Educational Responses to Artificial Intelligence (AI) Applications: Problems and Promise

Terry Hyland[1]

1 Free University of Ireland

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.

Abstract

The recent appearance of Artificial Intelligence (AI) applications such as Google's Bard tool and OpenAI's ChatGPT series has brought forth a flurry of reactions from both academic and popular commentators ranging from doom-laden pessimism to futuristic hyperbolic optimism. Both ends of this spectrum are revealed in the educational literature with the sharp contrast between critical accounts predicting the decline of academic standards and those welcoming the new applications as a means of enhancing teaching and learning. In the attempt to steer a middle way between these extremes, this article examines the merits and demerits of AI applications in education alongside the suggestion of possibilities for working with the new tools in a productive manner. In relation to ethical considerations raised by the recent developments, it is suggested that AI applications are best conceptualised as powerful entities – perhaps ideally regarded as non-human persons – which need to be utilised pragmatically and regulated ethically in partnership with humans in the best interests of all of us. Such a partnership model will be especially important to learners and teachers who are charged with making sense of a rapidly changing educational environment.

Professor Terry Hyland

Free University of Ireland, Dublin, Ireland –hylandterry@ymail.com

ORCID ID: https://orcid.org/0000-0001-8539-8211

Keywords: Artificial Intelligence (AI), ChatGPT, Consciousness, Educational technology, Learning / teaching, Professional ethics, Non-human persons.

1. Introduction

The appearance in recent times of OpenAI's ChatGPT series and Google's Bard application has resulted in a spate of articles analysing such developments which reflect views ranging from catastrophic conspiracy theories to incredulous debunking. As Renaud Foucart (2023) has commented:

> *AI is expected to affect every aspect of our lives – from healthcare to education, to what we look at and listen to, and even how well we write. But AI also generates a lot of fear, often revolving around a god-like computer becoming smarter than us, or the risk that a machine tasked with an innocuous task may inadvertently destroy humanity. More pragmatically, people often wonder if AI will make them redundant (p.1).*

Many of the concerns about the new AI tools have been expressed by educators who fear that teaching and learning will be damaged by the easy access to the large language models (LLMs) like ChatGPT which can write essays and answer

assignment questions in a matter of minutes (Heaven, 2023). A letter to *The Times* signed by leading educators from the state and private sector described AI developments as 'bewildering' and announced the launch of a review body of experts to advise schools on which areas are 'beneficial, and which are damaging' since 'we have no confidence that the large digital companies will be capable of regulating themselves in the interests of students, staff and schools' (Shad, 2023, p.1). In a similar vein, the *MIT Technology Review* recently reported that as soon as ChatGPT became publicly accessible in November 2022:

> *Los Angeles Unified, the second-largest school district in the US, immediately blocked access to OpenAI's website from its schools' network. Others soon joined. By January, school districts across the English-speaking world had started banning the software, from Washington, New York, Alabama, and Virginia in the United States to Queensland and New South Wales in Australia. Several leading universities in the UK, including Imperial College London and the University of Cambridge, issued statements that warned students against using ChatGPT to cheat (Heaven, 2023, pp.1-2).*

Against this, some educators and policymakers, whilst urging the need for control and regulation of the new applications, have sought to incorporate them in their programmes as a means of gaining insights and enhancing teaching and learning at all levels. Christine Elgersma (2023), for example, outlines a wide range of creative educational uses for AI in schools, and Karen Lancaster (2023) recommends strategies by which universities can use AI essay-writing software in imaginative ways to improve research and presentation skills.

Key themes in this debate about AI in education will be examined in more depth in later sections. At this stage, it would be useful to look at some general issues surrounding the rapid growth of technology in the machine learning field in recent times.

2. Emerging concerns about AI in general

Alongside the warnings of educators about the dangers to learning and teaching posed by the rapid development of AI tools, there has been a response from experts and representatives of the very technological field that has spawned the new generative machine learning devices. A key example is the recent open letter signed by, amongst other leading figures in the digital technology world, Elon Musk and co-founder of Apple, Steve Wozniak, which called for a pause to current machine learning AI developments whilst the wider implications are evaluated carefully. The letter published by the *Future of Life Institute* (2023) warns that:

> *AI systems with human-competitive intelligence can pose profound risks to society and humanity, as shown by extensive research and acknowledged by top AI labs****... Powerful AI systems should be developed only once we are confident that their effects will be positive and their risks will be manageable...****Therefore,* ***we call on all AI labs to immediately pause for at least 6 months the training of AI systems more powerful than GPT-4.***

> *This pause should be public and verifiable, and include all key actors. If such a pause cannot be enacted quickly, governments should step in and institute a moratorium.* ***(p.1, original italics)***

In a similar vein – in an announcement that stunned the AI tech world – Geoffrey Hinton, the so-called "godfather of AI" quit his job at Google warning of the dangers of unregulated and unsupervised trends in the field. Among his comments, he observed that:

> *Some of the dangers of AI chatbots were "quite scary", warning they could become more intelligent than humans and could be exploited by "bad actors". "It's able to produce lots of text automatically so you can get lots of very effective spambots. It will allow authoritarian leaders to manipulate their electorates, things like that. But, he added, he was also concerned about the existential risk of what happens when these things get more intelligent than us. I've come to the conclusion that the kind of intelligence we're developing is very different from the intelligence we have. So it's as if you had 10,000 people and whenever one person learned something, everybody automatically knew it. And that's how these chatbots can know so much more than any one person. (Taylor & Hern, 2023, pp.1-2)*

All of these matters – especially those being expressed by experienced experts and practitioners in the AI field – merit serious consideration. Before examining issues linked to intelligence and consciousness relevant to AI in education, it would be useful to highlight some general concerns about the use of AI in the workplace and the key issues surrounding the control and regulation of tools and applications.

2.1. AI and Work

The chief concern in this area is that AI tools and applications will replace many jobs currently performed by humans thus leading to mass redundancy and unemployment. British Telecom announced recently that it would be cutting its workforce by 55,000 with 11,000 of these jobs replaced by AI (BBC News, 2023). On a more dramatic scale, a recent report by Goldman Sachs predicted that around 300 million jobs would in the future be lost or degraded by forms of AI (Kelly, 2023).

