A Horizon Scan of AI and Higher Education: Fourteen Issues That May Transform the Sector

SJ Beard[1], Katherine Roehrick[2], James Beadle[3], Gareth A. Brinkworth[4], Martin Compton[5], Kate Crawford[6], Megan Ennion[4], Lance Gharavi[7], Claire Gordon[8], Velislava Hillman[9], David Madigan[10], Alexandru Marcoci[1], Alice E. Marwick[11], Alondra Nelson[12], Sebastián Otero[13], Matthew Pittinsky[14], Simon Walker[8], Dan Wang[15], Gabriela C. Zapata[16], and Matthew Connelly[2]

[1] Institute for Technology and Humanity, University of Cambridge

[2] Department of History, Columbia University

[3] Cambridge University Press & Assessment

[4] Faculty of Education, University of Cambridge

[5] Centre for Excellence in Learning and Teaching, University of East London

[6] School of Communication and Journalism, University of Southern California

[7] School of Music, Dance and Theatre, Arizona State University

[8] Eden Centre for Education Enhancement, London School of Economics and Political Science

[9] Goldsmiths, University of London

[10] Khoury College of Computer Sciences, Northeastern University

[11] Data & Society

[12] Institute for Advanced Study

[13] Department of Economics, Columbia University

[14] The Sanford School, Arizona State University

[15] Columbia Business School, Columbia University

[16] School of Education, University of Nottingham

Author Note: This paper represents a collaborative foresight exercise. Individual authors may not endorse all analyses or conclusions presented.

Abstract

AI, and especially Generative AI, are profoundly affecting higher education, challenging longstanding philosophical, pedagogical, cultural, and economic assumptions and reshaping its core practices. Much current research is trying to understand and assess these impacts; however, future trends and developments could be even more disruptive for both faculty and students. This paper provides a longer-term perspective in the form of a horizon scan, using the IDEA protocol with a group of 27 education and AI experts from the UK and USA to identify and evaluate future issues that may significantly impact the sector but have yet to reach prominence in discussions within the sector. It identifies 14 issues (from an initial longlist of 72) with the highest potential impact over both the near-term (such as data privacy, shadow curricula, and personalized education) and the longer-term (such as AI learning companions, the dematerialization of universities, and the need to prepare for existential risk). It also identifies important cross-cutting themes that arise from these issues: reflecting on the real value of education, understanding social aspects of learning, and the social and political context of AI deployment. The paper seeks to support educators in moving from a reactive approach to handling AI impacts towards a more proactive approach of planning for possible future social and technological transitions and providing students with the skills and resources needed to navigate them.

Keywords: Generative AI; Transformative AI; Higher Education; Horizon Scanning; Foresight.

A horizon scan of AI and Higher Education: Fourteen Issues That May Transform the Sector

Generative AI (GenAI) is transforming both the educational experience of students and the job market they are entering. A majority of senior academic leaders in the United States of America (USA) do not believe their institutions are adequately preparing students for the future or helping faculty use GenAI effectively (Association of American Universities and Colleges, 2025), while only 42% of students in the United Kingdom (UK) feel university staff are well-equipped to support them with GenAI (Freeman, 2025: 10). The literature on AI in education is growing rapidly, but much of this focuses on specific present challenges rather than AI's future potentialities. We urgently need to explore the opportunities and challenges that will shape higher education over the coming decades.

One scenario is that AI develops transformative capabilities, leading to irreversible change across most important aspects of life that radically challenge our assumptions about human progress and well-being (Gruetzemacher & Whittlestone, 2022). Among some experts, this has long been talked of as constituting a 'medium-term' for AI development trajectories (Baum, 2020; Cave & Ó hÉigeartaigh, 2019); however, many are now revising expectations due to the recent acceleration of progress in AI capabilities and applications. Even simply projecting current trends forward in a broadly linear fashion suggests that societies and economies could dramatically change. For instance, the largest survey of AI researchers to date suggested that by 2047, machines could feasibly surpass humans on dozens of tasks, from retail sales to writing best-selling fiction (Grace et al., 2025; see also Bengio et al., 2026).

