students effectively neglects the importance of 'learning how' in order to 'learn that', that teachers are also role models, and the social nature of education. To the extent that students need to learn how to use AI, they should do so in specialized study skills units. Rather than creating a market for dodgy educational AI by lowering their ambitions about what they can offer, universities should invest in smaller class sizes and teachers who are passionate about their disciplines. To flourish in the future, just as much as they do today, societies will need people who have learned to think and not—or not just—intelligent machines.

2 The brave new world of AI educators

The release of ChatGPT in 2022 led to an eruption of interest in pedagogical applications of Generative AI (Bearman et al. 2023; Farrelly and Baker 2023; Jensen et al. 2024; Korseberg and Elken 2024). As Martin Weller (2014) has argued compellingly, the "Edtech industry" is deeply committed to the idea that education as we know it is fundamentally "broken," and that it is "visionary" technology—not

✉ Robert Sparrow
robert.sparrow@monash.edu

Gene Flenady
gene.flenady1@monash.edu

[1] Monash University, Melbourne, Australia

[2] ARC Centre of Excellence for Automated Decision-Making and Society, Melbourne, Australia

[1] We have confined our criticisms here to the use of AI in university settings because that is where we are located. However, it should go without saying that we expect that our arguments will also have force for the use of AI in primary and secondary education.

deeper institutional change—that will provide the solution. It is, therefore, little wonder that GenAI has been heralded as a time-saving 'teaching assistant' able to instantly create lesson plans (van den Berg and du Plessis 2023) and to automatically generate quizzes on course content (Pesovski et al. 2024). Any number of companies are rolling out 'intelligent tutoring systems' able to provide students with timely and 'personalised' feedback (Kasneci et al. 2023) throughout the learning process and on submitted assessment (Sifaleras and Lin 2024), thus ostensibly building students' felt connection to subject content (Fuchs 2023), and promoting behavioral engagement and motivation (Ng et al. 2024). The ease with which machine learning systems learn in response to the features of particular datasets suggests that it will soon be possible to create individualized AI tutors for each and every student (Khan 2024)—a fantasy that, as Watters (2021) demonstrates, has animated the pursuit of educational technologies since the 1920s. Insofar as one would hope that there is a connection between feedback and grading—both because, in ordinary circumstances, students should not be guided in ways that would not improve their grades and because grades themselves are a form of feedback—provision of feedback by AI is likely to be followed, reasonably quickly, by grading by AI. Indeed, AI grading is already being promoted as overcoming scalability limitations of traditional grading processes in higher education (Fagbohun et al. 2024).[2] Finally, improvement in computer graphics engines means that it will soon be possible to create near photo-realistic 'Digital Humans', which, combined with GenAI, might serve as tutors or even deliver online lectures (Liu 2024; OIAI 2023). It is all too possible to imagine these technologies being combined into a system that purports to provide students with an education without them having any contact with human teachers.

All these innovations are being explored and introduced in a period in which students can use—and are using—AI to complete the work set for them by universities. AI built into *Adobe Acrobat* summarizes articles so our students do not have to read them. *Scite* and *Elicit* will survey the scientific literature and perform bibliographic research so students no longer need to go the library. *Mathway* and *Socratic* will answer math questions. *ChatGPT* now writes as good an undergraduate essay as most students—and a much better essay than the vast majority of students writing in a second language (Herbold et al. 2023). There is almost nothing that students are permitted by universities to do on a computer that cannot be done just as well, if not better—and much, *much*, faster—by AI.

If educators begin using AI to design classes and assess work, they will struggle to explain why their students should not use AI to write it. If we are not careful, then, we will shortly live in a world in which machines are marking the work of other machines.

3 How did we get here?

The traditional way to get machines to succeed at a task that could previously only be performed by human beings is to first make the work of the humans more mechanical (Braverman 1998). That usually means standardizing what was previously bespoke, so that there is less role for judgment and skill in the production of the good, breaking the task up into subtasks, which machines might be able to take on, and lowering expectations. Success often requires transforming the economics of the setting, such that inferior, but much cheaper, products, outsell and ultimately displace what was previously thought to be worth buying—and manufacturing.

Unfortunately, in the last several decades, governments and universities have done all of this, unwittingly or not paving the way for machines to replace teachers (Farnell 2023). Larger class sizes and lower *per* capita funding have increased the amount of marking and made it more difficult for teachers to provide individualized feedback on submitted work (Henderson et al. 2019). More and more classes are taught by staff in ongoing or 'long-term' precarious employment (Bone 2021), who often have little time to prepare for classes, few opportunities to engage in research, and are routinely requested to teach subjects outside of their area of expertise (Lopes and Dewan 2015), with the result that teaching has gradually become more standardized. Assessment has also become increasingly rote. If they are not answering multiple-choice questions, the majority of students submit work that consists of a pastiche of sources, minimally rewritten so as to avoid accusations of plagiarism: by-and-large their essays avoid taking controversial, let alone original, positions, which would require that they argue for them themselves.[3] Staff then use the quality of 'writing' in the result as a proxy for the extent of the student's understanding and their capacity to construct an argument. Students get good grades for stringing an argument together out

[2] This is despite the lack of clarity in current discussions of fairness and transparency in AI grading noted by Memarian and Doleck (2023).

[3] That students take an increasingly instrumental approach to learning and to assessment is not a moral failing on their behalf, but rather an understandable response to profound structural changes in the university, society, and the relationship between the two (Tao 2021). In addition to the slow dismantling of teaching and learning conditions discussed above, the widespread understanding of the purpose of mass higher education as the provision of employable skills tends to promote passive and instrumental approaches to learning (Molesworth et al. 2009).

of the claims of others regardless of its conclusions, as long as the latter are within some acceptable range. This dynamic is especially present in humanities disciplines because for various reasons—some good, some bad—teachers have become reluctant to engage with, or criticize, the substantive intellectual content of their students' essays or exam answers; the emphasis on student evaluations has meant that university teachers are incentivized to 'massage' relations and affirm student identity, rather than uphold rigorous disciplinary standards (Means 2019). Seldom is a student told that what they claim to believe is wrong or untrue—or given a poor grade on that basis. The advent of the internet encouraged this style of writing as it became easier for students to find vaguely related sources and access examples of similar pastiche.

