Rethinking the Ethics of GenAI in Higher Education: A Critique of Moral Arguments and Policy Implications

Journal of
APPLIED PHILOSOPHY

Journal of Applied Philosophy
doi: 10.1111/japp.70026

Rethinking the Ethics of GenAI in Higher Education: A Critique of Moral Arguments and Policy Implications

KARL DE FINE LICHT

ABSTRACT *This article critically examines the moral arguments for restrictive policies regarding student use of generative AI in higher education. While existing literature addresses various concerns about AI in education, there has been limited rigorous ethical analysis of arguments for restricting its use. This article analyzes two main types of moral arguments: those based on direct difference-making (where individual university actions have measurable impacts) and those centered on non-difference-making participation (where symbolic participation in harmful systems matters regardless of direct impact). Key concerns examined include environmental harm from AI energy consumption, exploitative labor practices in AI development, and privacy risks. Through careful analysis, the article argues that these arguments face significant challenges when examined in depth. The difference-making arguments often fail to establish that individual university actions meaningfully contribute to claimed harms, while the non-difference-making arguments lead to impractical conclusions when applied consistently across university operations. Rather than supporting blanket restrictions, the analysis suggests universities should focus on fostering responsible AI engagement through ethical guidelines, licensed tools, and education on responsible use. The article concludes that a balanced approach considering both moral and practical factors is more effective than restrictive policies in addressing ethical concerns while preserving educational benefits.*

1. Introduction

AI systems have long been used at universities for a variety of purposes. For example, coordinators use AI tools to streamline administrative tasks, such as scheduling classes or analyzing student enrollment patterns.[1] AI systems can help predict course demand based on historical data and student preferences.[2] They can also be used by university counselors to detect students who are at risk of failing,[3] to name just a few examples. Since late 2022, debate has intensified around *student use* of generative AI (GenAI) in higher education.[4] Supporters argue these tools enhance learning through instant feedback and varied explanations,[5] while providing guided experience with AI could improve ethical use post-graduation.[6] Critics worry about academic dishonesty, diminished critical thinking,[7] and broader ethical concerns including climate impact, privacy, labor exploitation, algorithmic bias, and digital inequality.[8]

Even though there is significant discussion about the use of GenAI and the policies to adopt in higher education, surprisingly little *in-depth moral reasoning* is presented on the topic.[9] Some papers attempt to address this issue, but most are either reviewing the debate or, at best, offering conclusions as an afterthought.[10] These papers do not build a solid case either for or against specific rules or policies. The same is true for the papers that examine student and faculty attitudes towards GenAI.[11] Even those that do explore moral arguments or attitudes for or against GenAI in higher education often limit themselves to niche or controversial notions that may not apply universally across educational contexts. For instance, some focus on a single Kantian principle,[12] which makes them problematic to use in policy situations since these individual normative principles are quite controversial. Similarly, Paglieri engages with various ethical concerns, but the analysis lacks depth in terms of directly addressing the most pressing moral issues specific to higher education.[13] My earlier article on the subject of GenAI takes a more comprehensive approach, but the depth of analysis could much be improved as the article does not focus exclusively on the moral arguments, thereby leaving key ethical considerations underexplored, and does not elucidate the underlying structure of the arguments.[14]

To address this gap in the literature, I aim to analyze the structure and validity of moral arguments concerning GenAI in higher education with a focus on the case presented in my earlier article to make the discussion more transparent and rigorous.[15] The structure of the argument proceeds in three main steps. First, I analyze two distinct ways moral arguments against GenAI can be interpreted: one focused on direct difference-making (where individual university actions have measurable impacts) and another centered on non-difference-making participation (where, for example, symbolic participation in harmful systems matters regardless of direct impact). Second, I examine the specific content of arguments including environmental harm, exploitative labor practices, and privacy concerns, together with student learning. Here I argue that these arguments fail to show that individual universities have strong moral reasons to ban or heavily restrict student use of GenAI. Specifically, I will challenge whether these moral arguments provide sufficient grounds for adopting very restrictive policies – or for discouraging the use of GenAI, as is already the case at over 27% of R1 universities in the United States.[16] Finally, I argue that even if we accept these moral arguments, they fail to justify heavily restrictive policies due to both that the universities might face genuine moral dilemmas and that these arguments become a *reductio* against themselves when applied consistently across university operations. Universities will not be able to do much if they are to abide by these standards.

The article is structured as follows. First, I clarify the structure and outline the content of the moral arguments that support very restrictive policies and bans regarding GenAI, focusing on the environmental, privacy, labor, and academic concerns often cited by critics. Then I engage in a critical examination of these arguments, questioning their feasibility and whether they can be applied consistently across higher education settings without leading to absurd or impractical outcomes. Lastly, in the conclusion, I suggest a balanced approach that acknowledges both the moral concerns and the practical advantages of GenAI, advocating for policies that enable responsible engagement with these tools rather than prohibition.

2. The Moral Arguments Against the Use of GenAI and the Case for Restrictive Rules

A range of arguments can be advanced in favor of banning or heavily restricting the use of GenAI tools. These fall into two broad categories: indirect and direct. Indirect arguments claim that students act wrongly by using these tools – such as by participating in morally problematic processes – which justifies a ban. Aylsworth and Castro, for example, argue that using tools like ChatGPT to write student papers, especially in the humanities, is morally objectionable.[17] In contrast, direct arguments assert that universities themselves act wrongly by allowing such use, regardless of whether students individually do anything wrong. This article focuses exclusively on direct reasons for banning or heavily restricting GenAI use – cases where institutional responsibility is central. While indirect concerns are important, my analysis centers on whether universities, as institutions, have a moral duty to restrict GenAI due to broader concerns such as environmental harm, labor exploitation, and privacy risks. This focus enables a clearer ethical evaluation of whether restrictive institutional policies are justified independently of student behavior.

Furthermore, the direct arguments can generally be divided into two groups: those where the university has a moral obligation to implement very restrictive policies or completely ban student use of GenAI because it is harmful to others, such as the environment, and those where the use of GenAI is harmful to the students themselves. While these arguments are interconnected, they can, to some extent, be conceived as separate and may need to be addressed in different ways. The first two arguments discussed here are other-regarding, focusing on harm to others, while the third is self-regarding, focusing on harm to students. Even though I only address three of the most prominent arguments in the debate,[18] a similar line of reasoning can be applied to other arguments in the debate by utilizing the same logic, such as those about scraping of data.[19]

The first argument against the use of GenAI in higher education that is going to be analyzed in this article is its association with exploitative labor practices.[20] The training phase of these AI models often involves human workers who are tasked with moderating content or interacting with early, problematic versions of these systems, frequently under poor working conditions and for minimal compensation.[21] For instance, workers in regions such as Nigeria and Kenya have been required to engage with offensive, traumatic, or otherwise harmful content to refine AI models, while earning wages as low as $1.50 per hour and lacking adequate labor protections. Regilme conceptualizes this dynamic as a form of 'AI colonialism', in which the Global South provides undervalued labor and bears the environmental and social costs of AI development, while wealth and technological benefits accrue in the Global North.[22] He identifies a dual structure of injustice – *necroexportation*, where environmental and labor harms are displaced onto marginalized populations, and *necrostratification*, where the life prospects of AI workers are systemically diminished. These harms go beyond distributive injustice and raise concerns about the erosion of fundamental human rights. If universities adopt these AI tools without critically engaging with the conditions of their production, it is argued, they risk legitimizing and reproducing these structural injustices, undermining their commitments to equity, justice, and ethical integrity in education.