Conversely, others argue that, while remarkable, the scale of progress that has been achieved so far is over-hyped (Eriksson et al., 2025), and the gap between current capabilities and genuinely transformative AI remains vast (Narayanan & Kapoor, 2025). In other words, we could be standing on the brink of a rapid and extreme social transformation, or we could be standing near the peak of an AI bubble and facing an imminent backlash against the technology and institutions that had uncritically embraced it.

Regardless of which scenario comes to pass, the job of universities and colleges is to prepare students and societies for the future, and thus, we must consider how we could adapt to different possibilities for the future development and deployment of AI, and what decisions need to be made now to manage the many risks and opportunities we might face going forward. This paper uses horizon scanning to explore and assess some of the issues these possibilities could pose for higher education. Horizon scanning is a structured expert elicitation process designed to combine knowledge and experience from diverse stakeholder communities to identify and evaluate future issues that have not yet achieved widespread prominence (Wintle et al., 2024). Our horizon scan included a series of anonymous surveys and participatory workshops with 27 experts on AI in higher education drawn from senior academic leadership, researchers, and the private sector across the UK and USA (listed in Appendix B). It was facilitated by Dr SJ Beard and Professor Matthew Connelly with the assistance of Dr Katherine Roehrick. After discussing and evaluating a longlist of 72 initial submissions, participants identified 14 emerging issues with the potential to substantially impact higher education across the near-term (<5 years time) and longer-term (>5 years). We also considered relationships between these issues, both in terms of potential causation between them and cross-cutting themes that emerge from them.

The next section sets out the background to our study. Section 2 summarizes our horizon scanning process. Section 3 describes the 14 highest priority issues we identified. Finally, section 4 reviews the connections between issues and suggests key lessons for those involved in higher education. A detailed explanation of our methodology and a list of participants are set out in appendices, with full data from each stage of the exercise provided in supplementary material.

1 Background

Responses from higher education to AI have so far been disjointed. Survey data shows enormous mistrust of AI companies among educators (Oxford University Press, 2023; Ruediger et al., 2023). Some academics feel the prospects for AI in higher education are over-hyped and that its use neglects the social and critical aspects of education (Sparrow & Flenady, 2025). Yet many are also experimenting with using GenAI to augment grading, monitoring, course delivery, discussion moderation (Bowen & Watson, 2024), formative assessment (Tzirides et al., 2024, 2025; Zapata et al., 2024, 2025), and even course design (Tran et al., 2024). Others propose combining generative AI with augmented reality to create new forms of immersive education (Dalziel et al., 2024; Kılıçkaya & Kic-Drgas, 2025). Universities have also started using AI to support, or even automate, a variety of administrative and leadership tasks, from resource allocation to policy development (Sposato, 2025).

One reason for disjunction is the dearth of evidence. At the time of writing, there have only been a few dozen experimental studies to measure the effect of AI on learning, showing mixed results. The vast majority had a sample size of fewer than 100 and lasted less than a single semester, making the longer-term effects on the full range of educational experiences highly uncertain (Jin & Seren, 2025). While some researchers have demonstrated specific ways in which AI can enhance learning, such as improving student motivation (Kestin, 2025) or formative feedback (Tzirides et al., 2025), others find that students using AI do not read as carefully (Fan et al., 2024) and write with diminished accuracy and originality (Niloy et al., 2024). Studies also show that students do not realize when AI is diminishing or enhancing their learning (Bastani et al., 2024).

Another factor hampering effective response to GenAI is the conservative nature of most institutions. University governance typically requires multi-year approval cycles for programme revisions, assessment changes, and policy updates. Regulatory frameworks, such as the UK Competition and Markets Authority rules obliging early publication of assessment details, further constrain rapid change. This poorly matches the pace at which AI is altering disciplinary content, professional standards, and graduate skill demands, or the pace of social, cultural, and economic change. Even if not for AI, colleges and universities would still need to innovate forms of education that depart from their disciplinary traditions and industrial models to advance human capabilities and societal wellbeing (Schleicher, 2018). While college graduates maintain an advantage over non-graduates in compensation and job stability once hired; the gap in unemployment rates has been narrowing since the 2008 Financial Crisis, and their relative advantage is at the lowest level since the 1970s (Cline & Kaymak, 2025). This is mainly because unemployed college graduates are no more likely to be hired than non-graduates. University has traditionally been seen as the primary pathway to a career in white-collar fields. If entry-level jobs disappear as AI automates tasks traditionally done by early-career staff, it might further erode the perceived value of degree programs (Brynjolfsson et al., 2025). While not an inevitable implication of GenAI adoption (Acemoglu et al., 2023), this could cause a serious drop-off in college enrollments.