There are, however, alternative views of these developments that offer a less catastrophic and more hopeful vision of the future. A report by the *World Economic Forum* (2023), for example, makes the following important points about what commentators are calling the 'fourth industrial revolution:

- Around 40% of all working hours could be impacted by AI large language models (LLMs) such as ChatGPT-4, says a report from Accenture.
- Many clerical or secretarial roles are seen as likely to decline quickly because of AI, according to the World Economic Forum's *Future of Jobs Report 2023.*
- But roles for AI and machine learning specialists, data analysts and scientists, and digital transformation specialists are expected to grow rapidly, the report adds.
- Reskilling people to use AI effectively will be the key to companies being able to use the technology successfully, says

Accenture.

The key message here is that:

> *Success with generative AI requires an equal attention on people and training as it does on technology...This means both building talent in technical competencies like AI engineering and enterprise architecture, and training people across the organization to work effectively with AI-infused processes (p.1).*

In a similar vein, Jonathan Aitken (2023) reminds us that:

> *The development of technology and its associated impact on job security has been a recurring theme since the industrial revolution. Where mechanisation was once the cause of anxiety about job losses, today it is more capable AI algorithms. But for many or most categories of job, retaining humans will remain vital for the foreseeable future.*

Aitken goes on to suggest that:

> *This means that, as workers, we need to look to harness the capability of AI systems and use them to their full potential. This means always questioning what we receive from them, rather than just trusting their output blindly...If we apply a sceptical mindset to how we use this new tool, we'll maximise its capability while simultaneously growing the workforce – as we've seen through all the previous industrial revolutions (pp.1-2).*

This message of careful regulation, focus and strategic planning will need to be applied to all impacts of AI technology and, as discussed below, this will be particularly important in working with AI tools in educational environments.

2.2. Control and Regulation

The flurry of concerned and critical activity referred to earlier surrounding the appearance of ChatGPT and other LLMs has been accompanied by feverish government action to keep pace with the rapid developments. The White House has issued a *Blueprint for an AI Bill of Rights*(2023) which includes the following five fundamental principles to guide future implementation and use of the new AI applications (pp.3-6):

- **You should be protected from unsafe or ineffective systems**
- **You should not face discrimination by algorithms and systems should be used and designed in an equitable way**
- **You should be protected from abusive data practices via built-in protections and you should have agency over how data about you is used**
- **You should know that an automated system is being used and understand how and why it contributes to**

outcomes that impact you

- **You should be able to opt out, where appropriate, and have access to a person who can quickly consider and remedy problems you encounter**

A similar set of policy principles have been proposed recently by the UK Government which have been described as a 'pro-innovation approach' designed to ensure that AI regulation does not interfere with investment in new technologies. Albert Sanchez-Graells (2023) has pointed out that the 'plans have been criticised for being too lax, already outdated, and lacking in meaningful detail' (p.1). He suggests that the policy proposals have more to do with ensuring a post-Brexit AI technology market for Britain rather than protecting the public from potential harm, and concludes that:

> *Only by implementing effective protections and showing strong and decisive action domestically can the UK government hope to build the credibility needed to lead international efforts of AI regulation (ibid., p.2).*

All such proposals seem to be rather light-touch - relying too on the goodwill of AI technology companies and the ability of the public to understand the new developments – and there is a growing consensus that regulation will need to be tighter. In terms of public understanding, the role of schools and colleges discussed below will be crucial. On the matter of state legislation in this field, the *European Commission's* legal framework governing medicinal products for human use (2023) might serve as a potentially effective model in this respect. In addition, all the current legislation for public protection in the sphere of medical practice and pharmaceutical development – which all countries seek to implement and monitor – can be effectively applied to AI policy alongside the regulation of the leading companies in the machine intelligence field. Such models – and the lessons learned from implementation over many years – will be foregrounded in the current moves to regulate and control AI on a global scale (Gomes, 2023).

Bostrom (2016) examines the 'control problem' in granular detail as the central element of AI safety, and outlines a range of methods (pp.175-6) that could be used to ensure that tools and applications are in alignment with our common values and goals. Anticipating some of the very worst scenarios that could result from a misalignment of superintelligent systems with human values, Bostrom (2003) observes:

> *It seems that the best way to ensure that a superintelligence will have a beneficial impact on the world is to endow it with philanthropic values. Its top goal should be friendliness. How exactly friendliness should be understood, how it should be implemented, and how the amity should be apportioned between different people and nonhuman creatures is a matter that merits further consideration. I would argue that at least all humans, and probably many other sentient creatures on earth should get a significant share in the superintelligence's beneficence (p.3).*

This values framework highlighted by Bostrom will be important in considering the role of AI applications within teaching and learning environments.

3. Teaching, Learning, and AI Perspectives

In addition to generic warnings about the dangers of unregulated AI issued by education workers in the machine learning field mentioned earlier, there are some critical commentators who dispute the power and potential of AI completely, arguing that its significance has been over-exaggerated. In this respect, Noam Chomsky's observations on recent developments are worth noting. In a recent article, Chomsky (2023) refers to the 'false promise of ChatGPT' arguing that – although such applications are 'marvels of machine learning' – the science of linguistics and epistemology indicate that 'they differ profoundly from how humans reason and use language (p.14). Similarly, Philip Goff (2023) argues forcefully that 'ChatGPT can't think – consciousness is something entirely different to today's AI (p.1). Against this, Bostrom's comments on contemporary developments are more favourable to AI sentience. In a recent interview, he claimed a fair degree of sentience for LLMs such as ChatGPT, and went on to say that:

> *I also think it's not doing them justice to say they're simply regurgitating text...They exhibit glimpses of creativity, insight and understanding that are quite impressive and may show the rudiments of reasoning... If an AI showed signs of sentience, it plausibly would have some degree of moral status... This means there would be certain ways of treating it that would be wrong, just as it would be wrong to kick a dog or for medical researchers to perform surgery on a mouse without anesthetizing it (Al-Sibai 2023, 1-2).*

3.1. AI and Consciousness

Commentators like Bostrom and educators taking a more positive line on AI developments recommend various strategies that will enable teachers to collaborate productively with the new tools and applications (Alhazmi, et al, 2023; Chen, 2023). Before examining aspects of such collaborative work, it is important to clarify the nature of AI uses in education within the framework of philosophical perspectives on consciousness, superintelligence and the nature and purposes of AI applications, and such an examination goes to the heart of the controversy about the nature of AI represented by the critical comments of Chomsky and Goff.