Another argument against the use of GenAI in higher education centers on its environmental unsustainability. Large language models require substantial computational resources for both training and deployment, resulting in significant energy consumption and associated carbon emissions – contributing to climate change at a time when universities increasingly emphasize sustainability.[23] Training a single model, such as GPT-4, can consume around 50 GWh of electricity, equivalent to the annual energy use of 2,000 American households.[24] These systems also rely on large-scale data centers that continue to draw energy post-training, further compounding their environmental impact. In addition, data centers consume vast quantities of fresh water for cooling – training GPT-3 alone required an estimated 700,000 liters – raising concerns about water depletion and resource competition.[25] Beyond energy and water use, GenAI depends on hardware built from rare earth elements and other critical minerals. The extraction and processing of materials such as cobalt, lithium, and nickel are linked to deforestation, toxic waste, and biodiversity loss. Mining operations in regions like the Democratic Republic of the Congo and the Philippines have been associated with severe soil contamination and water pollution. Taken together, the high energy demands, intensive water use, and environmentally destructive material inputs of GenAI pose serious sustainability concerns. Promoting its widespread use in higher education risks aligning institutions with practices that contradict their climate goals and environmental commitments.

The final key moral argument against the use of GenAI in higher education to be analyzed here relates to concerns over privacy.[26] These AI tools often require data input from users, which may include sensitive or personal information. The collection and use of such data, particularly without explicit informed consent, is ethically problematic, as it puts students' privacy at risk. Students might unknowingly share information that could be used for purposes beyond their control, especially since many AI companies prioritize market advantage over user privacy. This lack of transparency and control over how data is processed poses a significant ethical issue, particularly in academic environments where privacy should be paramount. Limiting the use of GenAI tools could thus be justified to protect student data and prevent privacy violations, ensuring that the educational environment remains a safe and ethical space for learning. This argument is also closely tied to the argument from student learning. According to this argument we should have very restrictive regulations to protect students' learning.[27]

When evaluating these arguments, we need to be careful about identifying the actor or agent, as this will significantly affect whether the effects of using GenAI can plausibly be described as negative. It will also influence the extent to which we actually participate in something morally objectionable. In this article, we focus on the *singular university* as the agent. Therefore, for the first strand of arguments to succeed, the actions of the university must have a direct negative impact, such as contributing to environmental harm or enabling exploitative practices, just to mention two examples. It is not enough to argue that GenAI used by universities in general, or by most universities in the world, has this effect; it must be the case that the individual university in question and its actions contribute to these negative outcomes. According to the second argument, the university should not take part in exploitation or environmental destruction, even if it does not directly cause these harms. Thus, from this perspective, the moral duty to refrain from participation exists regardless of whether the university's actions alone make a significant difference. This duty is grounded in the principle that universities, as moral agents, should avoid benefiting from or being complicit in unethical practices, irrespective of their individual contribution to the broader harm. Proponents of the argument against complicity would emphasize that participation is inherently wrong because, for example, it symbolically endorses harmful or unjust practices. Therefore, the ethical argument against complicity remains potent, as it is rooted in the commitment to moral integrity rather than the measurable effect of the action.

If we assume that it is morally wrong to utilize these tools, the next step is to show that we have strong reason to ban or be very restrictive when it comes to their use. This means demonstrating that being heavily restrictive towards GenAI is the preferable option for policymakers when compared to other alternatives in the given context. Perhaps the strongest case in favor of this position would be to argue that a ban directly addresses and mitigates the harms identified – such as exploitation, environmental impact, and privacy violations – in a manner that other solutions cannot achieve as effectively.[28] A very restrictive approach, it could further be argued, would prevent participation in and benefit from practices that contribute to exploitation and environmental degradation, thereby upholding the university's moral duty to avoid complicity in harmful activities. Moreover, by preventing GenAI use entirely, the risks associated with privacy breaches could be eliminated, creating a clearer and more ethically straightforward policy environment for both students and educators.

The argument in favor of a (very) restrictive approach to student use of GenAI in higher education can be summarized as follows:

1. The university has a strong moral reason to adopt a very restrictive policy regarding student use of GenAI.
2. When the university has a strong moral reason to adopt a very restrictive policy regarding x, it has a strong *pro tanto* reason to adopt a very restrictive policy regarding x.
3. Therefore, the university has a strong *pro tanto* reason to adopt a very restrictive policy against student use of GenAI.

While premise 2 might be criticized on the grounds that it overlooks non-moral considerations – such as those highlighted by Bernard Williams and Joseph Raz, who argue that moral reasons do not always override other practical concerns[29] – I will set aside this complexity here. For the purposes of this argument, I will proceed as if strong moral reasons should lead to very restrictive policies, without further exploring the interplay of non-moral considerations.

3. Critique of the Arguments

In this section, I will discuss the difference-making arguments and the non-difference-making arguments in turn.

3.1. The Difference-Making Arguments

Beginning with the argument from climate change, the idea is that if the university allows or is not very restrictive of student use of GenAI, this will contribute to climate change. This is because of the energy use when the models are trained and when they are up and running.[30] The underlying assumption is *not* that climate change would cease if the university enforced stricter policies. Rather, the argument suggests that lenient policies could accelerate the onset of catastrophic weather events or worsen the adverse effects of climate change – that is, leading to earlier onset and greater harms.[31]

This argument, however, has some weaknesses. First, companies *do not* create models specifically for each university; rather, they build large, general-purpose models that are accessible across institutions. Even if universities were to limit student use, it would not necessarily reduce the demand for, or the training frequency of, GenAI models globally. If universities were each to build and train their own language models, there might indeed be a stronger environmental argument against the use of GenAI. Yet, since they typically rely on pre-trained models, the training phase energy costs can, to some extent, be discounted. Additionally, the energy usage of GenAI is relatively minor compared to other significant sources of carbon emissions. For example, while the energy required to train large language models like GPT-4 is substantial, this is a small portion of the total energy footprint in the broader context. In the United States alone, there are approximately 145 million housing units compared to the 2,000 it took to train GPT-4, and the energy demands of industrial processes, transportation, and traditional energy-intensive technologies far exceed those of GenAI models. Furthermore, newer models like GLaM have shown that it is possible to drastically improve efficiency; despite having seven times the parameters of GPT-3, it used 2.8 times less energy to train.[32]

Moreover, the current focus on making AI systems more energy-efficient could further limit their environmental footprint, making concerns about their impact on climate change less pressing when it comes to training, and even more so with regard to running the systems.[33] Many developers are working on improving the efficiency of training processes, optimizing model architectures, and developing hardware that consumes less power. For instance, newer GenAI models, such as Meta's Llama, are designed to operate with relatively modest energy needs during deployment, which means that while the training process is energy-intensive, the day-to-day use is becoming increasingly efficient.[34] The same is true for the new Chinese model DeepSeek. These improvements suggest that the overall environmental impact of GenAI could decrease over time, making it even less of a concern in terms of contributing to climate change. In addition, the substitution effects of AI may result in greater net energy savings. For example, generating text or images with AI can be 130 to 2,900 times *less* carbon-intensive than human labor in equivalent tasks, which could outweigh direct energy costs in specific applications.[35] Therefore, while the energy consumption of GenAI is indeed significant in absolute terms, the argument that it has a major impact on climate change is, at least currently, overstated. Again, compared to larger-scale factors contributing to carbon emissions, GenAI's role is relatively small, and ongoing improvements in energy efficiency will likely make it even less of a contributing factor. This perspective casts some doubt on the argument that the environmental effects of GenAI are a strong reason for the individual university to adopt heavily restrictive policies, as its actual impact on climate change may be marginal at best.

Similarly, today GenAI systems rely on data centers that consume large volumes of fresh water for cooling, especially during peak computational operations. While these concerns are valid, they may not be as severe as they first appear. The environmental impact of water use varies by region; in countries with ample renewable energy and water resources, such as Sweden or France, the burden is considerably lower than in water-scarce or fossil-dependent regions. Major cloud providers are also investing in water-efficient cooling systems, water recycling, and relocating operations to more sustainable environments.[36] Similarly, the material demands of GenAI are not unique, as much of the hardware is shared across digital infrastructures. Mitigation efforts – including recycling programs, ethical sourcing, and low-impact chip design – are already underway to reduce these effects.[37] Taken together, these considerations suggest that while the environmental impacts of GenAI merit attention, they do not on their own justify highly restrictive university policies. The marginal contribution of student use – particularly via centralized, pre-trained models – is likely negligible in the context of global emissions and resource consumption. Given continuing improvements in efficiency and sustainability, the environmental argument alone appears insufficient to support broad limitations on GenAI regarding the individual university and use of GenAI in higher education.