More recently, the possibility of placing sources and lecture notes on e-learning or Learning Management System platforms, and especially the evolution of the technology to stream video, has encouraged administrators and some educators to think of education as fundamentally about the delivery of 'content', which can be provided remotely, or even asynchronously in order to provide flexibility (Watts 2016). When lectures and tutorials are delivered online, students have less opportunity to interact with academic staff: indeed, it is hard to see why lectures *need* to be written and delivered by staff employed by the university—or, in the coming Age of AI, by human beings at all.[4]

4 What computers still cannot do

Any assessment of the potential of AI to contribute to education must begin with an accurate understanding of the nature of the outputs of AI. The most important reason to resist the use of AI in universities is that its outputs are fundamentally bullshit—indeed, strictly speaking, they are meaningless bullshit.

4.1 AI-generated output is bullshit

Harry Frankfurt famously defined bullshit as speech characterized by a '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 his/her audience and does not care whether what he/she says is true or not.

The outputs of Generative AI are—not to put too fine a point on it—bullshit in Frankfurt's sense (Hicks et al. 2024; Sparrow et al. 2023). As we will argue further below, one reason why this is true is that GenAI systems do not care about the truth: insofar as they are unfeeling machines, they do not care about anything at all. However, another reason is that GenAI has no access to the world, and thus to truths about it, but only to things that human beings have said about the world, typically in the form of collections of text written online. Notoriously, not everything on the internet is true. Moreover, except in the rare case in which an AI has sensors, AI is unable to test what it 'knows' against the world. Instead, AIs 'learn' whether they should repeat a particular claim by seeing how human beings respond to it. That is to say, they are motivated by the desire to produce a particular reaction in their audience—to *convince* them of the soundness of their response—rather than by the desire to 'say' something true (Hicks et al. 2024: 38). They are, then, fundamentally bullshitters.

Importantly, that the outputs of AI are bullshit does not necessarily mean that they are untrue nor even that they cannot sometimes serve to justify beliefs. As Frankfurt puts it: 'although [bullshit] is produced without concern for the truth, it need not be false' (2005: 47). Just as a bullshitter may—despite not caring about the truth—sometimes achieve his/her rhetorical goals by uttering only true claims, so too may AI sometimes, perhaps even mostly, produce outputs that reliably correspond with states of affairs in the world.

4.2 AI-generated output is meaningless

One might have thought that the fact that AIs do not care about whether their outputs are true or not would already be sufficient to disqualify them from use in most educational settings. However, the problem with the outputs of AI is deeper still, for reasons that are yet to be widely recognized. To understand why, we need to take a brief detour into social epistemology—the science of how we come to have knowledge.

Descartes famously wondered whether he was justified in believing the evidence of his own senses (Descartes 2017). However, most of what we know comes from other people. We are able to learn about things of which we have no personal experience—distant places, far off times, rare phenomena—because other people tell us about them, we believe them, and we are justified in believing them (Coady 1992). That is to say, much of what we know we learn from the testimony of others.

Unfortunately, AI cannot testify: it cannot tell us anything (Sparrow and Flenady 2025). This is the case because telling someone something is an act, for which a person may be held responsible. If someone provides false testimony,

[4] For a discussion of concerns about how much teacher–student interaction is 'enough' in higher education, see Ellis and Romano (2008).

or fails to exercise due epistemic care in the formation of beliefs that they then communicate to others, then we will typically criticize them for it. It is wrong to lie, to falsely claim knowledge, or to be cavalier about the truth. However, as is almost universally acknowledged, machines are not moral agents (Hakli and Mäkelä 2019; Johnson 2006; Sparrow 2021a; Véliz 2021). Being unable to act, computers cannot testify. They do not care about the truth and cannot be held responsible for the outputs they produce.

That computers cannot testify is obscured by the predominance of empiricism, and a consequent focus on the role played by evidence, in our thinking about the justification of belief. Although most of what we learn from others we learn from their testimony, it is also possible to treat what other people say as evidence, which provides a different route to the justification of belief. Just as I can learn that it is raining outside by the fact that the people entering the building are carrying wet umbrellas, I can learn that it is raining by the fact that people are complaining about the rain to each other even when they are not telling *me* anything. Indeed, even if I know that someone is lying to me, I can form beliefs on the basis of what they say (Moran 2006: 292–293). Here testimony plays no role in the justification of belief.

As we noted above, despite being bullshit, the outputs of well-designed AIs may track, even if they are not motivated by a concern for, the truth. That is to say, they may serve as evidence in favor—or against—beliefs. If an AI has only ever told you that it is raining when it is raining, then the fact that it tells you that it is raining can serve as evidence of rain.

However, importantly, that machines are not moral agents has a further, and profound, implication for the proper understanding of the nature of their outputs and for the extent to which these can serve as justifications for belief. When someone tell us something, we interpret and evaluate their words in the context of the other things they say and do. If someone tells us that they just saw a crow and then mentions its beautiful song and light brown plumage, we may quickly realize that they do not know what a crow is and that they are talking about a nightingale. It is the connections between sentences—and between sentences and behavior—that gives words their meaning.

Philosopher Robert Brandom has developed this idea at length in his 'inferentialist' semantics. According to Brandom—and we believe that he is right in this—the meanings of words are determined by their use and, in particular, by the inferences that they support (Brandom 1994). To apply a concept, then, is to legitimize and endorse a particular set of inferences that others might draw from one's using it, and also to commit oneself to uttering certain other sentences—and to acting in certain ways—in appropriate circumstances.

Brandom provides an illustration that is helpful here. He imagines a young child who walks into a room and announces confidently that the house is on fire. However, the child does not run, shows no fear, and, having made this announcement, begins playing with a teddy bear. In such a case, we would say that the child does not know what she is saying: she does not understand the concept of fire (Brandom 2002: 360). Saying 'the house is on fire'—and meaning it—commits one to various other actions and dispositions, including the disposition, all other things being equal, to flee.

That inferences to action are included in the inferences that give concepts their content has profound implications for the epistemic status of the outputs of machines. As we observed above, machines cannot act, not being moral agents. This implies that (what appear to be) sentences produced by machines do not have conceptual content and are thus, strictly speaking, meaningless. The fact that AIs do not mean what they seem to say is obscured by the fact that their outputs often have the same form as sentences that *would* have content were they to be uttered by a human being. Like the Wizard of OZ, these machines appear clever to the extent that we do the work of granting them meaning by imagining that their outputs were the words of a human being.