Other challenges include demographic trends reducing college-age cohorts, funding cuts, an unsustainable rise in costs, and political attacks on academia's influence and prestige (Levine & Van Pelt, 2021). AI can be seen as compounding higher education's preexisting problems, but many believe it could be turned into an opportunity if institutions took a more proactive approach: moving from a performance-oriented mindset focused on grades and sanctions, towards a mastery-oriented one in which greater freedom to develop, explore, and apply new ideas enhances students' intrinsic desire to learn (Gallant & Rettinger, 2025; Kalaitzidis, 2025). AI seems to provide both needs and opportunities to redefine education towards the cultivation of human faculties like curiosity and creativity, which may be harder to automate, as well as making education more personalized and responsive to student needs and abilities (Aoum, 2017).

2 Summary of Horizon Scan

Responding to these challenges and opportunities requires a systematic approach to identifying, evaluating, and prioritizing among them. Horizon scanning offers one way of filtering diverse sources of information to seek weak signals that, when contextualized, indicate which issues have received insufficient attention (Wintle et al., 2024). This can highlight developments that are still in their early stages but might loom much larger in the future, while surfacing unspoken assumptions and unexpected connections that require closer scrutiny. Horizon scanning with the Delphi technique (Turoff & Linstone, 1975) has a long history, including several studies focused on developments in AI (Alon et al., 2025; Gokhale, 2025).

Our study used the Investigate-Discuss-Estimate-Aggregate (IDEA) protocol (Hemming et al. 2018). In our process (set out in Appendix A), participants were asked to investigate and submit candidate issues, privately and anonymously score a cross-section of these issues, and discuss their thinking with others. They then provided a second score for all issues, which was mathematically aggregated. This method shares features of the Delphi technique, in that the evaluation of issues is anonymous and iterative, and it draws on the collective wisdom of a diverse group while affording individuals the opportunity to give private judgments and revise them in light of information and reasoning provided by others. The main difference from standard Delphi studies is the presence of discussion, which can be a powerful epistemic tool as information sharing between participants improves the accuracy of their forecasts (Hanea et al., 2018). However, aside from seeking a shared understanding of terms and reducing linguistic ambiguity, consensus is not sought during discussion, and scores are kept anonymous throughout the exercise. This is done to avoid undesirable group dynamics and peer pressure distorting individual judgments. The IDEA protocol has performed well relative to prediction markets (Hanea et al., 2017) and has been successfully applied to a range of areas. We drew directly on experience from using this method to explore AI applications (Reynolds et al., 2025), conservation (Sutherland et al., 2025), and biotechnology (Kemp et al., 2020). The five stages of our horizon scan are shown in Figure 1.