We should also remember that the alternatives are not only to either be lenient regarding student use of GenAI or to be very restrictive full stop. Instead, there is a wide array of possible alternatives at hand where GenAI use is not highly restricted. For example, another possible approach to mitigating the environmental impact of GenAI is through carbon compensation or carbon offsetting initiatives. Universities that adopt GenAI tools could commit to compensating for the carbon emissions generated during both the training and use of these models. Carbon compensation involves investing in projects that reduce greenhouse gases, such as reforestation, renewable energy initiatives, or energy efficiency programs, to balance out the emissions produced by AI usage. By implementing carbon offset programs, universities can take responsibility for their environmental footprint without resorting to very restrictive policies that limit the adoption of beneficial technologies. This approach allows institutions to balance the advantages of GenAI in education with the need to address climate concerns, showing that meaningful environmental action can be taken even while using energy-intensive tools. Furthermore, carbon compensation could also serve as an educational opportunity for students, highlighting the importance of sustainability practices in the context of emerging technologies and helping to cultivate awareness of the broader impact of their choices. Thus, rather than heavily restrict GenAI, universities can use carbon compensation as a proactive solution to minimize the negative environmental effects while embracing the potential benefits of AI in higher education. Similarly, universities could demand that their models be based in countries and climates and utilize cloud services that are efficient and thereby push companies in the direction of 'green AI'.[38]

Of course, it could be argued that this reasoning is simply utilizing a faulty sense of moral mathematics, as discussed by Parfit, among others.[39] Parfit criticizes the idea that small, seemingly imperceptible contributions to a larger harm, such as climate change, should be dismissed as morally irrelevant. He introduces the concept of imperceptible harms, arguing that even if an individual agent's actions – such as using GenAI – make no noticeable difference by themselves, they still contribute incrementally to the collective harm. In Parfit's view, the moral weight of these contributions cannot be overlooked merely because their individual effects are small. He illustrates this with examples such as a group of torturers each contributing a fraction of the pain inflicted on a victim, where the collective result is undeniable harm despite the imperceptibility of any one torturer's contribution. Similarly, in the context of GenAI, even though a single university's energy consumption might *seem* negligible, when aggregated across all universities and institutions, the environmental impact could be *significant*. This reasoning suggests that we cannot dismiss the environmental harms associated with GenAI on the basis that individual contributions are small. Instead, we must acknowledge that collective responsibility plays a key role in addressing such issues.

Yet, I would contend, there is no such faulty reasoning when it comes to the moral mathematics in this context. It is simply not reasonable to think that the contribution from a single university, through its use of GenAI, would significantly affect the outcome of climate change. Even if we apply Sinnott-Armstrong's notion regarding the idea that a university's actions might contribute to the risk of climate change, this contribution would still be, in practical terms, effectively zero.[40] The marginal impact of one institution's energy consumption or emissions from GenAI, when considered in isolation, is too small to meaningfully influence the broader environmental outcome. Similarly, if we utilize any of the more recent views, such as John Broome's expected harm principle,[41] that states that even if an action does not always make a difference, it can still be morally significant if it raises the probability of harm, the results are basically the same. Even though this probabilistic reasoning avoids the problem of causal impotence that plagues many difference-making arguments, its application to heavy GenAI restrictions is not straightforward, and there are several reasons why it fails to justify very restrictive policies in this case. One issue with using expected harm to justify heavy GenAI restrictions is that it assumes the probability of harm is sufficiently high to warrant institutional intervention. Unlike climate change, where the link between carbon emissions and global warming is well established, the relationship between GenAI use in higher education and increased environmental harm is less clear. Universities primarily rely on pre-trained GenAI models, such as OpenAI's GPT or Google's Gemini, meaning their use does not directly contribute to the energy demands of AI training in a significant way. A similar claim can be regarding minerals for producing the hardware necessary for GenAI: this will not be affected by what the students do at a specific university. The case might be different, however, when it comes to water consumption and GenAI use. Here, use by individual students might make a difference. But then the argument for e.g. strict policies against GenAI use is reliant on students using these tools less under these conditions, which might not be true.[42]

Finally, while the environmental costs of GenAI – such as energy consumption, water use, and mineral extraction – are well documented, emerging research highlights its potential contributions to sustainability in specific applications, suggesting it may even yield a positive net impact in the long term. For instance, Yorke and colleagues argue that 'with the development of protein databases and more powerful model architectures such as transformers, large language models and diffusion models, end-to-end protein structure and function prediction from primary sequences has become realistic and has taken a leading role in the evolution of the field'.[43] If applied effectively, this could support the development of more energy-efficient materials produced through sustainable processes. Furthermore, even small AI-driven efficiency gains can translate into significant energy savings across industries, including mobile networks, where AI is expected to reduce power use by 10–15%.[44] In architecture, generative design has been shown to reduce energy consumption by 23% and cooling needs by 28%, while also lowering embedded carbon emissions.[45] In climate modeling, diffusion-based generative models developed at the University of California San Diego can simulate century-scale climate scenarios 25 times faster than conventional methods, enabling more granular emissions sensitivity analyses. GenAI is also advancing conservation and biodiversity monitoring. DeepMind's DGMR model, for example, improves short-term rainfall prediction and was preferred by meteorologists in 89% of comparisons,[46] while generative image augmentation techniques have boosted species recognition accuracy for rare wildlife to over 92%.[47] In battery recycling, generative learning is used to estimate the state of health (SOH) of retired batteries, reducing data requirements and potentially avoiding 35.8 billion kg of CO_2 emissions globally by 2030.[48] More broadly, GenAI can enhance ecological research by augmenting data-scarce datasets and extending ecological observations,[49] and support ESG (Environmental, Social, and Governance) performance through improved supply chain collaboration in small and medium-sized tourism enterprises.[50] Additionally, GenAI adoption is shown to increase sustainability-oriented entrepreneurial intentions by aligning psychological motivators such as feasibility and desirability.[51] These findings *do not* negate the environmental costs of GenAI, but they suggest that under the right conditions, these tools may also serve as catalysts for sustainability, particularly when deployed strategically and in contexts where their efficiency gains outweigh their resource demands.

When it comes to the question of making a difference, it is important to note that the production of GenAI models would likely proceed *regardless* of whether individual universities choose to adopt them. The development and deployment of AI systems are driven by powerful market and industry forces, meaning that a single university's decision to use – or not use – these models would have little impact on the overall production processes or the associated exploitation. In this sense, the connection between university adoption of GenAI and moral complicity is weakened, as institutional use is unlikely to meaningfully influence the labor practices involved in AI development. However, if universities and other institutions were to coordinate their efforts, they might exert meaningful pressure on AI companies to improve ethical labor practices – perhaps more easily than in industries involving more complex hardware production. These companies often already maintain formal ethical policies and recognize that universities can be valuable clients, sources of future talent, and institutions guided by normative commitments distinct from those of private corporations. This positioning may enable universities to influence corporate practices in ways that many other actors cannot. By pooling resources and advocating for ethical standards, universities could contribute to improvements in both the sustainability and fairness of GenAI technologies. Thus, while universities may not have strong grounds for prohibiting the use of GenAI outright, they plausibly do have a moral responsibility to work against exploitation. The same line of reasoning applies to environmental concerns such as climate change.[52] In both cases, while individual adoption may not have a direct negative impact, collective institutional action may offer a meaningful opportunity to drive positive change. Something similar can be argued when it comes to all other-regarding moral arguments.