As Generative AI becomes more sophisticated, and its output—superficially at least—more impressive, it can be hard to keep in mind that there is nothing behind the curtain: there is no mind, no perspective on the world, and no commitment to the truth. If we are cautious, we can learn something from these outputs, which, as we observed above, might well track the claims that would be made by a human expert. A history of reliability can establish that the outputs of AI can serve as a justification for belief. However, insofar as the outputs of AI need to be endorsed, either explicitly or implicitly, by a human being before they deserve to be taken seriously, there are, as we discuss further below, important limits to the extent to which we can learn from AI (Sparrow and Flenady 2025).

5 Computers in social context

Perhaps the problematic status of the outputs of AI would matter less if we could confine the use of AI in education to roles where a reliable source of strings-of-text-that-would-have-conceptual-content-were-they-uttered-by-a-human-being might make a useful contribution. Enthusiasts for AI in education typically insist that AI should—and will—only be used to enhance or supplement the vital work of human educators rather than to replace teachers or reduce the amount of contact that students have with human staff (Kasneci et al. 2023).[5] The history of automation in other

[5] This is not to deny that the 'dystopian extreme' of the extant literature on AI and higher education includes those who fear the replacement of teachers (Bearman et al. 2023). Selwyn (2019) is an important precursor to our argument here, insofar as it explicitly

contexts suggests that this is naïve. Computers, like other technologies, are shaped by the socio-economic contexts of their use, while at the same time shaping human behavior (Howcroft and Taylor 2023; Sparrow 2021b; Winner 1980), and thus there are important limits on the ways in which one might realistically anticipate that AI will be used.

5.1 Automation bias

One way that computers shape human behavior, which is highly relevant to the prospects of their success in educational contexts, is by encouraging human beings to rely upon them. Notoriously, users of computers often come to suffer from what is known as 'automation bias'—an over-trust in the outputs of a computer (Skitka et al. 1999). In particular, users will often come to trust that a computer that works most of the time will work all of the time. This is especially the case for systems that are generally reliable, even if still fallible. In part, this is because the task of monitoring a computer, or its outputs, is usually more psychologically—and often intellectually—demanding than performing the task oneself: people struggle to pay adequate attention to temporally extended sequences of events if they are not required to have regular and direct input into them (Cummings 2017). However, it is also because human beings suffer from a number of well-known biases when it comes to their ability to think about risk, such that we tend to overestimate the salience of experiences that are immediately available to us, such as when computers work, and underestimate the impact of low probability events, such as when the computer fails when it encounters some unusual set of conditions.

In educational contexts, this means that the idea that AI could perform tasks previously performed by human beings but that human beings could check them is naïve. If the grades or feedback provided by AI are usually accurate, teachers will quickly stop reading the work of their students and/or checking the 'reasoning' of the computer. The task will effectively be handed over to AI completely, even though, for institutional and legal reasons, the human staff member is likely to be held responsible should anything go wrong.

5.2 Displacement of labor

Rather than—or, at least, as well as—reducing the amount of labor involved in the production of some good, historically, computers have tended to change the nature of the labor involved. That is, computers, in substituting for human beings in some part of a production process, create new and different tasks for human beings (Autor 2015; Rosenberg 2013). We can see this in, for instance, the increase in the number of IT staff and 'educational designers' at universities, as well as the amount of time that teaching staff are required to spend creating websites, since the introduction of 'learning management systems' like *Moodle* and *Blackboard*. Tools that were supposed to be labor saving generated a surprising amount of work outside of the classroom. Similarly, as AI enters universities, we can expect that staff will need to spend time learning how to use AI, providing and curating data for AI, and checking the outputs of AI.[6] The time spent in these activities will not be spent in face-to-face interaction with students.

5.3 The economics of automation

Of course, sometimes computers do eliminate certain forms of work, even if their tendency to reduce the overall amount of labor involved in the production—broadly conceived—of goods is often exaggerated (Autor 2015). Machines will tend to displace humans from particular sorts of work as soon as it becomes cheaper to use a machine than to employ a human being (Brynjolfsson and McAfee 2014, 180–181; Frey and Osborne 2017). There is, therefore, a profound tension between the claim that machines can now—or will soon be able to—perform tasks that are currently performed by human teachers and the claim that AI will only be used to supplement—rather than replace—human educators.[7] We should recognize the argument about AI 'helpers' for what it is: a Trojan Horse to bring about full automation of education.

6 Why education resists automation

Unfortunately, contemporary enthusiasm for the use of AI in educational settings is not solely a product of naïvete about the nature of AI and about human–computer interaction: it is also a function of a shallow account of what education is and of how students learn (Biesta 2007). Understanding why education resists automation requires reminding ourselves of some hard truths about what makes for a good education (Biesta 2016a, 2016b).

Footnote 5 (continued)

acknowledges that it is in fact *possible* for computers to replace teachers, but then argues that they ought not on normative grounds.

[6] Which is to say, the introduction of AI is also likely to increase the number of "bullshit jobs", in Graeber (2018)'s sense, in education.

[7] A full account of the economics of automation is beyond the scope of our discussion here. Our point is simply that "money talks", in education as elsewhere, and thus that it is naïve to think that teachers will continue to be paid to perform tasks that machines can do more cheaply.

6.1 Learning how

Educating students requires both inculcating them with a body of truths about the world, including facts about the history of particular disciplines, and teaching them skills essential to the disciplines they are studying and/or the profession to which they aspire. That is, it involves teaching 'knowledge that' and 'knowledge how' (Ryle 1945). Students, in turn, must 'learn that' *and* 'learn how.'

To the extent that one is worried about students learning various facts, the philosophical concerns above may seem irrelevant. Who cares whether machines are committed to the truth or stand behind their words? What matters is whether their outputs are reliable. Students can memorize a list of a facts provided to them by a machine just as well as they can a list enumerated by a human being. Within certain limits, machines can identify when students have made errors in their work and remind them of the correct answers.

This is too swift. In various, albeit subtle, ways 'learning that' is parasitic on 'learning how' (Hetherington 2006)—and there are important limits on the extent to which machines can teach 'knowledge how'. One cannot learn how to ride a bicycle—or how to carry out surgery, or perform a titration, or to speak in public—from an AI. To learn practical, worldly, skills we need to watch closely as other people demonstrate them. More importantly, to develop skills, we need to practice them ourselves and receive feedback from other people who already possess them as we do so. Sometimes verbal advice or criticism will suffice; sometimes we need someone else to reposition our body or demonstrate correct technique alongside us (Kremer 2021). This is not to suggest that computers can make no contribution to learning how: it is, for instance, possible to learn some skills by watching videos, which presumably could be curated, or even generated, by AI.[8] However, developing a skill still requires practicing it and there are real limits to what most people can learn without contact with experts in real life.