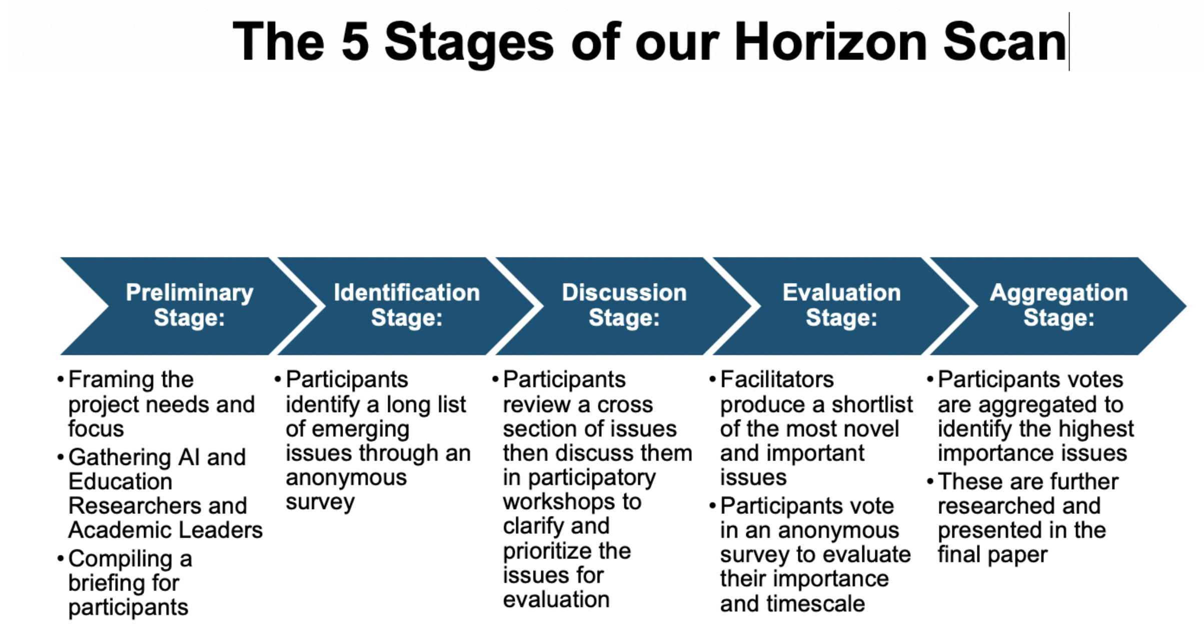

Figure 1: *Summary of our horizon scanning procedure*

Our methodology was thus optimized for surfacing neglected issues, especially those that will not fully emerge in the near-term. We specifically sought to remove more immediate concerns where there is already a broad consensus and high-level attention. Still, it is worth noting that, from the first anonymous survey, it was clear that all participants agreed that AI is already bringing about a profound transformation in higher education. To be sure, agreeing to commit significant time to discussing the subject introduces an unavoidable selection bias. But one might have expected at least one participant to take an opposing view, if only as a devil's advocate. Instead, a highly diverse group of participants, representing a broad range of policy experiences and different kinds of expertise, debated the issues detailed below, but not one disagreed with the idea that AI already has great and growing importance for the future of higher education.

Below, we provide a high-level summary of the top 14 issues (while acknowledging that the number of topics covered means that there has to be some sacrifice of depth for breadth). We hope to spur further research into these issues and discussion of their implications by researchers, policy-makers, and the wider public. To avoid giving a false sense of forecasting precision or overemphasizing minor differences in scoring, we present these issues not in order of their evaluated importance but the timescale over which participants expected them to emerge (Table 1).

Table 1: *The top 14 issues identified by our horizon scan*

Near-term (emerging in < 5 years)
Data privacy, reputational, and professional risks to students if educational ‘sandboxes’ are not provided
Lack of basic skill acquisition for students offloading cognitive tasks to AI
Resource-strapped universities and time-poor staff struggle to make the best use of AI, potentially triggering feedback loops of institutional decline
GenAI-supported learning pathways outside formal instruction, i.e., emergence of shadow curricula
Equity gaps in access to AI Literacy and customization
Personalized AI enables the kind of tutoring that has long been the holy grail of learning science
Declining human mastery of core academic practices
Longer-Term (emerging in > 5 years)
Emergence of AI learning companions as primary educational relationships
Developing programme-focused approaches to advance curriculum design
Persuasive design in AI tools as a scalable vector of influence over vulnerable learners
Transfer of Pedagogical Influence from Educators to AI Systems
Institutional competition with AI-native alternative education providers
The elimination of college and university as a material/spatial experience owing to the use of AI
Students need to develop the capabilities to help society mitigate existential risk, such as from AI and nuclear war

The full list of nominated issues and the results of each stage of the horizon scan are provided in supplementary material.

3 The Issues

3.1 Near-term (emerging in < 5 years)

3.1.1 Data privacy, reputational, and professional risks to students if educational ‘sandboxes’ are not provided

Much of the AI use by students is currently informal and via corporately owned, under-regulated models, creating significant risks for students. Without protected environments, students face data privacy violations as their personal information, academic work, and learning patterns get collected by commercial AI platforms with unclear data handling practices (Ismail & Aloshi, 2025). Reputational risks emerge when students' experimental AI interactions or mistakes become permanently recorded or discoverable, potentially affecting future academic or career opportunities. Professional consequences include academic integrity violations when students unknowingly cross institutional boundaries around AI use, since policies are often reactively designed, ad hoc, and lag behind technology adoption. This issue is critical because students are naturally gravitating toward AI tools regardless of institutional guidance. Without “secure sandboxes that allow safe experimentation while protecting their data and academic standing,” one participant proposed, institutions could risk either stifling beneficial AI learning or exposing students to long-term harm.