The final issue we will discuss is student integrity and learning. A common argument for banning GenAI in higher education centers on protecting student privacy.[53] Privacy, in this context, refers to a student's control over how their data is collected, used, and shared. Understanding the implications of sharing data is challenging, especially since today's powerful AI systems can use even small data points to predict actions, influence behavior, or uncover hidden details, such as undiagnosed diseases. Given that GenAI relies on extensive datasets, it is reasonable to expect companies to continue collecting data. The leading GenAI tools are produced by commercial entities whose interests may conflict with user privacy when innovation and market expansion take priority.[54] Even though this is starting to change, high-quality AI tools in general require large non-synthetic datasets, and student-generated data, being non-synthetic and high-quality, provides companies with a strong incentive to use it. Therefore, it is crucial to secure and manage the data fed into these systems carefully to prevent misuse. Without stringent guidelines and strong data protection measures, mishandling this data could lead to major privacy breaches, risking unauthorized sharing of students' information.[55]

The concern is that relying on AI tools may undermine students' privacy. However, an outright ban or heavy restrictions on GenAI may not be the best solution to this issue. Instead of prohibiting its use, universities could provide students with access to licensed GenAI tools that have been vetted for privacy protections, ensuring that the tools they use align with principles of transparency and data security. This approach would allow students to use AI responsibly while safeguarding their personal information. By establishing clear privacy guidelines and offering secure GenAI tools, universities can protect student privacy without imposing prohibitive measures. Additionally, protecting student privacy could involve educating students about the risks and responsibilities associated with using GenAI. Teaching students how to use AI tools in ways that respect their own privacy, and that of others, would help them develop a critical understanding of these technologies, not using them haphazardly, and so on. Rather than avoiding GenAI, universities could guide students in how to engage with it responsibly, equipping them with the skills to maintain their privacy while benefiting from the technology. This approach is plausibly more likely to foster proactive engagement with GenAI, preparing students for a world where these tools are increasingly prevalent, without compromising their personal privacy or ethical standards. Thus, instead of banning GenAI in the name of privacy protection, universities can promote responsible AI use, ensuring that students are both informed and empowered to navigate the ethical challenges posed by these technologies.

In support of this, studies have found that when GenAI is banned or heavily restricted, people are more likely to use it covertly, which has been seen in the industry,[56] which undermines the very notion of protecting student integrity. If GenAI is prohibited in higher education, students – particularly the most vulnerable ones – may be inclined to use these tools in secret, without the guidance or ethical oversight that could ensure their proper use. Vulnerable students, such as those with less funding, limited time, or weaker academic backgrounds, are more prone to *using AI covertly*.[57] These students are also less likely to be able to afford the kind of licensed tools that provide better privacy and data protection, leaving them at greater risk of compromising their personal data or relying on untrustworthy sources. By creating a climate in which students feel they must hide their use of GenAI, universities may inadvertently push them into unethical behavior and expose them to additional risks. Instead of fostering honesty and integrity, a ban could lead to increased academic misconduct and make it more difficult for students to receive meaningful support for their use of generative tools. When students use GenAI covertly, they miss the opportunity to learn how to integrate these technologies responsibly and ethically, ultimately doing a disservice to their education and personal development. This is also why it is unclear what the relationship is between water consumption and policies regarding GenAI use. It is possible that banning or heavily restricting the use of GenAI makes the students use it less, and hence this has a positive impact on water consumption. Nonetheless, it is possible that they use the tools as much as they would otherwise and that water consumption is thereby unaffected by the policy at hand.

A similar line of reasoning applies to student learning more broadly. Just as universities must consider their role in systemic harms when adopting AI policies, they must also recognize their responsibility in shaping how students develop critical academic and ethical skills.[58] If institutions ban or severely restrict the use of GenAI, students may still engage with these tools outside formal education settings, but without the necessary guidance, reflection, and ethical considerations that responsible AI use requires.[59] If universities fail to integrate AI into their pedagogical frameworks, they may contribute to a larger systemic issue where students enter professional and academic environments without a critical understanding of AI's implications. Rather than prohibiting AI in an attempt to preserve traditional learning methods, universities should ensure that students learn how to use these tools thoughtfully, assess their limitations, and develop AI literacy.[60] If institutions take a restrictive rather than an educational approach, they risk reinforcing a hidden curriculum where students engage with AI in ways that are unregulated, inconsistent, and potentially unethical. Just as ethical responsibility in AI adoption cannot be reduced to banning tools without considering alternatives, student learning should not be framed as a choice between restriction and misuse, but rather as an opportunity to teach responsible AI engagement in a structured academic setting. This in combination with aligning examinations with the reality of the utilization of AI should lead to better outcomes. For example, if a learning objective is that the students are supposed to be able to write a coherent and well-written text or a code, then this should be tested under proctored circumstances, just to mention one example.

There might also be *unintended side effects* from using such technologies and these are also relevant from an ethical perspective where difference-making is key to what is the morally right thing to do. While direct contributions to environmental and ethical harm have been central arguments, it is important to consider how using AI tools might indirectly influence behavior within the broader university community and society. These side effects can manifest in various forms. For instance, if a single university allows its students to use AI with minimal or no restrictions, this could lead other universities to adopt a similar stance to remain competitive. As students who graduate from these programs enter the workforce, they may carry with them an uncritical, positive attitude toward GenAI. This, in turn, could lead to the private and public sectors becoming dominated by individuals who do not fully appreciate the ethical implications of GenAI use, potentially resulting in the normalization of practices that might have detrimental social and environmental impacts.

This concept is similar to the 'snowball effect' described by for instance Regan and Nefsky,[61] among others, where small actions – seemingly insignificant in isolation – can, when accumulated, lead to significant collective outcomes. The argument here is that the unrestricted use of AI within one university can set off a chain reaction, where the normalization of AI use becomes morally significant due to its broader consequences. Thus, the key argument against the use of GenAI in higher education extends beyond direct effects to include these indirect but morally relevant side effects. When evaluating whether a university should heavily restrict or ban AI usage, it is crucial to assess the broader impact of its normalization and the environmental and ethical footprint it collectively leaves behind. The university's role, in this context, is not solely about preventing direct harm, but also about setting an ethical precedent and preventing the normalization of behavior that, in the aggregate, contributes to greater societal harms. This approach reflects perhaps a more comprehensive understanding of moral responsibility – one that takes into account both direct consequences and the indirect influences that a university may have on broader social and environmental wellbeing.[62]

However, the problem with this argument is that it rests on several assumptions that may not necessarily hold. First, it assumes that there are clear negative side effects in terms of how other universities respond to one institution's policies. It presupposes that the actions of a single university would influence others in such a way that they too would adopt lax restrictions on AI usage. This might be the case with large and well-renowned universities such as MIT or Harvard, but hardly those universities that are less prestigious. This means that some universities might have more stringent responsibilities than others, but it is still unclear what these are. Second, the argument assumes that students will carry uncritical attitudes toward AI into their professional lives, neglecting the possibility that universities could provide critical education and foster ethical awareness in their students. Finally, it assumes that universities will not actively work to mitigate these side effects by promoting sustainability, improving AI's environmental footprint, or advocating against exploitative practices within the AI industry. If universities focus on embedding ethical considerations and sustainability into both the education they provide and their own AI-related activities, they could help ensure that students graduate with a more critical and responsible attitude toward AI, mitigating the potential for these unintended side effects to manifest.[63]

3.2. *The Non-Difference-Making Arguments*

Even if individual universities have little direct impact, participating in morally dubious practices – such as harm or exploitation – may still be wrong. The Complicity Principle[64] and similar arguments by Singer and Regan suggest that contributing to harmful outcomes can be morally problematic, even without direct causation.[65] Julia Nefsky expands on this by arguing that individuals and institutions share responsibility for systemic harms – such as climate change, labor exploitation, and digital privacy violations – since these harms result from the cumulative effects of many small contributions.[66] A university allowing unrestricted AI use may thus be complicit in reinforcing harmful structures. Similarly, group agency theorists argue that institutions, not just individuals, bear moral responsibility when their policies shape social norms and industry standards.[67] If universities collectively normalize uncritical AI adoption, they may contribute to an AI-driven future that disregards ethical concerns about privacy, labor rights, and environmental impact. Finally, Kantian-inspired arguments, such as the generalization test,[68] hold that a moral principle should be universalizable, meaning that if every university excused AI adoption on the basis of causal impotence, ethical considerations surrounding AI, environmental responsibility, and digital privacy would become meaningless.