To learn the skills necessary to become a surgeon, a wildlife biologist, a chemical or mechanical engineer, then, requires that one practice these skills in the relevant real-world contexts, and also receive guidance and feedback from people who are expert in these skills (Crawford 2009). Nor is it the case that the importance of learning how is confined to physical skills, or education in disciplines that require them. Even philosophy students must learn 'how': they must learn how to read philosophical texts, how to write, how to persuade a reader or an audience, how to construct an argument, and how to respond to objections. Indeed, thinking and learning—and the capacity to eventually take ownership over one's thinking and learning—are themselves both skills requiring concerted 'training' (Ryle 1967). We learn these skills by watching others demonstrate them, practicing them ourselves, and then receiving feedback from others more expert in these skills than ourselves. Again, AI can make some contribution to providing relevant feedback on students' performance at some tasks—for instance, Large Language Models are surprisingly good at offering suggestions as to how one might improve a piece of writing—but cannot show students how they should overcome the intellectual, motivational, and psychological challenges they face as individual human beings. What students *might* learn from a machine is how to be like a machine—but in fact human beings cannot perform tasks in the same way that AI performs tasks.[9]

The case of writing—and the implications of the use of AI to write—is particularly instructive here. Generative AI systems write so well now that it is tempting to think that students no longer need to learn how to write themselves but only how to 'prompt' an AI system to express their thoughts. If one already knows the thought that one wants to express, it is possible to assess whether text generated by AI captures it correctly. However, this hope neglects that the process of writing is not—or not usually—a mechanical one, whereby we put on paper thoughts that we already possess. Rather, writing is often itself the means by which we discover what we think (Menary 2007). This is particularly the case in subjects like philosophy and law, wherein expressing a thought precisely is the *sine qua non* of the discipline. Only by choosing the right words, and constructing the right sentence, can we work out what we mean to say—and thus what we think. The current generation of pundits pontificating about the educational potential of AI grew up having learned to write themselves. They can—sometimes—tell when AI-generated output goes wrong and thus are inclined to believe that it is possible to use AI as a tool to write more quickly and effectively. However, a student who learns only how to prompt an AI, and not how to choose their words themselves, is unlikely to know precisely what they want to say—and thus whether the AI has expressed their own thoughts or put words into their mouths. Students who do not learn how to write, because they let AI write for them, will struggle to learn how to think.

Unless students learn key skills, they will struggle to absorb the body of knowledge associated with the subject that they are studying. Most of the 'content' that students learn, they learn outside of the classroom, from textbooks—or now, depressingly, from watching YouTube videos—and

[8] For a recent, and nuanced, discussion of the impact of an AI system on the extent to which students learned 'how' to provide feedback on each other's work, see Darvishi et al. (2024).

[9] Not least because, notoriously, we have surprisingly little understanding of how Generative AI systems work (Burrell 2016).

they need to learn how to interpret and assess this material in order to do so. Educational use of AI is highly likely to reduce the opportunities to learn how and thus, eventually, to learn that.

A popular vision of the future of the education has students using GenAI to work more quickly and effectively (Khan 2024).[10] To do so, though, given the tendency of AI systems to 'hallucinate' (Ji et al. 2023), students must be able to check the quality of the outputs of the AI in line with disciplinary standards (Bearman and Ajjawi 2023). As we have already observed, in the context of our discussion of automated grading, the phenomenon of automation bias calls into question the idea that such checking is likely to be effective. Our discussion here suggests a more profound problem with this approach: the more students are encouraged the use AI in the course of their education, the less likely they are to develop the skills and knowledge necessary to do so effectively.

6.2 Education and example

As we noted above, a previous generation of technologies—the internet and streaming video—encouraged managers of (and in) universities to conceptualize teaching as the delivery of 'content', which in turn encouraged the idea that learning could occur remotely and asynchronously. Unfortunately, this misses most of what goes on in the classroom (Biesta 2016a). We have already noted the role played by teaching how in teaching skills. However, equally important is the example provided by teachers when it comes to demonstrating why someone might care about the material being taught.

A good teacher teaches a way of being-in-the-world. They show students what it is like to be animated by, and called to obedience to, the regulative ideals of an intellectual, the academy, and of a particular discipline (Gaita 2004: 283–330). Importantly, commitment to these ideals is a moral commitment—a commitment to honesty, to a certain form of selflessness, and to the norms of a community of seekers and knowers.[11] This is one of the reasons why contact with a good teacher can be so transformative. Students discover that someone they admire cares about the material they are studying and what it means in practice to care about it. They learn what it means to be a philosopher, a historian, a mathematician, a doctor, or an engineer—and are given the opportunity to imagine that they might become such a person. Bad teachers also teach their students something, albeit inadvertently: what it is to fall short of the standards internal to their discipline. Even if an AI were to perform perfectly when it came to the production of educational 'content', it would not serve as an example to students. It could not teach students why they should care about, or how they should relate to, the material that they are studying.[12] The problems that AI needs to overcome to perform well in science, mathematics, history, or philosophy, etc., are not the challenges that students need to confront and overcome in order to learn and flourish as scientists, mathematicians, historians, or philosophers.

6.3 Education and others

Whenever, and wherever, AI is used in an educational setting, a machine will be doing something that might otherwise have been done by a human being. As we suggested above, whether advocates intend this or not, the economic imperatives informing enthusiasm for the educational use of AI tend towards a world in which students only interact with AI. This grim vision of the future of education neglects that education is social in ways that extend from the mundane to the profound.

One important thing that students gain from attending a university is an opportunity to meet new people, to encounter others who might challenge their own beliefs, and to make new friends. The relationships students establish on campus are a form of social capital. Classmates today are philosophical interlocutors, artistic collaborators, or business partners tomorrow.

Students also learn social skills in the classroom. They learn how to speak in public, to frame an argument so it will convince an audience, and to offer, and to receive, criticism of the arguments and work of others. These are skills that they will need in the workforce as well as in the pursuit of knowledge. It is not possible to learn these skills except in relationships with other people.

In the course of practicing and learning these skills, students also enter into the social world in a more profound way. They experience what Hegel termed 'recognition' [*Anerkennung*] (Hegel 2019). That is, they discover their own moral and political status, and their authority in relation to knowledge claims, through their relationships with others, who recognize them as having such. This recognition has a subjective and an objective component. Subjectively, students grow more confident, which allows them to assert themselves when they leave university and enter the workforce. Objectively, students are acknowledged as a source

[10] For the claim that this is future of work more generally, see Mollick (2024a).