Differing positions on the consciousness or otherwise of AI will be substantially determined by the prior perspectives of commentators in relation to attempts to deal with the so-called 'hard problem' in this domain. David Chalmers (1996) outlines the 'easy' problems of consciousness, that is, how to map brain functions onto human thinking and behaviour. Such 'easy' problems include the integration of information by a cognitive system, the focus of attention, and the reportability of mental states, but such essentially functional processes leave us with the question of 'why the performance of these functions is accompanied by experience' (p.5). This is labelled by Chalmers the 'central mystery' (ibid) of consciousness and gives rise to the 'hard problem' of how to understand and explain the undisputed existence of subjective mental states in a world that science tells us consists only of physical objects. Panpsychism – the notion of consciousness, experience or awareness as a fundamental part of nature – has emerged as a prominent philosophical candidate for a realistic solution to the hard problem, and there are materialist and idealist versions of this general approach (Hyland, 2021).

To make headway in relation to the hard problem, Galen Strawson (2006) argues that it is necessary to introduce some notion of subjective experience into existing physical theories. Real physicalists according to Strawson, 'must accept that experiential phenomena are physical phenomena' (2006, p.1), and support the assertion concerning the emergence of experiential or consciousness properties from physical, non-experiential characteristics through, *inter alia*, the analogy of the emergence of the liquidity of water from non-liquid H2O molecules. A core aspect of this speculative thesis is that we do not know enough about the nature of the physical to argue – as dualists since Descartes and most post-Cartesian philosophers have held – that the physical and the mental are irrevocably distinct and irreconcilable. Making use of arguments by Eddington and Russell, Strawson asks 'on what conceivable grounds do so many physicalists simply assume that the physical, in itself, is an essentially and wholly non-experiential phenomenon?'(ibid., p.3).

Although physicalist materialism has been the foundation of science since the Enlightenment it has not gone unchallenged within philosophy where idealist theories of knowledge, truth and reality have been around since the Ancient Greeks. Shan Gao (2014) has produced a fascinating philosophical history of panpsychism which demonstrates how thinkers from the pre-Socratics, through Plato and Aristotle, and down through the Renaissance and Enlightenment periods to the current philosophy of science have advanced theories that propose that the natural world is imbued with, and indeed dependent upon, some form of conscious or mental element. However, in order to avoid the mind/body dualist black hole some form of monism needs to be considered, and Occam's Razor has led many thinkers – notably Leibniz, Berkeley and, in more recent times, Russell and Whitehead – to consider seriously the notion that, as Philip Goff (2019) puts it, 'consciousness is a fundamental and ubiquitous feature of physical reality' (p.112).

A principal materialist move is to assert that – since it is generally assumed that consciousness is generated by the brain – it is simply a matter of time before cognitive neuroscientists provide data that will solve the hard problem. However, as Steve Taylor (2019) has argued at length, there are no satisfactory models of how the mind/brain link can be supported, and he outlines the range of implausible claims – from epiphenomenalism to illusionism (pp.58-64) – which have failed to solve the principal problems. In addition, there is now a good range of neuroscientific data that indicates that – contra the physicalist assumptions – certain anomalous states of awareness (such as those produced by brain impairment, hallucinogenic episodes, or near-death experiences) result in *reduced* brain activity (ibid., pp.67ff.).

Along with the glaringly obvious implausibility of the notion that there might be neural correlates of the taste of coffee, the smell of a flower or the sound of falling rain, the reduction of brain activity in transcendent states of awareness is the exact opposite of what is entailed by the materialist assumption that all experience is generated by the brain. The realisation that metaphysical materialism has to be abandoned as an explanation of consciousness represents a courageous step but such a move has been made by Francis Crick's former colleague, Christof Koch (2014) who argues that the 'emergence of subjective feelings from physical stuff appears inconceivable' and that, rather than being produced by the circuitry of the brain, consciousness is 'inherent in the design of the universe' (p.28).

Moreover, as Bernardo Kastrup (2014) points out, there is a crucial difference 'between materialism as a*metaphysics* and scientific theories as *models'* (p.10). Scientific materialism observes patterns and regularities in nature and constructs

models that explain objects and forces – such as subatomic particles and negative electric charge – in terms of their relationship to other cognate constructions and issues only in *quantities* not the *qualities* of phenomenal experience. Explaining and predicting how aspects of the material world operate relative to other aspects reveals nothing about the fundamental aspects of nature. The analytical idealism proposed by Kastrup (2015) as a more cogent alternative is claimed to solve, or rather, dissolve the hard problem by positing a form of idealist panpsychism by which consciousness is the ultimate primitive.

The explanation of why we seem to be separate from the world and other beings is couched in terms of the idea of dissociated mind states drawn from well-established psychological studies. The brain, rather than generating experience, receives and canalizes information from the transpersonal world of the mind. Like whirlpools in the stream of consciousness, individual minds are a 'partial localization of the flow of experiences in the stream' (2014, p.82). This idea of subjective experience as individualised representations of transpersonal consciousness is further elaborated by Donald Hoffman (2019) in his theory of conscious realism.

Following Occam's simplest is best doctrine, the next logical step is to posit the idea that, as Hoffman (2019) prefers to say, it is consciousness itself – not spacetime, forces or material objects – that forms the fundamental basis of the cosmos. Hoffman argues that 'space, time and physical objects are not objective reality. They are simply the virtual world delivered by our senses to help us play the game of life' (p.xv). His ultimate claim – justified in terms of mathematical arguments rooted in evolutionary facts – is that, contra the physicalist case, it could be that 'consciousness does not arrive from matter...instead matter and spacetime arise from consciousness' (p.xviii). Labelled 'conscious realism', this theory 'claims no central role for human consciousness' but 'posits countless kinds of conscious agents with a boundless variety of conscious experiences' (p.201).

Applying the physicalist/idealist dichotomy to contemporary AI developments helps to clarify a number of key issues. Degrees and levels of consciousness outlined by a number of commentators on the hard problem allow us to place AI applications such as ChatGPT and Bard somewhere along a spectrum of consciousness. The precise location of various forms of AI on a consciousness spectrum will necessarily be an open-ended question informed by the theoretical frameworks outlined above. If we are inclined towards a more physicalist account, the placement might be of the following:

micro-organisms – lower animals – plants - primates (chimpanzees/bonobos) and mammals (dolphins/whales) – AI – humans

Idealist versions, on the other hand, might place AI lower on the spectrum:

micro-organisms – lower animals – plants – AI – primates and mammals - humans

However, noting the dangers of anthropocentrism mentioned earlier, it is feasible to place certain animals and plants at a level higher than humans in terms of animate intelligence (Lent, 2022) though, obviously, humans would rate the highest place on conceptual intelligence since this is a specifically human construction.