Even though there is truth to these arguments, they are perhaps not as straightforward as they may initially seem. First, the Complicity Principle assumes that all AI systems inherently contribute to exploitation or environmental harm, but this is an oversimplification. While it is true that some producers of frontier models today – such as OpenAI and Meta – engage in problematic practices, this is not necessarily intrinsic to AI development. These companies, like others, can adopt fair labor standards and pursue ambitious environmental goals. Given that model training and deployment do not require large numbers of employees, implementing decent working conditions may even be more feasible for AI firms than for many other industries. Moreover, many AI developers are already working toward more ethical and transparent supply chains, with some actively seeking to reduce exploitation by adhering to fair labor practices and responsibly sourcing materials.[69] Since many of these companies have already established ethical frameworks and sustainability goals, they can be held accountable when they violate their own standards – potentially making it easier to push for improvements in how they operate. There is also a growing trend, particularly as open-source foundation models such as DeepSeek, Mistral, and Llama improve, for developers to build new tools based on these models that are more suitable for academic use. Even if these foundational models were produced under questionable labor or environmental conditions, tools derived from them are further removed from the original wrongdoing. According to some theories of moral complicity, this distance may reduce the extent of moral taint. Therefore, the moral concerns attributed to AI are not inherent to the technology itself but contingent upon the practices of specific companies. Similarly, as mentioned earlier, it is not entirely clear whether GenAI will exacerbate climate change or help mitigate it. As such, it remains uncertain to what extent the use of these tools contributes to either exploitation or environmental harm. Consequently, it is at least not evident that allowing more permissive policies regarding student use of GenAI is morally wrong.

Second, many of these views emphasize proportional responsibility, arguing that institutions should address ethical concerns in ways that align with their core mission.[70] In the case of universities, where the primary function is education, the integration of GenAI should be evaluated in terms of its potential to enhance learning while mitigating associated risks. If universities can implement AI in ways that uphold ethical standards – such as ensuring student privacy, especially for their vulnerable students, minimizing environmental impact, and advocating for fair labor practices – then they may have a *greater obligation* to *engage with AI responsibly* rather than prohibit its use outright. This perspective suggests that rather than banning GenAI, universities should adopt constructive approaches that incorporate AI ethically and effectively. One way to fulfill this responsibility is by advocating for fair labor practices in AI development, rather than simply rejecting the technology. Universities may have the capacity to influence industry standards by supporting AI tools developed under ethical conditions, encouraging companies to adopt transparent, fair employment practices in data annotation and content moderation. Similarly, universities can in many ways implement AI in controlled, ethical ways, such as by providing students with access to vetted AI tools that provide strong privacy protections, thereby reducing the risks associated with unregulated AI use. Additionally, instead of imposing outright prohibitions, institutions can develop ethical AI policies that guide student engagement, ensuring that students learn how to use AI critically and responsibly rather than avoiding it entirely.

These alternatives suggest that, for example, the group agency argument does not uniquely support heavily restrictive AI policies – it can equally justify an approach focused on ethical engagement with AI rather than outright bans. If universities act as moral agents in shaping responsible AI use, they contribute to a broader ethical framework that ensures AI serves educational objectives without reinforcing systemic harm. Thus, proportional responsibility does not necessarily require prohibition but instead calls for thoughtful integration, oversight, and advocacy to align AI use with institutional values. However, it should be noted that these are demanding expectations, and many universities today are likely to fall short of meeting them – particularly in areas such as license procurement and advocacy for socially and environmentally responsible AI practices. Given that high-tech AI companies are highly dependent on attracting talent from universities – and that their success partly depends on how they are perceived by students – universities may be especially well positioned to push for higher standards in social and environmental sustainability. In fact, this unique position may imply that universities bear even more stringent responsibilities than other actors when it comes to demanding and modeling ethical AI practices.

Third, the universalization argument suggests that universities should avoid participation in any ethically questionable system. Yet the moral implications of not imposing very restrictive regulations on GenAI use depend largely on how one interprets the nature of the process involved, which is itself influenced by the underlying intentions (or maxims) and corresponding actions. For instance, universities can educate students on responsible use of these tools while being transparent about their efforts to purchase licenses for technologies produced under fair working conditions and with minimal environmental impact. Additionally, universities can demonstrate their commitment to improving AI practices by advocating for better employment standards within the AI industry and supporting the development of less resource-intensive algorithms. This approach not only provides students with valuable skills but also reflects the university's dedication to seeking ethical alternatives within the AI industry. Such an approach represents a significant departure from the original argument that using AI inherently entails complicity in unethical practices. Rather than passively endorsing harmful practices, the university takes an active role in balancing the educational needs of its students with a commitment to ethical progress. By collaborating with other institutions and pushing for advancements in AI development and usage, universities can avoid moral compromise. Instead, they become agents of responsible AI use and drivers of positive change in the industry, aligning their actions with both practical and ethical objectives. This argument generalizes to the other versions of the non-difference-making argument as well.

Furthermore, even if we assume that universities can make a difference and that participating in these processes is morally wrong and blameworthy, thus giving them strong – perhaps even decisive – moral reasons to avoid involvement, it does not necessarily follow that they *only* have a decisive moral reason to strictly limit the use of GenAI tools at their institutions. Nor does it mean that this strong moral reason is, in fact, decisive. Instead, there may be situations where a university faces a genuine moral dilemma: it might have a decisive moral reason to be heavily restrictive of these tools while simultaneously having compelling reasons not to be so restrictive. Alternatively, the notion that the justification above necessarily leads to a very restrictive stance could itself serve as a *reductio* against that justification. This is true when it comes to the difference-making arguments as well as the non-difference-making arguments, even though I am only discussing it for non-difference-making arguments here.

Beginning with the potential dilemma for universities: on the one hand, they may have a moral duty to avoid contributing to collective harms such as environmental degradation. On the other hand, banning the use of GenAI might limit students' educational opportunities, placing them at a disadvantage compared to their peers at institutions that embrace AI tools, and harming the most vulnerable student groups. This creates a scenario where the university risks either compromising its moral responsibility to the environment or undermining its educational mission and the care for their students. Assuming that faculty members have enough resources to learn the tools, the decision depends on what you believe about the strength of the arguments above. It also depends on the strength of the positive arguments for using AI. However, if you doubt the strength of the negative arguments but not the positive ones, you might find that these are equally strong, creating a genuine dilemma.

In addition, if we accept premise 1 and 2 in the original argument described in Section 2 – that the university has a strong moral reason to adopt a heavily restrictive policy regarding student use of GenAI due to ethical concerns such as environmental harm or exploitation and when it has such a reason it needs to act on it through heavily restrictive measures – then we must be prepared to extend this reasoning to other areas of university operations which we alluded to above. For instance, many of the materials used to manufacture computers, servers, and other essential technology in higher education are sourced in ways that involve environmental degradation, unethical labor practices, and exploitation.[71] If we apply the same strict moral reasoning consistently, then the university would also have a strong moral reason to ban the use of these technologies. Yet, given how integral computers and technology are to the functioning of universities – from research to teaching – this would severely impede their operations. This suggests that the logic behind premises 1 and 2, if fully followed, leads to a highly 'impractical' outcome: the university would be unable to engage in many basic functions necessary for modern education, which of course is a *reductio ad absurdum* against this view.[72]

So, while the arguments above strengthen the idea that universities should consider their role in systemic harm, and so on, even when their individual actions do not have an impact, this does not establish a strong reason to ban or heavily restrict GenAI use. Instead, it suggests that universities should adopt ethical AI policies rather than bans, engage in systemic solutions such as advocating for fair AI development, supporting green AI initiatives, and integrating AI in ways that benefit students without exacerbating harm. Additionally, universities must recognize that their obligations apply across multiple domains, not just GenAI use, making targeted bans inconsistent and impractical. Thus, while these ideas challenge the notion that universities have no responsibility, it ultimately fails to justify heavy GenAI restrictions as the best moral response. Instead, a more effective and balanced approach could probably involve responsible engagement with AI that aligns with universities' broader ethical commitments while preserving academic innovation and accessibility.