[11] Not every teacher—not even every great intellectual—demonstrates these virtues but the pursuit of knowledge would proceed more haltingly, if at all, if they were not ideals to which members of the relevant communities aspire and mostly live up to.

[12] For a survey of the findings of the existing empirical literature on the impact of AI on student engagement, see Lo et al. (2024).

of claims to whom respect is owed: we express that respect when we provide reasons for the way we treat them. We enter the kingdom of ends and become citizens in a democratic polity by means of the recognition others grant us as ends in ourselves (Korsgaard 1996).

Finally, education is social because language, concepts, and knowledge are social. This in turn means that our mental life is fundamentally social. It is through conversation and exchange with other moral agents with ends and perspectives different to our own that we learn what the language in which we express our thoughts and desires means—and thus what those thoughts and desires are.

Taking teachers out of the classroom, or—worse still—taking students out of the classroom and putting them in front of computers, would be disastrous for education and for society and culture more generally. It will deprive students of social capital, as well as of many of the experiences that once made attending university so rewarding. In the absence of interactions with staff and other students, students will not learn social skills that are essential to intellectual inquiry and to participation in shared projects. A world in which 'content' is generated, and essays and exams are marked, by AI would be a world in which no-one takes students, and what they think, seriously: if nobody takes my work seriously, why should I do so? A world in which what is taught and learnt at universities is chosen or generated by AI would also be a world in which the expression and progress of human thought is outsourced to mindless machines.

7 Responses to objections

Inevitably, there are any number of objections that might be made to the argument we have sketched out above: if the reasons to resist the use of AI in education were obvious, the question of whether we should do so would not arise. In this section, we respond to four objections that, we acknowledge, are worth taking seriously.

7.1 Book learning

There is an obvious riposte to the line of argument we have been developing here, which points to the role played by another educational technology that was itself controversial in its time—the book! Our emphasis on the importance of learning how, on the role played by example of commitment to the truth and to the critical standards of the discipline, and on the fundamentally social process of education might be held to imply that it is impossible to gain an education by reading books.

The idea that we should restrict the use of books at universities is difficult to countenance. Nevertheless, it is worth pointing out that there *are* many things that one cannot learn from books. In particular, as we observed above, one cannot learn 'knowledge how' from books alone. Nor can one learn what it is, in practice, to be someone who is committed to the pursuit of knowledge and to the critical standards internal to a particular discipline. To the extent that one can learn some of these things from books, it is because, in a well-written book, the *persona* of the author shines through.

This highlights a fundamental difference between books and AI: books have authors in a way that the outputs of AI do not. The authority of the text is derived from the expertise and moral standing of the author. We are justified in forming beliefs on the basis of what we read in books because the author stands behind the work: if the claims in a book are false, or morally pernicious, we rightly blame the person who wrote it (and sometimes the person who published it as well).

Trying to learn from an AI is like reading a book without an author. At most, chatbots and other AI systems have designers. Designers do have some responsibility for the outputs of AI. Elon Musk's 'Borg' is noticeably different in tone and style than Microsoft's Copilot or Google's Gemini. Moreover, if the outputs of an AI play a role in bringing about bad consequences, the designer is one of the people who might plausibly be held responsible for these consequences. Nevertheless, the designers of AI do not—and cannot—stand behind the words generated by AI in the same way that an author stands behind the words in a book. An author, after all, chooses his/her words. In an important sense, their words are their 'own'. A designer of a Generative AI creates a machine that produces outputs that he or she may never see and perhaps could never have imagined.[13]

Students can learn from textbooks because books contain and express the thoughts of experts in the field. The authors of these books stand behind their words and the claims therein count as a form of testimony. As we have seen, the outputs of AI do not express anything at all. At most, they parrot the words of others: even when they do so accurately, their claims do not count as testimony and do not place us under the same obligation to take them seriously.

7.2 Access and equity

Much of the contemporary enthusiasm for the use of AI in education arguably derives from a techno-fetishism present in the broader culture, accompanied by the desire of universities and educators to be seen as 'progressive' and 'on-side' with progress in science and engineering. More cynically, some administrators clearly see AI as a mechanism

[13] A previous generation of educational software that functioned like a 'choose your own adventure book' did allow the programmer/designer to stand in the role of an author.

for reducing staff numbers and cutting costs. However, there are at least two morally admirable reasons for enthusiasm for AI. Many people, especially in the global South but also in First Nations and other marginalized communities in the global North, are currently unable to access higher education, either because of geography, or as a result of poverty, or both. Online classes enhanced with AI have the potential to radically expand access to education (Kasneci et al. 2023; Khan 2024).[14] Systems that use machine learning are also well suited to identifying strengths and weaknesses in the performance of an individual student, or a group of students, and structuring an educational program to suit their learning needs (Escotet 2023; Fuchs 2023).

These are genuine virtues of AI. However, assessment of their implications must begin with acknowledging the force of the critique we developed above. AI is likely to greatly expand access to, and/or individually tailor, an education that is significantly worse than that which traditional educational institutions aspire to provide. Indeed, any education provided by AI would lack key features of a traditional education, requiring students to teach themselves essential skills as well as to develop a love of their discipline in the absence of role models.[15] Whether it is appropriate to offer those who are already socially and economically disadvantaged an inferior education because the alternative would be none at all is an open—and difficult—ethical question. That embracing AI will degrade the quality of education offered to those currently privileged enough to be able to access it also needs to be factored into this deliberation.

It is also worth noting that there are alternative ways to expand access to education that do not involve AI—for instance, making textbooks open access, investing in education locally and/or correspondence courses—and that may in fact be more accessible to potential students in the global South, as well as to marginalized communities, including First Nations, in the global North, by virtue of not requiring students to have access to a computer, a reliable source of electricity, and the Internet. Given the historical failures of universities in the Global North to make education cheaper and more widely available via such low-tech methods, one might reasonably conclude that contemporary enthusiasm for the idea that AI will improve access to education represents the attraction of a technological fantasy where a solution to a political problem is required.

7.3 The future of work

Another objection to resisting the introduction of AI into educational settings points to the likelihood that students will need to be able to use AI effectively when they enter the workforce. If the sorts of jobs that graduates are likely to get will involve using AI, then it is important that they learn how to use it and also how to think critically about what AI can and cannot do (Al Naqbi et al. 2024; Mollick 2024b; Rudolph et al. 2023).