The lower position of AI on the consciousness spectrum for idealist positions may be explained by the problematic nature

of AI claims when considered within the monistic perspectives developed by Hoffman, Kastrup and similar commentators such as Taylor and Goff. For Kastrup – who conceives the subjective consciousness of humans and animals as individual segments of a larger mental membrane – the computational /information-processing power of AI tools must always fall short of consciousness since such material objects are just constructions of human minds which themselves are the only recipients of universal consciousness. Kastrup, like Hoffman, does not dispute the fact that AI applications will potentially outstrip human capabilities in all spheres – including art and poetry as well as science and mathematics – but, as a computer scientist, he suggests that there is absolutely no justification for entertaining the hypothesis that objects made from silicon chips running on algorithms designed by us can achieve the consciousness associated with Nagel's 'what-is-like-to-be' forms of subjective experience (Institute of Art and Ideas, IAI, 2023).

For Hoffman, we assess the consciousness of other beings through the portal provided by our evolutionary history. Thus, we gain hardly any information about the consciousness or otherwise of, say, rocks and minerals, a little information from invertebrates, even more from observing primates and mammals, and quite a lot from interacting with other humans as conscious agents. Hoffman is clear that the question does not turn on the difference between organic living organisms and inorganic synthetic objects, or between the intentions and motivations of humans as against the programmes of AI tools – all of this may be achievable through the complex algorithms of machine intelligence – but, rather, our interactions with other potential sources of consciousness in a universe of conscious agents (IAI, 2023). AI applications will be, like all other aspects of reality, part of our user interface which provides only partial and limited information about the ultimate nature and source of experience.

Broadly similar sentiments inform Goff's panpsychist perspective which allows for many types and degrees of consciousness throughout the natural world. In terms of the mainstream Integrated Information Theory of consciousness the processing of computers might not 'be integrated enough to give rise to consciousness' in terms of subjective awareness yet, Goff speculates, 'might such computers not be programmed to believe that they have feelings and experiences?' (2019, p.103). His recent comments which claim that ChatGPT is not an example of consciousness at work indicate that he thinks that current developments have not quite reached that stage (2023) but – especially in the light of the enormous speed of machine learning increments and in anticipation of major developments in quantum computing (Hyland, 2019) – both General AI and the subjective, experiential aspects of consciousness are potentially achievable in ways which satisfy the Turing test for emulating human capabilities (Kleppen, 2003; Cave, 2007).

3.2. AI, Intelligence and Searle's Chinese Room

A useful portal to the nature and purpose of AI implementation in educational contexts may be provided by a consideration of the philosopher, John Searle's, celebrated 'Chinese Room' thought experiment which is proposed as a powerful argument against 'strong AI' or general artificial intelligence (2004). The key thrust of the argument is explained clearly by John Horgan (2021):

> *[Searle] asks us to imagine a man (sic) who doesn't understand Chinese sitting in a room. The room contains a manual that tells the man how to respond to a string of Chinese characters with another string of characters.*

> *Someone outside the room slips a sheet of paper with Chinese characters on it under the door. The man finds the right response in the manual, copies it onto a sheet of paper and slips it back under the door. Unknown to the man, he is replying to a question, like "What is your favourite colour?," with an appropriate answer, like "Blue." In this way, he mimics someone who understands Chinese even though he doesn't know a word. That's what computers do, too, according to Searle. They process symbols in ways that simulate human thinking, but they are actually mindless automatons (p.2).*

Searle (2004) argues that the Chinese Room argument – claiming that mere computation or information-processing, no matter how complex or stunningly fast – 'strikes at the heart of the strong AI project (p.63), and he has been able to defend his position against some key philosophical objections (ibid., pp.69-71). However, it is worth noting that the argument here leans heavily on understanding language – particularly the differences between syntax (grammar) and semantics (meaning) – and, though applicable to the LLMs discussed earlier, might not be as relevant to the more general and advanced forms of machine learning currently emerging.

Also, in identifying language as the *sine qua non* of human reasoning the argument might be said to fall foul of the anthropocentric mistake of exalting human capabilities over those of the animal world in general. Jeremy Lent (2021) illustrates the folly of this mistake in his recent work in which he makes the useful distinction between what he describes as 'animate intelligence...the original AI' found in abundance in the natural world and in the lives of our ancestors, and 'conceptual intelligence' which is broadly those forms of reasoning distinctive of human behaviour and measured by IQ tests (pp.33-4). Lent goes on to explain that animate intelligence can be detected even at the level of microscopic cells which are observed to be 'acting purposefully to maintain and propagate' their lives. Similarly, sentient awareness of surroundings and intelligent behaviour is fully on display in the 'networked intelligence of plants' and animals throughout the natural world (pp.40ff). In the valuing of conceptual intelligence and consciousness based on human cognition – what is criticised as 'our mainstream culture's self-congratulatory obsession with humaniqueness' (p.50) – to the neglect of animate forms found in abundance throughout nature, we are, Lent argues, overlooking so much that can benefit the whole planet. What is required is an integration of animate and conceptual forms of awareness; as he puts it:

> *Perhaps the greatest challenge to human intelligence today is not how to accomplish the next technological breakthrough or build the most advanced AI, but how to integrate human ingenuity with our own animate intelligence and that of the natural world (p.55).*

Such an integration of different forms of intelligence can imaginatively inform the educational uses of AI within the framework of the human/AI partnership models discussed later. In addition, there is a particular response to Searle's argument that offers some creative insights into the nature and purpose of such potential strategies.

It will be recalled that Searle's argument turns on the contention that – although the person in the Chinese Room shows every sign of understanding Chinese because he/she can manipulate symbols and satisfactorily answer all questions in Chinese – that person does 'not understand a word of Chinese' (2004, p.63). This is because, like a computer, the

process of asking and answering questions in Chinese by means of algorithmic rule for translation consists only of the manipulation of symbols in a programme and this does not amount to understanding Chinese. Thus, though Searle freely admits that the person in the Chinese Room may indeed 'pass the Turing test for understanding Chinese' (ibid), the process still amounts to the computation of information, rather than genuine understanding.