4. Conclusion

The moral arguments for adopting bans or heavy restrictions on GenAI use in higher education face significant challenges when examined in depth. Many of the concerns presented – such as environmental harm, labor exploitation, and threats to privacy – are undoubtedly important. However, the analysis presented here suggests that these concerns do not necessarily provide a strong reason for banning or heavily restricting the use of GenAI. On closer examination, the arguments against GenAI often lead to *reductio* outcomes and are inconsistent when applied broadly across other technologies and university practices. Additionally, the notion that moral reasons should automatically lead to extremely restrictive policies ignores the complexity of practical decision-making within educational institutions, where competing responsibilities – such as fostering innovation, ensuring equitable access, and supporting student learning – must also be considered.

Instead of outright bans, heavy restrictions, and such, universities should probably focus instead on fostering responsible engagement with GenAI. Providing students with ethical guidelines, access to licensed AI tools, and opportunities to learn how to use these technologies thoughtfully and not unnecessarily, can address many of the concerns raised without sacrificing the educational benefits that AI can bring. Furthermore, by working collaboratively to address ethical issues – such as supporting carbon offsetting for energy consumption and advocating for fair labor practices – universities can engage meaningfully with the challenges associated with GenAI without resorting to extreme measures that may hinder educational progress. Ultimately, a balanced, context-sensitive approach that recognizes both moral and non-moral considerations is likely to be more effective in addressing the complex ethical landscape of GenAI in higher education.

Karl de Fine Licht, Department of Technology, Management and Economics, Chalmers University of Technology, Gothenburg, Sweden. karl.definelicht@chalmers.se

Acknowledgments

I would like to thank the reviewers for all the effort they have put into reading my paper and for their valuable suggestions on how to improve it.

NOTES

1 Zawacki-Richter *et al.*, "Systematic Review"; Ounejjar *et al.*, "SmartBlendEd."
2 Domingo-Alejo, "AI."
3 Bañeres *et al.*, "Early Warning System"; Pan *et al.*, "Gathering"; Embarak and Hawarna, "Enhancing."
4 de Fine Licht, "Generative"; Hagendorff, "Mapping."
5 Mollick and Mollick, "New Modes"; Mollick and Mollick, "Using AI."
6 Ericsson, "Deliberate Practice"; Gulseren *et al.*, "What."
7 Ferdman, "Human Flourishing"; Paglieri, "Expropriated Minds"; Williams, *Moral Luck*.
8 De Fine Licht, "Generative"; Hagendorff, "Mapping."
9 De Fine Licht, "Generative."
10 See e.g. Rudolph *et al.*, "ChatGPT."
11 Chan and Hu, "Students' Voices"; Schei *et al.*, "Perceptions"; Stöhr *et al.*, "Perceptions."
12 Cassinadri, "ChatGPT."
13 Paglieri, "Expropriated Minds."
14 De Fine Licht, "Generative."
15 Ibid.
16 McDonald *et al.*, "Generative."
17 Aylsworth and Castro, "Should I Use ChatGPT?"
18 Cf. de Fine Licht, "Generative."
19 Goetze, "AI Art."
20 See e.g. Crawford, *Atlas*; Kneese, "Climate Justice"; de Fine Licht, "Generative"; Driessens and Pischetola, "Danish"; Regilme, "Artificial Intelligence"; Sahebi and Formosa, "Artificial Intelligence."
21 E.g. Munn, "Digital Labor."
22 Regilme, "Artificial Intelligence."
23 Bender *et al.*, "On the Dangers"; Kneese and Young, "Carbon Emissions."
24 Minde, "GenAI."
25 Kirkpatrick, "Carbon Footprint."
26 See e.g. Veliz, *Privacy*; Huang, "Ethics"; de Fine Licht, "Generative"; Shukla and Taneja, "Ethical Considerations."
27 De Fine Licht, "Generative."
28 Cf. ibid.
29 Williams, *Moral Luck*; Raz, *Engaging Reason*.
30 Bender *et al.*, "On the Dangers"; Kneese and Young, "Carbon Emissions."
31 De Fine Licht, "Generative."
32 Castro, 2024. "Rethinking."
33 See e.g. Alzoubi and Mishra, "Green."
34 Meta, "Introducing."
35 Castro, 2024. "Rethinking."
36 See Alzoubi and Mishra, "Green," for an overview.

37 Ibid.
38 Ibid; Stein, "GenAI."
39 Parfit, 1984; see e.g. Sinnott-Armstrong, "It's Not My Fault"; Kagan, "Do I Make a Difference?"; Broome, *Climate Matters*; Nefsky, "Consequentialism"; Nefsky, "How You Can Help."
40 Sinnott-Armstrong, "It's Not My Fault."
41 Broome, *Climate Matters*.
42 Mollick, "Co-intelligence."
43 Yorke *et al.*, "Design."
44 Thormundsson, "Power."
45 Suphavarophas *et al.*, "Performance-Based."
46 Ravuri *et al.*, "Skillful precipitation."
47 Zhang *et al.*, "Few-Shot."
48 Tao *et al.*, "Generative."
49 Rafiq *et al.*, "GenAI."
50 Wang and Zhang, "Promoting."
51 Duong, "How."
52 cf. Chen *et al.*, "Survey"; Stein, "GenAI."
53 Williams, *Moral Luck*; Ghimire and Edwards, "From Guidelines."
54 Véliz, *Privacy*.
55 See e.g. Bhutoria, "Personalized."
56 Mollick, "Co-intelligence."
57 E.g. Li *et al.*, "Ethical."
58 E.g. Schwartz, "Science."
59 See e.g. Dell'Acqua *et al.*, "Navigating"; Mollick and Mollick, "Using AI"; Chiu, "Future."
60 E.g. Kasneci *et al.*, "ChatGPT."
61 Regan, *Utilitarianism*; Nefsky, "Consequentialism"; Nefsky, "How You Can Help."
62 See also e.g. Sunstein, "On the Expressive Function"; Shell, "Limits," on symbolic and expressive harms.
63 See e.g. Ericsson, "Deliberate Practice"; Gulseren *et al.*, "What," for studies examining how students trained in ethics in higher education act in professional settings after graduation.
64 Kutz, "Complicity."
65 Singer, "Act-Utilitarianism"; Regan, *Utilitarianism*.
66 Nefsky, "Consequentialism"; Nefsky, "How You Can Help."
67 May, *Sharing*; Pettit, "Responsibility"; Pettit, *Robust Demands*; List and Pettit, *Group Agency*.
68 Glover, "It Makes No Difference"; O'Neill, *Acting*.
69 Alzoubi and Mishra, "Green."
70 E.g. Pettit, "Groups"; see also de Fine Licht, "Generative."
71 Hampton *et al.*, "From Mining."
72 There are other moral arguments that have to do with the digital divide (see de Fine Licht, "Generative"). These can be dealt with in the same way as the arguments listed in this article.

References

Alzoubi, Y. I., and A. Mishra. "Green Artificial Intelligence Initiatives: Potentials and Challenges." *Journal of Cleaner Production* 468 (2024): 143090. https://doi.org/10.1016/j.jclepro.2024.143090.