3.1.2 Lack of basic skill acquisition for students offloading cognitive tasks to AI

There is growing evidence that students' use of GenAI involves cognitive offloading, i.e., delegating cognitive tasks to external aids in a way that reduces engagement in deep, reflective thinking (Gerlich, 2025; Shukia et al., 2025). Over time, cognitive offloading is likely to degrade students’ ability to learn key skills and improve them as they progress through education and their careers. How this impacts higher education will depend on what students seek to get out of it; however, some basic academic skills, like writing, reading, and research, are crucial to developing critical thinking, intellectual ability, and curiosity. The effects of this process might not be immediate, but we could soon see the loss of these skills degrading people's ability to engage effectively and autonomously in society as workers and citizens. How can educators ensure that students are still developing their abilities when punitive or abolitionist measures are unlikely to work, given the widespread use of GenAI and the difficulty in detecting its use? As vocations and subjects will likely be impacted differently, AI integration must align with pedagogical goals, ethics, and disciplinary expectations. So universities will need to adopt critical policies to guide classroom instruction and assessment, based on what Corbin et al. (2025, p. 11) call contextual flexibility, in which "[w]hat counts as appropriate AI use is not universal but varies across disciplines, assessment types, and learning objectives."

3.1.3 Resource-strapped universities and time-poor staff struggle to make the best use of AI, potentially triggering feedback loops of institutional decline

Many academic and professional staff recognize the need to develop their AI literacies and review, rethink, and redesign their approaches to curriculum design, teaching, learning, and assessment. Yet as a participant observed, they have neither the time nor the necessary resources to “rebuild the ship at sea.” For instance, a recent report by the American Association of University Professors (2025) found that work intensification and devaluation (rather than pedagogical objectives) were the main reasons for faculty using AI to assist with academic tasks; but that implementing AI actually added to faculty and staff workload. Additionally, resource scarcity can create a defensive mentality among staff that makes it harder for them to work in partnership with students on AI integration. This is likely to become increasingly important given the rapid ongoing development of GenAI tools, their ever-increasing integration into every aspect of life, and the role students expect AI to play in their own futures. Moreover, this situation could be further impacted by the effect of AI on university and college finances, reducing the time and resources that can be allocated to making AI integration work well, and creating the potential for reinforcing feedback between future educational crises that demand immediate reform and the ability of faculty to drive the necessary change to address them.

3.1.4 GenAI-supported learning pathways outside formal instruction, i.e., emergence of shadow curricula

Students are increasingly learning through AI-generated explanations, writing assistants, and prompt-based inquiry that take place without instructor awareness. This 'shadow learning' may develop parallel to, or even replace, aspects of formal instruction, resulting in 'shadow curricula' (Kim & Jung, 2019), as well as challenging academic integrity, equity, and the university’s role in knowledge mediation (Fulsher et al., 2025). GenAI-supported learning processes are often unobservable, personalized, and shaped by algorithmic design, complicating institutional attempts to evaluate or guide student development. They also differ from previous self-learning tools, such as on-line resources or social media, because AI generates more personalized outputs, meaning there is less transitional work (like summarising, redrafting, applying, or translating material) for students to undertake in using what they encounter, and reducing the scope for critical engagement with the shadow curriculum. Additionally, AI allows more diverse learning pathways to support a much wider range of objectives, making its use even more pervasive and atomizing. AI systems create the illusion of a well-curated educational resource, but without human oversight or intention, making it hard to assess their validity and reliability as learning aids. Students who rely on AI systems may also isolate themselves from shared educational experiences.