However, though ultimately dismissed by Searle, one rejoinder to these claims is what he refers to as the 'Systems Reply' which goes as follows:

> *The idea of the Systems Reply is that though the man (sic) in the room does not understand Chinese, the man is only part of a larger system, consisting of a room, rule books, windows, boxes, programme, etc., and it is not the man but the whole system that understands Chinese (ibid., p.70).*

Although Searle claims that this scenario still amounts to the mere manipulation of symbols, this systems approach does accord with current philosophical perspectives on the inter-dependence of all aspects of experience and meaning (Hoffman, 2019; Lent, 2021) and, more to the point, offers valuable insights into the implementation of AI tools in educational settings.

As in the systems operating in the Chinese Room, in educational settings, teaching and learning are achieved by the interdependence of teachers, learners, curricula, classrooms, technology in the form of computers and audio-visual equipment, and all the infrastructure involved in the work of schools and colleges in their communities. The inclusion of AI tools in this interdependent network of resources simply expands the range and power of meaning and access to the knowledge, understanding and moral/affective elements which together contribute to positive educational experiences and development.

3.3. AI, Education and Non-Human Persons

The interdependent systems approach to AI in education requires the establishment of a particular ethical stance on the nature and purpose of the human/AI interface in teaching/learning contexts. With this in mind, I would suggest that we can gain much by looking at Peter Singer's conception of 'non-human persons'. Initially, Singer wanted to argue that certain animals – higher primates and mammals in particular – should be brought into the moral community on the grounds that they were sentient beings who could suffer (2009). In recent work, he has examined the status of AI machines and robots in light of the wider questions of rights and responsibilities surrounding the human/AI interface. Anticipating contemporary developments, Singer (2023) asks 'if machines can and do become conscious, will we take their feelings into account?' This is considered a pertinent question since our treatment of the only non-human sentient beings we have encountered so far – animals – gives no ground for confidence that we would recognize sentient robots not just as items of property, but as beings with moral standing and interests that deserve consideration (p.382).

After all, like corporations and similar artefacts, there are solid grounds for granting the legal personhood of AI tools and applications. As Visa Kurki (2019) argues, there is no reason to doubt that an 'AI can function as a legal person, it can be

granted legal personhood on somewhat similar grounds as a human collectivity' (p.175). It is well worth acknowledging all the increments in learning and development – in just about every sphere of activity – that have been made in recent decades through AI tools. The fact that humans have improved at the fiendishly complex game of Go since DeepMind's AlphaGo finally defeated the world's best players (Rosenblum, 2023) is just one small, perhaps emblematic, indication of what educators and policymakers might gain by working in partnership with AI applications.

However, any human/AI partnership – as Bostrom admits in his call for the recognition of the 'moral status' (Al-Sibai 2023, p.2) of superintelligent AI – must be regulated by the sort of control and alignment procedures referred to earlier. Bostrom provides a comprehensive analysis of control methods and offers what he calls the 'common good principle' for the regulation of the human/AI interface which is that 'superintelligence should be developed only for the benefit of all humanity and in the service of widely shared ethical ideals' (2016, p.312). Such ideals, stemming from the programming of 'friendliness' and 'philanthropic values' into advanced AI applications will be crucial in ensuring that 'all humans, and probably many other sentient creatures on earth...get a significant share in the superintelligence's beneficence' (Bostrom, 2023, p.3).

4. Implementing AI in Educational Settings: A Few Practical Ideas

As mentioned in earlier sections, there are a good many recommendations for the productive utilisation of AI applications in educational settings along with some worthwhile suggestions for practical implementation.

4.1. AI in Education: Some Suggestions for Practice

Christine Elgersma (2023, p.1), for example, suggests that, in dealing with concerns about plagiarism, ethical biases and misinformation over the use of AI, teachers might adopt such strategies as:

- Play with AI platforms yourself to understand what they can and can't do. Plug in your prompts and test them.
- Get to know students' writing as much as possible.
- Do some writing in class, staying mindful of some students' limitations to do so.
- Use formative assessment to get snapshots of progress over time.
- Evaluate your prompts and try to include a student-centred approach that features emotional intelligence and experience. As one teacher reported, "If you don't want generic answers, don't ask generic questions."
- Test AI tools together and discuss their limitations.
- As a last resort, there are tools you can use to detect AI, but be aware that they aren't perfect and can even sometimes flag plagiarism where it doesn't exist.

On the basis of research with teachers incorporating AI tools into their programmes, Elgersma came up with the following suggestions (ibid., p.2):

- Have students use coding to write their own algorithms.

- Have students write to a prompt and compare their results with those of an AI programme. Dig into the differences.
- Use AI to generate art, noting the origins of its materials.
- Get into the math behind algorithms and machine learning.
- Discuss the many uses AI has in science: How can something like an algorithm or generative AI help move scientific research forward?
- Talk about the impacts of new technologies—including AI—on civilizations and societies in history and today.

In a similar vein, Claire Chen (2023, pp.1-2), reporting on a conference on AI and education held at Stanford University recently, commented that – among all the usual warnings about the potential dangers in the misuse of AI – there was a general consensus that 'AI presents an opportunity to support teachers as they refine their craft at scale' through applications such as:

- **Simulating students:** AI language models can serve as practice students for new teachers. Percy Liang, director of the Stanford HAI Centre for Research on Foundation Models, said that they are increasingly effective and are now capable of demonstrating confusion and asking adaptive follow-up questions.
- **Real-time feedback and suggestions:** Dora Demszky, assistant professor of education data science, highlighted the ability of AI to provide real-time feedback and suggestions to teachers (e.g., questions to ask the class), creating a bank of live advice based on expert pedagogy.
- **Post-teaching feedback:** Demszky added that AI can produce post-lesson reports that summarize the classroom dynamics. Potential metrics include student speaking time or identification of the questions that triggered the most engagement. Research finds that when students talk more, learning is improved.
- **Refreshing expertise:** Sal Khan, founder of online learning environment Khan Academy, suggested that AI could help teachers stay up-to-date with the latest advancements in their field. For example, a biology teacher would have AI update them on the latest breakthroughs in cancer research, or leverage AI to update their curriculum.