If one really believed that AI would revolutionize the nature of work, then you might have thought that it followed that universities should hold off changing what—and, more importantly, how—they teach until the nature of this revolution became clear. We suspect that, whatever the future of work, the core skills taught by the traditional university will remain relevant and valuable (Royer 2024)—and perhaps especially, if we may be forgiven for spruiking our own discipline, those taught in philosophy.

That said, we agree that it is important that students learn how to use AI and use it well. However, that AI is a tool that students will need to learn how to use, as well as a powerful technology that it is appropriate to study in, for instance, classes in philosophy of technology or the history and philosophy of science, does not establish that students should be taught by AI. Moreover, teaching students how to use AI is best thought of as akin to teaching them how to use a library catalog or to construct a bibliography—other important skills that are most appropriately taught in specialized study skills units rather than in the course of education in particular disciplines. In classes in philosophy, literature, sociology, etc., student should be studying the work of the key thinkers in each discipline, rather than wasting their time on the meaningless outputs of machines.

7.4 Actually existing universities

There is a final objection to our criticisms of the use of AI in education that accuses us, with some justification, of utopianism about what universities offered before the advent of AI. It has been many years since most Australian universities offered students an education that proceeded via talking at length in small groups after exposure to experts who were deeply committed to the truth and whose intellectual lives were governed by the standards internal to

[14] Although, as we acknowledge below, many in the global South, and some—for instance, members of First Nations and other marginalized communities—in the global North, still struggle to access online resources owing to a lack of the physical and digital infrastructure necessary to support this.

[15] There are also well-known issues with cultural, gender, and racial bias in AI (Benjamin 2019; Campolo et al. 2017; O'Neil 2016), which would loom all the larger were education to be made available to members of marginalized communities mostly in the form of AI tutors.

their discipline—if they ever did.[16] Moreover, even when it comes to the most passionate and engaged students, there are some things that they need to learn by rote or skills at which they need only an uninspired competence, such that it seems plausible that AI might teach these things. A more realistic account of the education being offered by actually existing institutions makes it more plausible to believe that AI could improve upon it.

It is worth observing that this is a prime example of the phenomenon to which we drew attention above, whereby people make it more possible to imagine automation succeeding in some task by lowering their standards. Nevertheless, there is truth in the charge that we want to hold AI to a standard that much contemporary education often fails to meet—and to reject AI on that basis. Where others are optimistic about AI because they are pessimistic about the future of education, we choose to be optimistic about what societies could achieve if they invested in education itself rather than AI.

8 Back to the future

The best thing that can be said about enthusiasm for the use of AI in education is that it provides an opportunity to think seriously about the significance of the outputs of AI and about how—and what—students learn at university. Unfortunately, this is an opportunity that seems to be mostly being missed in the current moment.

Despite the impressive advances in Generative AI over recent years, the outputs of these systems are neither oriented towards the truth nor connected to action in such a way as to have conceptual content. Employing AIs as teachers would undermine the authority of all teachers and of the body of established knowledge that they aim to impart to students. Moreover, students must learn how in order to learn that, be exposed to the personal examples of their teachers in order to understand what it means to be committed to a discipline, and interact with other students and staff in order to realize themselves as moral agents and citizens. Machines offer little in service of these goals. To the extent that students will need to learn how to use AI, they should do so in units dedicated to this purpose rather than across the whole of their degrees.[17]

Rather than investing in AI, good universities will invest in (human) teachers and students. They will reduce class sizes and ensure that they are taught by people who are passionate about their field of expertise. Given the ease with students can now complete any task that they can submit online by using AI, those who want to assess what their students know, and can do, will need to spend more time talking with students and assessing them on that basis and also require more work to be completed under invigilated conditions. They will prize and reward students who can think for themselves rather than those who rely on machines to pretend to think for them.

Developments in AI pose profound philosophical, ethical, social, and political challenges to the way we live, work, and learn today. Confronting these will require lengthy and difficult democratic deliberation among an educated citizenry (Formosa et al. 2024; Sparrow 2020). It would be a perverse outcome, and a terrible shame, if, in their race to embrace AI, universities made this less, rather than more, likely.

Acknowledgements The authors would like to thank Mark Andrejevic, Thomas Corbin, Lloyd Humberstone, Zala Volcic, and an anonymous referee for comments and discussion that have improved this paper. Professor Sparrow is an Associate Investigator in the Australian Research Council Centre of Excellence for Automated Decision-making and Society (Grant number CE200100005) and worked on this paper in this role.

Author contributions Both authors contributed to the conceptualization and design of the research. RS wrote the first draft of the paper. GF provided comments on and revised the initial draft and contributed to the second draft of the manuscript. Both authors then further revised and contributed to the manuscript. Both authors read and approved the final manuscript.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions. This research was not supported by external research funds.

Data availability No datasets were generated or analyzed during the current study.

Declarations

Conflict of interest One of the authors, Sparrow, is an Associate Editor of AI & Society.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

[16] We have confined this claim to the Australian context with which the authors are most familiar but we suspect that it is also true of many other contexts.

[17] Mollick (2024b), which makes suggestions as to when to—and when not to—use AI strikes us as a useful starting point in this context.

References

Al Naqbi H, Bahroun Z, Ahmed V (2024) Enhancing work productivity through generative artificial intelligence: a comprehensive literature review. Sustainability. https://doi.org/10.3390/su16031166

Autor DH (2015) Why are there still so many jobs? The history and future of workplace automation. J Econ Perspect 29(3):3–30. https://doi.org/10.1257/jep.29.3.3

Bearman M, Ajjawi R (2023) Learning to work with the black box: pedagogy for a world with artificial intelligence. Br J Educ Technol 54(5):1160–1173. https://doi.org/10.1111/bjet.13337

Bearman M, Ryan J, Ajjawi R (2023) Discourses of artificial intelligence in higher education: a critical literature review. High Educ 86(2):369–385. https://doi.org/10.1007/s10734-022-00937-2

Benjamin R (2019) Race after technology: abolitionist tools for the New Jim Code. Polity, Cambridge

Biesta G (2007) Why "what works" won't work: evidence-based practice and the democratic deficit in educational research. Educ Theory 57(1):1–22

Biesta GJJ (2016a) The beautiful risk of education. Routledge, London and New York

Biesta GJJ (2016b) Good education in an age of measurement: ethics, politics, democracy. Routledge, Abingdon and New York