3.1.5 Equity gaps in access to AI literacy and customization

Students’ background, discipline, and digital confidence play critical roles in their ability to navigate AI tools. There is empirical and anecdotal evidence that students' AI usage mirrors preexisting socio-economic divides (Dinker, 2024) and that more self-confident students use it with greater critical awareness (Lee et al., 2025). As personalized AI grows more prevalent, students with higher AI fluency or access to customized tools will enjoy greater benefits. On the other hand, students from marginalized groups are also more likely to be accused of cheating, but not more likely to have done so (Robey et al., 2022; Harrad et al,, 2024), which is making some groups (e.g. women, lower socioeconomic groups) more cautious about using AI, even when they are permitted to do so (Freeman, 2025). This threatens to widen existing educational inequities across embodied, symbolic, and material differences, unless institutions invest in scaffolding equitable access and AI literacy for all. While some universities have developed policies aimed at mitigating the impact of these inequalities, and levels of AI literacy among staff are generally increasing (Freeman, 2025), many institutions have scarcely begun to address this systematically. What makes this issue more complicated, and potentially more important, is the uncertainty as to whether and how AI fluency and access to customized tools will enhance or degrade education overall.

3.1.6 Personalized AI enables the kind of tutoring that has long been the holy grail of learning science

Personalized AI (pAI) combines general-purpose GenAI models with comprehensive academic content, such as curricula, lectures, and textbooks, as well as individualized private histories and projected future trajectories towards learning and career goals for each learner (assuming these had been mapped-out and aligned by course designers). One recent study has indicated that well-designed pAI tutors are capable of enhancing in-class active learning, both helping students learn and motivating them to engage with a subject (Kestin, 2025). However, education is not only about the quantity of students' learning, but what they are learning, why they want that kind of education, and whether it will help them continue learning in the future (Hendrick, 2025). pAI systems could provide 24/7 tutoring and both proficiency and formative assessment, as well as ongoing and highly personalized academic and career advice, services that higher education currently struggles to provide, and could even interact with each other to facilitate shared experiences such as study groups and group projects. They could even obviate the need for traditional forms of summative assessment, such as exams. However, they would require a large amount of data storage (potentially millions of tokens per student), with associated data management and governance challenges. They also raise important questions for educators about their role and significance, and for institutions about staffing requirements. Finally, pAI will pose many ethical questions, including its potential impact on learners' autonomy (Richter, 2025); If personalized instruction also appears to require dehumanizing instruction, it could have unforeseen and potentially negative consequences.

3.1.7 Declining human mastery of core academic practices

As GenAI systems gain domain-specific proficiency and increasingly automate complex intellectual tasks, there is a risk that early-stage reasoning becomes easier to outsource and human mastery of foundational academic practices will atrophy. Previous technological shifts, like computerized calculation and statistical software that enabled students to undertake modelling without the need to understand the underlying mathematics, have already reshaped curricula (Habibullon et al., 2024). Generative AI could extend similar automation into many other research tasks that many disciplines consider integral to creativity, problem-solving, and original thought. While experiments suggest AI is most effective when used in combination with traditional research, students often prefer using it as a replacement (Kreijkes et al., 2025), which could cause students to lose formative intellectual struggles that cultivate original insights and creative breakthroughs. At the same time, they may develop different, but still valuable skills, like promptcraft or critical awareness of model biases, which could, in turn, enable them to develop creativity and insight through their use of AI. This issue raises questions about what forms of intellectual development higher education should preserve and why. For instance, it may necessitate a greater focus on the cultivation of the mind (Bildung in the Humboldtian tradition), rather than on the acquisition of knowledge and skills. This is not the first time technology has posed such deep questions to the sector, and we should consider how academic practice developed around previous technological revolutions, like the printing press, which negated the need for scholars to memorize large quantities of text, or bureaucratic specialization, which disincentivized the cultivation of generalist and interdisciplinary knowledge (Harari, 2024). However, the changes that could be brought about by AI are likely to be a lot faster, and potentially even more dramatic.

3.2 Longer-Term (emerging in > 5 years)

3.2.1 Emergence of AI learning companions as primary educational relationships

With the onset of significant 'digital assistance' in all aspects of life for young people, advanced AI systems may evolve to become students' primary source of educational guidance and intellectual companionship, potentially displacing traditional faculty-student relationships. “The degree of personalization and interaction is extremely attractive to most,” a participant warned, “and is likely to become more so.” Because of this, learning behaviours of students who have used AI for a significant portion of their education may become dramatically different in 5-10 years than what educators are used to. This shift could fundamentally alter the social dimension of education, reducing opportunities for human modelling of critical thinking, ethical reasoning, and professional development. Yet, these are key to the kinds of cyber-social learning that appear most important for students to get the best out of AI tools (Tzirides et al., 2024; Zapata et al., 2024). The implications of this shift extend across faculty roles, institutional culture, and the development of the social skills necessary for professional collaboration. Universities may need to actively articulate and defend the irreplaceable value of human educational relationships, while determining how to integrate AI companions without undermining human connection.