Karen Lancaster's (2023) argument that universities should embrace AI essay-writing software draws interesting parallels between ChatGPT applications and the acceptance by maths teachers of the use of calculators in lessons, homework and exams. As she remarks:

> *By the time I was in secondary school in the early '90s, calculators were a common addition to the maths classroom, and with their help, pupils were doing mathematics which would have been far beyond what our parents could have done with a pen and paper. But the other pupils and I weren't* ***cheating*** *– we were using technology to help us grapple with more complex material. One still needed to learn the proper ways to use the calculator to obtain the right answer. In short, the pocket calculator changed some of the skills that were necessary in the maths classroom (p.1).*

She goes on to argue:

> *At the moment, universities consider the use of essay mills or AI essay writing to be plagiarism. However, these are not equivalents: students submitting work produced by essay mills have had* ***no part*** *in writing the text: they are not the ostensible author. However, if a student uses AI essay-writing technology appropriately to enable them to write an excellent essay (they edit, improve, and rework the essay that the AI produced) then this would seem to be a worthwhile skill, in the same way that using a calculator to solve complex equations or using the internet to research a topic is laudable (ibid.)*

Although I would want to make some sort of distinction between using a calculator as an ancillary tool to solve a larger maths problem, and the use of AI to complete a whole *assignment* – including research, sources, academic writing style etc. – as a *finished piece*, I would tend to agree with the conclusion that 'educational establishments accept that the software exists, and teach students how to use it effectively, because AI still – currently, at least – has its flaws' (ibid),

Similar points are made by Andy Phippen (2023) in his argument that:

> *Banning students from using ChatGPT, or expecting teachers to scour homework for its use, would be shortsighted. Education has adapted to – and embraced – online technology for decades. The approach to generative AI should be no different (p.1).*

He goes on to express the hopeful expectation that we will:

> *see generative AI tools being merged, eventually, into mainstream learning. Saying "do not use search engines" for an assignment is now ridiculous. The same might be said in the future about prohibitions on using generative AI... Perhaps the homework that teachers set will be different. But as with search engines, word processors and calculators, schools are not going to be able to ignore their rapid advance. It is far better to embrace and adapt to change, rather than resisting (and failing to stop) it (ibid).*

However, it is one thing to advance such proposals about embracing AI technology and incorporating it into mainstream teaching and learning, and quite another to work out specific strategies for achieving such inclusion in practice. In the following section, I offer a tentative recommendation for practice drawn from my own field of teaching basic philosophy courses in university adult education to students without formal qualifications.

4.2. Using AI in a Philosophy Ethics Module

The module would include a discussion of the nature and origins of ethical discourse – ranging from the Ancient Greek notions of virtue ethics to modern emotivist, contractualist Kantian and utilitarian theories – and, following this foundationalist groundwork, examine the famous 'Trolley Problem' thought experiment devised by Philippa Foot (1967) as an example of bringing out the subtleties, problems and nuances of modern ethical theories and moral discourse in general (Hyland, 2020).

Foot's original thought experiment was outlined in an article on abortion and the double effect principle. She notes:

> *Suppose that a judge or magistrate is faced with rioters demanding that a culprit be found for a certain crime and threatening otherwise to take their own bloody revenge on a particular section of the community. The real culprit being unknown, the judge sees himself as able to prevent the bloodshed only by framing some innocent person and having him executed. Beside this example is placed another in which a pilot whose aeroplane is about to crash is deciding whether to steer from a more to a less inhabited area. To make the parallel as close as possible it may rather be supposed that he is the driver of a runaway tram which he can only steer from one narrow track on to another; five men are working on one track and one man on the other; anyone on the track he enters is bound to be killed. In the case of the riots the mob has five hostages, so that in both the exchange is supposed to be one man's life for the lives of five. The question is why we should say, without hesitation, that the driver should steer for the less occupied track, while most of us would be appalled at the idea that the innocent man could be framed (1967, p.3).*

The principle of double effect comes into operation due to the fact that, in choosing the lesser of two evils, we cannot but avoid immoral acts. However, intentions are crucial in these cases, the pilot does not intend to crash the plane no more than the tram driver is aiming to kill a man working on the track. Unlike the judge in Foot's example, they are presented with a *fait accompli* – if they don't act, more people will be killed than if the agents choose other courses of action.

Foot's trolley problem has achieved unprecedented fame and celebrity within philosophical ethics leading to a specialised domain known as "Trolleyology" (Danaher, 2014). Danaher expresses the key dilemmas in his outline of the two 'classic' versions of the problem (p.1):

> ***Switch:*** *A trolley car is hurtling out of control down a train track. If it continues on its current course, it will collide with (and kill) five workers who are on the track. You are standing beside the track, next to a switch. If you flip the switch, the trolley will be diverted onto a sidetrack, where it will collide with (and kill) one worker. Do you flip the switch?*
>
> ***Footbridge****: A trolley car is hurtling out of control down a train track. If it continues on its current course, it will collide with (and kill) five workers who are on the track. You are standing on a footbridge over the track, next to a very fat man. If you push him off the footbridge, he will collide with the trolley car, slowing it down sufficiently to save the five workers. He, however, will die in the process. Do you push the fatman?*

In psychological experiments, 87% of respondents said they would flip the switch but only 31% would be willing to push the fat man (ibid., p.2). It seems to be that our moral intuitions make us reluctant to engage personally with obviously immoral actions whereas remote action at a distance is apparently more acceptable. Variations on the original experiment have involved making the one worker on the track a young person (or even a close relative) whereas the five on the other track are old people and the resulting responses vary predictably.

AI Assignment Protocol on the Trolley Problem

Given the current shortcomings and potential limitations of current AI tools, it will be important to address the issue of formulating different ways of asking questions about the ethical thought experiment. In addition, with the aim of fostering a critical engagement with different versions and capabilities of AI applications, the process will involve comparing and contrasting the responses of OpenAI's ChatGPT with those of Google's Bard and Socratic, and Microsoft's AI-powered Bing (Laurent, 2023). Such a critical approach should also take note of James Intrilligator's (2023) important point that LLMs such as ChatGPT are not search engines like the standard Firefox and Google applications. As he comments:

> *Unlike a search engine, with static and stored results, ChatGPT never copies, retrieves or looks up information from anywhere. Rather, it generates every word anew. You send it a prompt, and based on its machine-learning training on massive amounts of text, it creates an original answer (p.1)*

With this in mind, the correct way to use such tools is in an iterative fashion by asking a succession of questions on a particular topic by refining each successive question on the basis of previously AI-generated responses. Intrilligator, calls this 'gliding, not searching' and explains that:

> *The interaction with ChatGPT is best performed not as a simple or undirected question-and-answer session, but as an interactive conversation that progressively builds knowledge for both the user and the chatbot. The more information you provide to it about your interests, and the more feedback it gets on its responses, the better its answers and suggestions. The richer the journey, the richer the destination (ibid., p.2).*

Keeping all this in mind, the Trolley Problem assignment might incorporate the following stages:

Stage 1

Students are given the same title/question to ask, such as,

Describe Philippa Foot's 'Trolley Problem' and explain its relevance in ethical discourse.