Aylsworth, T., and C. Castro. "Should I use ChatGPT to Write My Papers?" *Philosophy and Technology* 37 (2024): 117. https://doi.org/10.1007/s13347-024-00809-w.

Bañeres, D., M. E. Rodríguez, A. E. Guerrero-Roldán, and A. Karadeniz. "An Early Warning System to Detect At-Risk Students in Online Higher Education." *Applied Sciences* 10, no. 13 (2020): 4427. https://doi.org/10.3390/app10134427.

Bender, E. M., T. Gebru, A. McMillan-Major, and S. Shmitchell. "On the Dangers of Stochastic Parrots: Can Language Models be too Big?" In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623. New York: Association for Computing Machinery, 2021.

Bhutoria, A. "Personalized Education and Artificial Intelligence in the United States, China, and India: A Systematic Review Using a Human-in-the-Loop Model." *Computers and Education: Artificial Intelligence* 3 (2022): 100068.

Broome, J. *Climate Matters: Ethics in a Warming World*. New York: W.W. Norton & Company, 2012.

Cassinadri, G. "ChatGPT and the Technology-Education Tension: Applying Contextual Virtue Epistemology to a Cognitive Artifact." *Philosophy & Technology* 37, no. 1 (2024): 14.

Castro, D. *Rethinking Concerns About AI's Energy Use*. Center for Data Innovation, 2024.

Chan, C. K. Y., and W. Hu. "Students' Voices on GenAI: Perceptions, Benefits, and Challenges in Higher Education." *International Journal of Educational Technology in Higher Education* 20, no. 1 (2023). https://doi.org/10.1186/s41239-023-00411-8.

Chen, Z., M. Wu, A. Chan, X. Li, and Y. S. Ong. "Survey on AI Sustainability: Emerging Trends on Learning Algorithms and Research Challenges." *IEEE Computational Intelligence Magazine* 18, no. 2 (2023): 60–77.

Chiu, T. K. "Future Research Recommendations for Transforming Higher Education with GenAI." *Computers and Education: Artificial Intelligence* 6 (2024): 100197.

Crawford, K. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press, 2021.

Dell'Acqua, F., E. McFowland, III, E. R. Mollick, H. Lifshitz-Assaf, K. C. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani. *Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality* (Working Paper No. 24-013). Harvard Business School, Technology & Operations Management Unit, 2023. https://doi.org/10.2139/ssrn.4573321.

de Fine Licht, K. "Generative Artificial Intelligence in Higher Education: Why the 'Banning Approach' to Student Use is Sometimes Morally Justified." *Philosophy and Technology* 37, no. 3 (2024): 113.

Domingo-Alejo, J. "AI Integrated Administration Tool Design with ML Technology for Smart Education System." In *2024 4th International Conference on Advance Computing and Innovative Technologies in Engineering (ICACITE)*, 1423–28. Piscataway: IEEE, 2024.

Driessens, O., and M. Pischetola. "Danish University Policies on GenAI: Problems, Assumptions, and Sustainability Blind Spots." *MedieKultur* 40, no. 76 (2024): 31–52. https://doi.org/10.7146/mk.v40i76.143595.

Duong, C. D. "How AI-Enabled Drivers Inspire Sustainability-Oriented Entrepreneurial Intentions: Unraveling the (In)congruent Effects of Perceived Desirability and Feasibility from the Entrepreneurial Event Model Perspective." *Sustainable Development* (2025). https://doi.org/10.1002/sd.3461.

Embarak, O. H., and S. Hawarna. "Enhancing Student Success with XAI-Powered RADAR: An Automated AI-Driven System for Early Detection of At-Risk Students." *Procedia Computer Science* (2025). https://doi.org/10.1016/j.procs.2023.12.187.

Ericsson, K. A. "Deliberate Practice and Acquisition of Expert Performance: A General Overview." *Academic Emergency Medicine* 15, no. 11 (2008): 988–994.

Ferdman, A. "Human Flourishing and Technology Affordances." *Philosophy and Technology* 37, no. 1 (2023): 1. https://doi.org/10.1007/s13347-023-00686-9.

Ghimire, A., and J. Edwards. "From Guidelines to Governance: A Study of AI Policies in Education." In *25th International Conference on Artificial Intelligence in Education (AIED 2024), Communications in Computer and Information Science*, Vol 2151 299–307. Cham: Springer, 2024.

Goetze, T. S. "AI Art is Theft: Labour, Extraction, and Exploitation: Or, On the Dangers of Stochastic Pollocks." In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 186–196. New York: Association for Computing Machinery, 2024.

Gulseren, D., N. Turner, and J. M. Weinhardt. "What Makes Ethics Education Effective? An Umbrella Review and Evidence-Led Best Practices." *Journal of Business Ethics Education* 18, no. 1 (2021): 5–24.

Hagendorff, T. "Mapping the Ethics of GenAI: A Comprehensive Scoping Review." Preprint, arXiv:2402.08323 (2024)

Hampton, Lelia, Madeline Schlegel, Ellie Bultena, Jasmin Liu, Anastasia Dunca, Mrinalini Singha, et al. "From Mining to E-Waste: The Environmental and Climate Justice Implications of the Electronics Hardware Life Cycle." In *MIT Case Studies in Social and Ethical Responsibilities of Computing*. Hoboken, NJ, USA: John Wiley & Sons, Ltd., 2024.

Huang, L. "Ethics of Artificial Intelligence in Education: Student Privacy and Data Protection." *Science Insights Education Frontiers* 16, no. 2 (2023): 2577–87.

Glover, Jonathan. "It Makes no Difference Whether or Not I do It." *Proceedings of the Aristotelian Society* 49 (1975): 171–190.

Kagan, S. "Do I Make a Difference?" *Philosophy & Public Affairs* 39 (2011): 105–141.

Kasneci, E., K. Seßler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, et al. "ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education." *Learning and Individual Differences* 103 (2023): 102274.

Kirkpatrick, K. "The Carbon Footprint of Artificial Intelligence." Communications of the ACM 2023. https://cacm.acm.org/news/the-carbon-footprint-of-artificial-intelligence/.

Kneese, T. "Climate Justice & Labor Rights." SSRN Scholarly Paper No. 4533853). SSRN 2023.

Kneese, T., and M. Young. *Carbon Emissions in the Tailpipe of Generative AI*. Harvard Data Science Review, 2024. Special Issue 5. https://doi.org/10.1162/99608f92.fbdf6128.

Kutz, Christopher. *Complicity: Ethics and Law for a Collective Age*. New York: Cambridge University Press, 2000.

Li, Z., A. Dhruv, and V. Jain. "Ethical Considerations in the Use of AI for Higher Education: A Comprehensive Guide." In *Proceedings of the 18th IEEE International Conference on Semantic Computing (ICSC 2024)*, 218–223. Piscataway: IEEE, 2024.

List, C., and P. Pettit. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press, 2011.

May, L. *Sharing Responsibility*. Chicago: University of Chicago Press, 1992.

McDonald, N., A. Johri, A. Ali, and A. Hingle. "Generative Artificial Intelligence in Higher Education: Evidence from an Analysis of Institutional Policies and Guidelines." Preprint, arXiv:2402.01659 (2024)

Meta. "Introducing Meta Llama 3: The Most Capable Openly Available LLM to Date." April 18, 2024. https://ai.meta.com/blog/meta-llama-3/.

Minde, B. "GenAI Does Not Run on Thin Air! Blogpost." *RI.SE*, 2023. https://www.ri.se/en/news/blog/generative-ai-does-not-run-on-thin-air.

Mollick, E. R., and L. Mollick. "New Modes of Learning Enabled by AI Chatbots: Three Methods and Assignments." SSRN 2022. https://ssrn.com/abstract=4300783

Mollick, E. R., and L. Mollick. "Using AI to Implement Effective Teaching Strategies in Classrooms: Five Strategies, Including Prompts." SSRN 2023. https://ssrn.com/abstract=4391243

Mollick, E. *Co-Intelligence: Living and Working With AI*. Penguin, 2024.