Bone KD (2021) Cruel optimism and precarious employment: the crisis ordinariness of academic work. J Bus Ethics 174:275–290. https://doi.org/10.1007/s10551-020-04605-2

Brandom RB (1994) Making it explicit: reasoning, representing, and discursive commitment. Harvard University Press, Cambridge

Brandom RB (2002) Tales of the mighty dead: historical essays in the metaphysics of intentionality. Harvard University Press, Cambridge

Braverman H (1998) Labor and monopoly capital: the degradation of work in the twentieth century. Monthly Review Press, New York

Brynjolfsson E, McAfee A (2014) The second machine age: work, progress, and prosperity in a time of brilliant technologies. WW Norton & Company, New York and London

Burrell J (2016) How the machine 'thinks': understanding opacity in machine learning algorithms. Big Data Soc. https://doi.org/10.1177/2053951715622512

Campolo A, Sanfilippo MR, Whittaker M, Crawford K (2017) AI NOW 2017 report. AI Now Institute, New York University, New York. https://ainowinstitute.org/publication/ai-now-2017-report-2. Accessed 31 Jul 2024

Coady CAJ (1992) Testimony: a philosophical study. Clarendon Press, Oxford

Crawford MB (2009) Shop class as soulcraft: an inquiry into the value of work. Penguin Books, New York

Cummings ML (2017) Automation bias in intelligent time critical decision support systems. In: Harris D, Li W-C (eds) Decision making in aviation. Routledge, London and New York, pp 289–294

Darvishi A, Khosravi H, Sadiq S, Gašević D, Siemens G (2024) Impact of AI assistance on student agency. Comput Educ. https://doi.org/10.1016/j.compedu.2023.104967

Descartes R (2017) Meditations on first philosophy, with selections from the objections and replies (trans and ed Cottingham J). Cambridge University Press, Cambridge

Ellis J, Romano D (2008) Synchronous and asynchronous online delivery: how much interaction in e-learning is enough in higher education? In: Bonk CJ, Lee MM, Reynolds T (eds) E-Learn: world conference on E-learning in corporate, government, healthcare, and higher education. Association for the Advancement of Computing in Education, Las Vegas, pp 2615–2620

Escotet MÁ (2023) The optimistic future of artificial-intelligence in higher-education. Prospects. https://doi.org/10.1007/s11125-023-09642-z

Fagbohun O, Iduwe NP, Abdullahi M, Ifaturoti A, Nwanna OM (2024) Beyond traditional assessment: exploring the impact of large language models on grading practices. JAIMLD 2(1):1–8. https://doi.org/10.51219/JAIMLD/oluwole-fagbohun/19

Farnell A (2023) AI will replace academics unless our teaching challenges students again. Times higher education, January 19, 2023. https://www.timeshighereducation.com/opinion/ai-will-replace-academics-unless-our-teaching-challenges-students-again. Accessed 31 Jul 2024

Farrelly T, Baker N (2023) Generative artificial intelligence: implications and considerations for higher education practice. Educ Sci. https://doi.org/10.3390/educsci13111109

Formosa P, Kashyap B, Sahebi S (2024) Generative AI and the future of democratic citizenship. Digit Gov Res Pract. https://doi.org/10.1145/3674844

Frankfurt H (2002) Reply to Cohen. In: Buss S, Overton L (eds) The contours of agency: essays on themes from Harry Frankfurt. MIT Press, Cambridge, pp 340–344

Frankfurt H (2005) On bullshit. Princeton University Press, Princeton

Frey CB, Osborne MA (2017) The future of employment: how susceptible are jobs to computerisation? Technol Forecast Soc Chang 114:254–280

Fuchs K (2023) Exploring the opportunities and challenges of NLP models in higher education: is Chat GPT a blessing or a curse? Front Educ. https://doi.org/10.3389/feduc.2023.1166682

Gaita R (2004) Good and evil: an absolute conception, 2nd edn. Routledge, London

Graeber D (2018) Bullshit jobs. Simon & Schuster, New York

Hakli R, Mäkelä P (2019) Moral responsibility of robots and hybrid agents. Monist 102(2):259–275. https://doi.org/10.1093/monist/onz009

Hegel GWF (2019) Phenomenology of spirit (trans Pinkard T). Cambridge University Press, Cambridge

Henderson M, Ryan T, Phillips M (2019) The challenges of feedback in higher education. Assess Eval High Edu 44(8):1237–1252. https://doi.org/10.1080/02602938.2019.1599815

Herbold S, Hautli-Janisz A, Heuer U, Kikteva Z, Trautsch A (2023) A large-scale comparison of human-written versus ChatGPT-generated essays. Sci Rep. https://doi.org/10.1038/s41598-023-45644-9

Hetherington S (2006) How to know (that knowledge-that is knowledge-how). In: Hetherington S (ed) Epistemology futures. Clarendon Press, Oxford, pp 71–94

Hicks MT, Humphries J, Slater J (2024) ChatGPT is bullshit. Ethics Inf Technol. https://doi.org/10.1007/s10676-024-09775-5

Howcroft D, Taylor P (2023) Automation and the future of work: a social shaping of technology approach. New Technol Work Employ 38(2):351–370. https://doi.org/10.1111/ntwe.12240

Jensen LX, Buhl A, Sharma A, Bearman M (2024) Generative AI and higher education: a review of claims from the first months of ChatGPT. High Educ. https://doi.org/10.1007/s10734-024-01265-3

Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, Ishii E, Bang YJ, Madotto A, Fung P (2023) Survey of hallucination in natural language generation. ACM Comput Surv 55(12):1–38. https://doi.org/10.1145/3571730

Johnson DG (2006) Computer systems: moral entities but not moral agents. Ethics Inf Technol 8(4):195–204. https://doi.org/10.1007/s10676-006-9111-5

Kasneci E, Sessler K, Küchemann S et al (2023) ChatGPT for good? on opportunities and challenges of large language models for education. Learn Individ Differ. https://doi.org/10.1016/j.lindif.2023.102274

Khan S (2024) Brave new word: how AI will revolutionize education (and why that's a good thing). Viking, New York

Korseberg L, Elken M (2024) Waiting for the revolution: how higher education institutions initially responded to ChatGPT. High Educ. https://doi.org/10.1007/s10734-024-01256-4