This exact question is then given to four groups of students, each group assigned a different LLM – ChatGPT, Bard, Socratic, or Bing. Responses are compared and contrasted with differences and inconsistencies noted and analysed. What can be learned from all this? In line with the 'gliding' iterative strategy referred to earlier, students are subsequently asked to pose further – specifically refined – questions about the various responses to the Trolley Problem already elicited from the AI applications.

Stage 2

Students are further asked to use ChatGPT in constructing assignments on the general topic of the Trolley Problem using different forms of words to interrogate the issue. Some examples might include:

What are the main points Foot wants to make with her thought experiment?

Does the Trolley Problem defeat utilitarian ethics?

Explain why so many people would sacrifice one person to save five yet not push one person off a bridge to achieve the same outcome.

Does virtue ethics or Kant's duty morality help to solve the Trolley Problem?

Given what commentators such as Lancaster, Chen, Elgersma, and others mentioned above observe about the shortcomings and current flaws of AI tools such as ChatGPT, the next stage of the learning process would ask students to compare and contrast the different answers to the questions posed on the Foot thought experiment. Such an examination will take into account what has been learned from the Stage 1 activity of asking the same question using four different AI applications. Are there any obvious flaws, inconsistencies and other weaknesses in the responses provided? Are the shortcomings the same as those brought out in Stage 1 and how can the whole process be improved? As Heaven (2023) noted in research with teachers using ChatGPT:

> *teachers are finding opportunities in the misinformation and bias that large language models often produce. These shortcomings can kick off productive discussions...The fact that it's not perfect is great (p.2).*

Stage 3

In terms of assessment, this comparative and critical analysis can itself function as a formative means of evaluating learner progress, understanding and development. In addition, some form of oral questioning – either individually or in groups – might supplement this process if there is a requirement for summative as well as formative assessment. This whole enterprise can be seen as an attempt – recommended by educators taking a positive stance on work in the AI field – to embrace the new technology and form an educational/AI collaborative partnership to enhance teaching and learning in all spheres. It can be seen as one way of implementing the all-encompassing systems strategy of fostering collective consciousness described in earlier sections.

5. Coda

In conclusion, a rational and pragmatic approach to the use of AI applications in education seems to be justified by the current evidence. A collaborative partnership – between teachers, learners and non-human AI persons (Hyland, 2023) – is well worth considering as a viable operational model, always providing the control and regulation elements recommended by Bostrom and others are built into the new applications so as to ensure that the desirable 'common good principles' are always front and centre in all activities.

References

- Aitken, J. (2023, May 25). AI could threaten some jobs, but it is more likely to become our personal assistant.*The Conversation*. Retrieved June 18, 2023, fromhttps://theconversation.com/ai-could-threaten-some-jobs-but-it-is-more-likely-to-become-our-personal-assistant-206297
- Alhazmi, A. K., Alhammadi, F., Zain, A. A., Kaed, E., & Ahmed, B. (2023). AI's Role and Application in Education: Systematic Review. In A. K. Nagar, D. Singh Jat, D. K. Mishra, & A. Joshi (Eds.), *Intelligent Sustainable Systems* (Vol. 578). Singapore: Springer. https://doi.org/10.1007/978-981-19-7660-5_1. Retrieved July 23, 2023
- Al-Sibai, N. (2023, April 13). Nick Bostrom says that AI Chatbots may have some degree of sentience.*The Byte*. Retrieved July 29, 2023, from https://futurism.com/the-byte/nick-bostrom-ai-chatbot-sentience
- BBC. (2023, May 18). BT to cut 55,000 jobs with up to a fifth replaced by AI*BBC News*. Retrieved June 9, 2023, from https://www.bbc.com/news/business-65631168
- Bostrom, N. (2016). *Superintelligence*. Oxford University Press.
- Bostrom, N. (2003). Ethical Issues in Advanced Artificial Intelligence. Retrieved July 7, 2023, from https://nickbostrom.com/ethics/ai
- Cave, P. (2007). *Can a Robot be Human?*Oneworld Publications.
- Chalmers, D. (1996). *The Conscious Mind*. Oxford University Press.
- Chen, C. (2023, March 9). AI will transform and teaching and learning. Let's get it right. Stanford University AI+ Education Summit. Retrieved June 7, 2023, from https://hai.stanford.edu/news/ai-will-transform-teaching-and-learning-lets-get-it-right
- Chomsky, N. (2023, March 8). The False Promise of Chat GPT.*New York Times*. Retrieved June 13, 2023, from https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgptai.html
- Danaher, J. (2014). Advanced Trolleyology and the Doctrine of Double Effect.*Philosophical Disquisitions*. Retrieved July 9, 2023, from https://philosophicaldisquisitions.blogspot.com/2014/02/advanced-trolleyology-and-doctrine-of.html
- Elgersma, C. (2023). ChatGPT and Beyond: How to Handle AI in Schools.*Commonsense Education*. Retrieved July 11, 2023, from https://www.commonsense.org/education/articles/chatgpt-and-beyond-how-to-handle-ai-in-schools
- European Commission. (2023). Legal framework governing medicinal products for human use in the EU. Retrieved July 23, 2023, from https://health.ec.europa.eu/medicinal-products/legal-framework-governing-medicinal-products-human-use-eu_en
- Foot, P. (1967). The Problem of Abortion and the Doctrine of Double Effect.*Oxford Review*, 5, 1-7.
- Foucart, R. (2023, April 27). Tech giants forced to reveal AI secrets – here's how this could make life better for all.*The Conversation*. Retrieved July 19, 2023, fromhttps://theconversation.com/tech-giants-forced-to-reveal-ai-secrets-heres-how-this-could-make-life-better-for-all-204081
- Future of Life Institute. (2023). Pause Giant AI Experiments – An Open Letter. Retrieved July 21, 2023, from https://futureoflife.org/open-letter/pause-giant-ai-experiments/
- Gao, S. (2014). *Consciousness Encounters Quantum Physics*. Amazon Kindle Direct Publishing.
- Goff, P. (2019). *Galileo's Error: Foundations for a New Science of Consciousness* Rider.