3.2.2 Developing programme-focused approaches to advance curriculum design

The modular structure of higher education often fragments learning and obscures the overarching disciplinary or professional vision of a programme. As AI increasingly automates routine academic tasks, including how instructors assess and grade student work, we need to reorient education around deeper, more integrated learning experiences supported through meaningful staff-student relationships with and without generative and agentic AI (i.e., AI that is capable of using a wide range of tools and operating for extended periods on a task without additional prompting). Programme-focused assessment can provide this type of education, with its focus on non-graded development meetings that reflect year-level and programme-level learning outcomes. This approach resonates with Boyer's vision of scholarship: the university as a site of academic apprenticeship, where students and staff engage in collaborative inquiry and application (Boyer, 1990). However, implementing it requires a fundamental shift in IT and business systems, including academic quality assurance and regulations, curriculum design, and pedagogical practice, moving from discrete module outcomes assessed separately to holistic programme-level assessment linked to the overarching vision for the programme (see, e.g., Walker, 2025). Such a change will likely take a long time; however, AI could accelerate this shift in mindset by facilitating the collection, evaluation, and transfer of information about students' performance across a wider range of contexts, but only if universities and educators have the space, resources, and incentives to make it.

3.2.3 Persuasive design in AI tools as a scalable vector of influence over vulnerable learners

AI tools are increasingly built with persuasive capacities, delivering responses that are informationally dense, emotionally resonant, rhetorically confident, and designed to steer user decisions (Hackenberg et al., 2025). These capabilities are already being harnessed to improve student engagement and motivation (Alsaiari & Baghaei, 2024), but they clearly raise important ethical issues (Rahman & Adaj, 2024). For example, what if students were persuaded to act in ways that are dangerous or damaging to their mental health by models that were owned or endorsed by an institution or teacher? More generally, persuasive AI could create a powerful channel for influence over students, especially vulnerable learners. When persuasive systems are designed or controlled by corporate, political, or ideological actors, a participant warned, they can serve as “subtle but highly scalable instruments of behavioural and cognitive influence”. The risk is not merely individual misuse, but structural manipulation: student beliefs, values, and learning paths may be shaped in ways that reflect the agendas of those who own or configure the systems. Without transparency, oversight, or critical safeguards, such influence may go unnoticed, entrenching inequality, distorting autonomy, and embedding hidden norms into the heart of education. This risk may be partially mitigated through developing students' (and instructors') critical AI literacy, but previous evidence from the internet and social media suggests that this will be insufficient unless ethical oversight mechanisms and transparency standards are also established for AI tools.

3.2.4 Transfer of pedagogical influence from educators to AI systems

The integration of GenAI into student feedback, study tools, and writing aids is shifting where students look for intellectual guidance. As AI systems increasingly suggest revisions, probe reasoning, and scaffold structure, they assume roles traditionally held by instructors, tutors, and even other students. This could alter students’ perceptions of authority, especially in resource-constrained institutions where human contact is already minimal. Additionally, it could erode the relational foundation of education, where trust, debate, and shared inquiry are essential. If resource pressures drive institutions to replace human tutoring and feedback with AI at scale, students may have even fewer opportunities to engage in authentic dialogue. Differential practices around the world create opportunities for cross-cultural learning about the impacts this is likely to have. For instance, China’s ‘Educational Modernization Plan 2035’ is already driving rapid integration of AI into higher education, including training teachers for, and with, AI-based education methods (Guo & Xu, 2025). On the other hand, different educational traditions also present a variety of social models for education, with many non-Eurocentric traditions offering potentially superior alternatives by prioritizing the development of ethics, cultural awareness, and relational thinking before technical skills (Younas & Zeng, 2025). The issues raised by the transfer of student trust and attention are not just functional, but political and epistemological. Over-reliance on AI guidance could reduce students' ability to question assumptions, as they adapt to values embedded in training data rather than cultivating independent judgment through human interaction.