Korsgaard C (1996) Creating the kingdom of ends. Cambridge University Press, Cambridge
Kremer M (2021) Gilbert Ryle on skill as knowledge-how. In: Fridland E, Pavese C (eds) The Routledge handbook of philosophy of skill and expertise. Routledge, Abingdon, Oxon, pp 100–112
Liu D (2024) How Sydney educators are building 'AI doubles' of themselves to help their students. Teaching@Sydney. https://educational-innovation.sydney.edu.au/teaching@sydney/how-sydney-educators-are-building-ai-doubles-of-themselves-to-help-their-students/. Accessed 31 Jul 2024
Lo CK, Hew KF, Jong MS (2024) The influence of ChatGPT on student engagement: a systematic review and future research agenda. Comput Educ. https://doi.org/10.1016/j.compedu.2024.105100
Lopes A, Dewan IA (2015) Precarious pedagogies? the impact of casual and zero-hour contracts in higher education. J Fem Scholarsh 7(Fall):28–42
Means AJ (2019) Precarity and the precaritization of teaching. In: Peters MA (ed) Encyclopedia of teacher education. Springer, Singapore. https://doi.org/10.1007/978-981-13-1179-6_117-1
Memarian B, Doleck T (2023) Fairness, accountability, transparency, and ethics (FATE) in artificial intelligence (AI), and higher education: a systematic review. Comput Edu Artif Intell. https://doi.org/10.1016/j.caeai.2023.100152
Menary R (2007) Writing as thinking. Lang Sci 29(5):621–632. https://doi.org/10.1016/j.langsci.2007.01.005
Molesworth M, Nixon E, Scullion R (2009) Having, being and higher education: the marketisation of the university and the transformation of the student into consumer. Teach High Educ 14(3):277–287. https://doi.org/10.1080/13562510902898841
Mollick E (2024a) Co-intelligence: living and working with AI. WH Allen, London
Mollick E (2024b) 15 times to use AI, and 5 not to: notes on the practical wisdom of AI use. One useful thing, 9 December 2024. https://www.oneusefulthing.org/p/15-times-to-use-ai-and-5-not-to. Accessed 2 Feb 2025
Moran R (2006) Getting told and being believed. In: Lackey J, Sosa E (eds) The epistemology of testimony. Oxford University Press, Clarendon, pp 272–306
Ng D, Tan CW, Leung JKL (2024) Empowering self-regulated learning and science education through ChatGPT: a pioneering pilot study. Brit J Educ Technol 55:1328–1353. https://doi.org/10.1111/bjet.13454
OIAI (2023) Meet the world's first AI digital human teacher. Otermans Institute. https://oiai.oiedu.co.uk/educators. Accessed 14 Jul 2024
O'Neil C (2016) Weapons of math destruction: how Big Data increases inequality and threatens democracy. Allen Lane, London
Pesovski I, Santos R, Henriques R, Trajkovik V (2024) Generative AI for customizable learning experiences. Sustainability. https://doi.org/10.3390/su16073034
Rosenberg RS (2013) The social impact of computers. Elsevier, Saint Louis
Royer C (2024) Outsourcing humanity? ChatGPT, critical thinking, and the crisis in higher education. Stud Philos Educ. https://doi.org/10.1007/s11217-024-09946-3
Rudolph J, Tan S, Tan S (2023) ChatGPT: bullshit spewer or the end of traditional assessments in higher education? J Appl Learn Teach 6(1):342–363. https://doi.org/10.37074/jalt.2023.6.1.9
Ryle G (1945) Knowing how and knowing that: the presidential address. P Aristotelian Soc 46:1–16
Ryle G (1967) Teaching and training. In: Peters RS (ed) The concept of education (International library of the philosophy of education volume 17), 1st edn. Routledge, London, pp 73–82
Selwyn N (2019) Should robots replace teachers? AI and the future of education. Polity Press, Cambridge
Sidorkin AM (2024) Embracing chatbots in higher education: the use of artificial intelligence in teaching, administration, and scholarship, 1st edn. Routledge, New York
Sifaleras A, Lin F (eds) (2024) Generative intelligence and intelligent tutoring systems: 20th international conference, ITS 2024, proceedings part I. Springer, Berlin
Skitka LJ, Mosier KL, Burdick M (1999) Does automation bias decision-making? Int J Hum Comput Stud 51(5):991–1006. https://doi.org/10.1006/ijhc.1999.0252
Sparrow R (2020) Artificial intelligence. In: LaFollette H (ed) International encyclopedia of ethics. John Wiley & Sons, Hoboken. https://doi.org/10.1002/9781444367072.wbiee870
Sparrow R (2021a) Why machines cannot be moral. AI Soc J Knowl Cult Commun 36(3):685–693
Sparrow R (2021b) How robots have politics. In: Véliz C (ed) Oxford handbook of digital ethics. Oxford University Press, Oxford, pp 292–311. https://doi.org/10.1093/oxfordhb/9780198857815.013.16
Sparrow R, Flenady G (2025) The testimony gap: machines and reasons. Minds and Machines. https://doi.org/10.1007/s11023-025-09712-5
Sparrow R, Koplin J, Flenady G (2023) Generative AI is dangerous—but not for the reasons you might think. ABC religion and ethics, 22 Feb 2023, updated 24 Apr 2024. https://www.abc.net.au/religion/why-generative-ai-like-chatgpt-is-bullshit/102010238. Accessed 31 Jul 2024
Tao Y (2021) The understanding and take on the blatant instrumentalism among university students: an early career academic perspective. Aust Univ Rev 63(2):62–66
Taulli T (2023) Generative AI: how ChatGPT and other AI tools will revolutionize business. Apress, Berkeley
Trumbore A (2023) ChatGPT could be an effective and affordable tutor. The conversation, February 22, 2023. http://theconversation.com/chatpt-could-be-an-effective-and-affordable-tutor-198062. Accessed 28 Feb 2023
van den Berg G, du Plessis E (2023) ChatGPT and generative AI: possibilities for its contribution to lesson planning, critical thinking and openness in teacher education. Educ Sci. https://doi.org/10.3390/educsci13100998
Véliz C (2021) Moral zombies: why algorithms are not moral agents. AI Soc J Knowl Cult Commun 36:487–497. https://doi.org/10.1007/s00146-021-01189-xch
Watters A (2021) Teaching machines: the history of personalized learning. MIT Press, Cambridge
Watts L (2016) Synchronous and asynchronous communication in distance learning: a review of the literature. Q Rev Distance Educ 17(1):23–32
Weller M (2014) The battle for open: how openness won and why it doesn't feel like a victory. Ubiquity, London
Winner L (1980) Do artifacts have politics? Daedalus 109(1):121–136

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.