Munn, L. "Digital Labor, Platforms, and AI." In *Introduction to Digital Humanism: A Textbook*, 557–569. Cham: Springer, 2023.

Nefsky, J. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy & Public Affairs* 39, no. 4 (2011): 364–395.

Nefsky, J. "How You Can Help, Without Making a Difference." *Philosophical Studies* 174 (2017): 2743–67.

O'Neill, O. *Acting on Principle: An Essay on Kantian Ethics*. Cambridge: Cambridge University Press, 2013.

Ounejjar, L. A., M. Lachgar, O. Ouhayou, M. D. Laanaoui, E. Refki, R. Makaoui, and A. Saoud. "SmartBlendEd: Enhancing Blended Learning through AI-Optimized Scheduling and User-Centric Design." *SoftwareX* 27 (2024): 101891. https://doi.org/10.1016/j.softx.2024.101891.

Parfit, D. *Reasons and Persons*. Oxford University Press, 1984.

Pan, L., Patterson, N., McKenzie, S., Rajasegarar, S., Wood-Bradley, G., Rough, J., Luo, W., Lanham, E., & Coldwell-Neilson, J. (2020). Gathering Intelligence on Student Information Behavior Using Data Mining. *Library Trends*, 68(4), 636–658. Johns Hopkins University Press. https://doi.org/10.1353/lib.2020.0015.

Paglieri, F. "Expropriated Minds: On Some Practical Problems of GenAI, Beyond Our Cognitive Illusions." *Philosophy and Technology* 37, no. 2 (2024): 55. https://doi.org/10.1007/s13347-024-00743-x.

Pettit, P. "Responsibility Incorporated." *Ethics* 117, no. 2 (2007): 171–201.

Pettit, P. "Groups with Minds of Their Own." *Social Epistemology: Essential Readings* 242 (2010): 167–193.

Pettit, P. *The Robust Demands of the Good: Ethics with Attachment, Virtue, and Respect*. Oxford: Oxford University Press, 2015.

Rafiq, K., S. Beery, M. S. Palmer, Z. Harchaoui, and B. Abrahms. "GenAI as a Tool to Accelerate the Field of Ecology." *Nature Ecology & Evolution* 9 (2025): 378–385. https://doi.org/10.1038/s41559-025-02345-9.

Ravuri, S., K. Lenc, M. Willson, D. Kangin, R. Lam, P. Mirowski, M. Fitzsimons, et al. "Skillful Precipitation Nowcasting Using Deep Generative Models of Radar." *Nature* 597, no. 7878 (2021): 672–77. https://doi.org/10.1038/s41586-021-03854-z.

Raz, Joseph. *Engaging Reason: On the Theory of Value and Action*. Oxford: Oxford University Press, 1999.

Regan, D. H. *Utilitarianism and Co-Operation*. Oxford: Oxford University Press, 1980.

Regilme, S. S. F. "Artificial Intelligence Colonialism: Environmental Damage, Labor Exploitation, and Human Rights Crises in the Global South." *SAIS Review of International Affairs* 44, no. 2 (2024): 75–92. https://doi.org/10.1353/sais.2024.a950958.

Rudolph, J., S. Tan, and S. Tan. "ChatGPT: Bullshit Spewer or the End of Traditional Assessments in Higher Education?" *Journal of Applied Learning and Teaching* 6, no. 1 (2023): 342–363.

Sahebi, S., and P. Formosa. "Artificial Intelligence (AI) and Global Justice." *Minds and Machines* 35, no. 1 (2024): 4.

Schei, O. M., A. Møgelvang, and K. Ludvigsen. "Perceptions and Use of AI Chatbots Among Students in Higher Education: A Scoping Review of Empirical Studies." *Education in Science* 14, no. 8 (2024): 922. https://doi.org/10.3390/educsci14080922.

Schwartz, B. "Science, Scholarship, and Intellectual Value: A Guide to What Higher Education Should be Like." *Journal of Moral Education* 51, no. 1 (2020): 61–72. https://doi.org/10.1080/03057240.2020.1772211.

Shell, S. M. "The Limits of Expression: Free Speech and the Moral Foundations of Liberty." *Harvard Law Review* 118, no. 6 (2004): 1923–78.

Shukla, R. P., and S. Taneja. "Ethical Considerations and Data Privacy in Artificial Intelligence." In *Integrating GenAI in Education to Achieve Sustainable Development Goals*, 86–97. New York: IGI Global, 2024.

Singer, P. "Is Act-Utilitarianism Self-Defeating?" *Philosophical Review* 81, no. 1 (1972): 94–104.

Sinnott-Armstrong, W. "It's Not My Fault: Global Warming and Individual Moral Obligations." In *Perspectives on Climate Change: Science, Economics, Politics, Ethics*, edited by W. Sinnott-Armstrong and R. B. Howarth, 285–307. Bingley: Emerald, 2005.

Stein, A. L. "GenAI and Sustainability." *University of Florida Levin College of Law Research Paper*. Forthcoming, 2024. SSRN. https://ssrn.com/abstract=4935208.

Stöhr, C., A. W. Ou, and H. Malmström. "Perceptions and Usage of AI Chatbots among Students in Higher Education across Genders, Academic Levels, and Fields of Study." *Computers and Education: Artificial Intelligence* 7 (2024): 1–12. https://doi.org/10.1016/j.caeai.2024.100259.

Sunstein, C. "On the Expressive Function of Law." *University of Pennsylvania Law Review* 144, no. 5 (1996): 2021–53.

Suphavarophas, N., T. Harputlugil, C. Iliescu, and M. P. Nielsen. "Performance-Based Generative Design Strategies for Sustainable Buildings: Reducing Energy and Cooling Demand through AI." *Energy and Buildings* 301 (2024): 113373. https://doi.org/10.1016/j.enbuild.2024.113373.

Tao, S., R. Ma, Z. Zhao, G. Ma, L. Su, H. Chang, and G. Zhou. "Generative Learning Assisted State-of-Health Estimation for Sustainable Battery Recycling with Random Retirement Conditions." *Nature Communications* 15 (2024): 10154. https://doi.org/10.1038/s41467-024-54454-0.

Thormundsson, B. "Power Consumption when Training Artificial Intelligence (AI) Based Large Language Models (LLMs) in 2023." Statista 2024. https://www.statista.com/statistics/1384401/energy-use-when-training-llm-models/.

Véliz, C. *Privacy Is Power: Why and How You Should Take Back Control of Your Data*. Hoboken: Melville House, 2021.

Wang, S., and H. Zhang. "Promoting Sustainable Development Goals through Generative Artificial Intelligence in the Digital Supply Chain: Insights from Chinese Tourism SMEs." *Sustainable Development* 33, no. 1 (2024): 1231–48.

Williams, B. *Moral Luck: Philosophical Papers 1973–1980*. Cambridge: Cambridge University Press, 1981.

Yorke, S. K., Z. Yang, E. G. Wiita, A. Kamada, T. P. Knowles, and M. J. Buehler. "Design and Sustainability of Polypeptide Material Systems." *Nature Reviews Materials* 10 (2025): 1–19. https://doi.org/10.1038/s41578-025-00793-3.

Zawacki-Richter, O., V. I. Marín, M. Bond, and F. Gouverneur. "Systematic Review of Research on Artificial Intelligence Applications in Higher Education – Where Are the Educators?" *International Journal of Educational Technology in Higher Education* 16, no. 1 (2019): 1–27.

Zhang, Q., X. Yi, J. Guo, Y. Tang, T. Feng, and R. Liu. "A Few-Shot Rare Wildlife Image Classification Method Based on Style Migration Data Augmentation." *Ecological Informatics* 75 (2023): 102237. https://doi.org/10.1016/j.ecoinf.2023.102237.