- Goff, P. (2023, May 17). ChatGPT can't think – consciousness is something entirely different to today's AI. *The Conversation*. Retrieved May 22, 2023, from https://theconversation.com/chatgpt-cant-think-consciousness-is-something-entirely-different-to-todays-ai-204823
- Gomes, A. M. (2023, June 13). Global trends in Artificial Intelligence regulation. Retrieved July 20, 2023, from https://www.cullen-international.com/news/2023/06/Global-trends-in-Artificial-Intelligence-regulation.html?gclid=EAIaIQobChMIMXf_vcgAMV5z0GAB3vfAFfEAAYAiAAEgK8NPD_BwE
- Heaven, W. D. (2023, April 6). ChatGPT is going to change education, not destroy it. *MIT Technology Review*. Retrieved May 19, 2023, from https://www.technologyreview.com/2023/04/06/1071059/chatgpt-change-notdestroy-education-openai/
- Hoffman, D. (2019). *The Case Against Reality: How Evolution Hid the Truth from Our Eyes* Penguin.
- Horgan, J. (2021, March 9). Quantum Mechanics, the Chinese Room Experiment and the Limits of Understanding. *Scientific American*. Retrieved August 1, 2023, from https://www.scientificamerican.com/article/ quantum-mechanics-the-chinese-room-experiment-and-the-limits-of-understanding/
- Hyland, T. (2019). *Adventures in the Matrix*. Scholars' Press.
- Hyland, T. (2020). Telling Moral Tales: Exploring Ways of Enhancing the Realism and Explanatory Power of Ethical Thought Experiments. *Open Journal of Social Sciences*, 8, 304-327.
- Hyland, T. (2021). *Perspectives on Panpsychism*. Scholars' Press.
- Hyland, T. (2023, June 13). AI Applications and Non-Human Persons. *Philosophy of Education Society of Great Britain Blog*. Retrieved August 1, 2023, from https://www.philosophy-of-education.org/ai-applications-and-non-human-persons/
- IAI. (2023). Consciousness in the Machine: Kastrup, Schneider, Hoffman Debate on Consciousness. *Institute of Art and Ideas*. Retrieved August 1, 2023, from https://www.youtube.com/watch?v=VmQXpKyUh4g
- Intrilligator, J. (2023, July 21). Gliding, not searching: Here's how to reset your view of ChatGPT to steer it to better results. *The Conversation*. Retrieved August 1, 2023, from https://theconversation.com/gliding-not-searching-heres-how-to-reset-your-view-of-chatgpt-to-steer-it-to-better-results-205819
- Kastrup, B. (2014). *Why Materialism is Baloney*. iff Books.
- Kastrup, B. (2015). *Brief Peeks From Beyond: Critical Essays on Metaphysics, Neuroscience, Free Will, Skepticism and Culture*. iff Books.
- Kelly, J. (2023, May 31). Goldman Sachs Predicts 300 Million Jobs Will Be Lost Or Degraded By Artificial Intelligence. *Forbes*. Retrieved June 12, 2023, from https://www.forbes.com/sites/jackkelly/2023/03/31/goldman-sachs-predicts-300-million-jobs-will-be-lost-or-degraded-by-artificial-intelligence/
- Kleppen, E. (2023, January 13). What is the Turing Test? *Built In*. Retrieved July 25, 2023, from https://builtin.com/artificial-intelligence/turing-test
- Koch, C. (2014). Is Consciousness Universal? *Scientific American Mind*, 25, 26-29.
- Kurki, V. A. J. (2019). *The Legal Personhood of Artificial Intelligences*. Oxford Academic. Retrieved July 29, 2023, from https://academic.oup.com/book/35026/chapter/298856312?login=false
- Lancaster, K. (2023, February 24). Why universities should embrace AI essay-writing software. *PESGB Blog*. Retrieved June 5, 2023, from https://www.philosophy-of-education.org/why-universities-should-embrace-ai-essay-writing-

software/

- Laurent, C. (2023). Discover Top AI Chat Bots: Alternatives to ChatGPT. Retrieved July 31, 2023, from https://filmora.wondershare.com/chatgpt/chatbots-like-chatgpt.html?gclid=EAIaIQobChMI-PTTwde4gAMVGrjtCh1jhwCxEAAYBSAAEgIgrvD_BwE
- Lent, J. (2021). *The Web of Meaning*. Profile Books.
- Phippen, A. (2023, July 28). ChatGPT isn't the death of homework – just an opportunity for schools to do things differently. *The Conversation*. Retrieved July 18, 2023, fromhttps://www.philosophy-of-education.org/why-universities-should-embrace-ai-essay-writing-software/
- Rosenblum, A. (2023, March 13). Humans have improved at go since AIs became best in the world*New Scientist*. Retrieved July 17, 2023, from https://www.newscientist.com/article/2364137-humans-have-improved-at-go-since-ais-became-best-in-the-world/
- Sanchez-Graells, A. (2023, June 7). The UK wants to export its model of AI regulation, but it's doubtful the world will want it. *The Conversation*. Retrieved July 15, 2023, fromhttps://theconversation.com/the-uk-wants-to-export-its-model-of-ai-regulation-but-its-doubtful-the-world-will-want-it-206956
- Searle, J. R. (2004). *Mind*. Oxford University Press.
- Shad, N. (2023). Schools bewildered by AI advances, say head teachers. *BBC News*. Retrieved May 20, 2023, from https://www.bbc.com/news/uk-65655010
- Singer, P. (2009). *Animal Liberation*. HarperPerennial Modern Classics.
- Singer, P. (2023). *Ethics and the Real World*. Princeton University Press.
- Strawson, G. (2006). Realistic Monism: Why Physicalism Entails Panpsychism.*Journal of Consciousness Studies*, 13(10/11), 3-31; also Strawson, G. (2016, May 16). Consciousness isn't a mystery. It's Matter. *New York Times*.
- Taylor, S. (2019). *Spiritual Science: Why Science Needs Spirituality to Make Sense of the World* Watkins Media.
- White House. (2023). Blueprint for an AI Bill of Rights. Retrieved June 11, 2023, from https://www.whitehouse.gov/ostp/ai-bill-of-rights/
- World Economic Forum. (2023, May 4). These are the jobs most likely to be lost – and created – because of AI. Retrieved July 12, 2023, from https://www.weforum.org/agenda/2023/05/jobs-lost-created-ai-gpt/