3.2.5 Institutional competition with AI-native alternative education providers

Universities face emerging competition from AI-powered learning platforms that offer highly personalized, skills-focused education at lower costs and faster timescales than traditional degree programs. The dislocation, particularly in the USA, between the cost of university education and the 'return on investment' for students is a real problem, particularly when employers are finding it increasingly difficult to differentiate student abilities due to pervasive 'grade inflation' problems. If universities cannot demonstrate clear advantages in cultivating critical thinking, adaptability, and deeper learning, they risk losing credibility and relevance in a competitive educational market increasingly shaped by technology. For instance, large graduate employers may lose confidence in the validity of university qualifications and look to increase their offer of degree-level apprenticeships, in collaboration with Edtech providers and consultants who are more adept at using AI to design and administer higher-level courses (Adeoye & Otemuyiwa, 2024). If accreditation systems begin to recognize AI-driven providers, the competitive landscape could shift dramatically, challenging universities to redefine their purpose and distinct contributions beyond credentialing. More agile universities may adapt to offer their own provision of similar courses, but those who are slower to change could lose a key source of funding, with only the most 'prestigious' being able to recoup some of this through endorsement or collaboration. Multiple models could emerge: universities might merge, some will cease to operate, and other forms of higher education could become more popular.

3.2.6 The elimination of college and university as a material/spatial experience owing to the use of AI

As AI tutors, virtual labs, and personalized learning platforms become more sophisticated, the traditional justifications for physical campuses, such as access to expert faculty, specialized facilities, and peer collaboration, weaken. Students may increasingly question the value of expensive on-campus experiences when AI can deliver customized education remotely at a fraction of the cost, especially in systems where residential campuses are currently the norm, as in the USA. It should be noted that university campuses have not been much impacted by the availability of on-line and massively open learning experiences, even after much of higher education became remote only during the 2020-21 COVID pandemic. However, student perceptions are shifting, while many still report preferring in-person education (Photopoulos et al., 2022), sentiments towards remote learning opportunities are improving (Turner et al., 2024). Future erosion of the material campus experience could eliminate crucial elements of human development: spontaneous intellectual encounters, embodied learning through hands-on experiences, and social capital formation through diverse peer networks. The university as a physical space fosters serendipitous discoveries, cross-disciplinary pollination, and the development of social skills essential for professional and personal success. For many students, higher education provides the physical and social space in which they discover and form their character, values, and social networks. As one participant emphasized, losing this risks creating “a generation of technically educated but socially isolated individuals,” while undermining universities' role as community anchors and cultural institutions.

3.2.7 Students need to develop the capabilities to help society mitigate existential risk, such as from AI and nuclear war

Humanity's creation, use, and failure to adequately regulate technologies like nuclear weapons and AI, along with our disruption to global climate systems, pose risks to human existence that are intensifying across multiple domains (Beard et al., 2023; Science and Security Board, 2026). AI poses particularly serious risks through multiple pathways, including: potential misalignment with human interests (Hadshar, 2023), exacerbation of other risks (Bucknall & Dori-Hacohen, 2022), and impacts on democratic institutions (Jungherr, 2023). However, higher education remains largely disconnected from preparing society for these civilizational threats. Universities are producing graduates equipped for a stable world that no longer exists, and tend to prioritise disciplinary specialisation and career and civic preparation. While experts and specialists are still undoubtedly needed, the cross-cutting capabilities required to navigate unprecedented global threats remain underdeveloped, such as critical thinking about complex systems, moral imagination and ethical reasoning under uncertainty, and collaborative governance skills. This represents a potentially significant gap in educational provision during a period of heightened existential risk. While the window for preventive educational intervention remains open, the accelerating nature of AI development and other threats makes timely curriculum reform increasingly critical.

4 Discussion

When presented in their expected chronological order of appearance, these issues begin to tell a kind of narrative about how we might expect AI to impact higher education. Of course, it would be foolish to try to predict just one such scenario given how much of the future depends upon decisions that are yet to be taken (we present a set of five divergent possible scenarios in Beard & Connelly, 2025). However, the ordering of these issues does indicate how a cross-section of the kind of people who will be making these decisions think the future might unfold over the next decade.

How we address these challenges and opportunities will make these scenarios more or less likely. For instance, if we act now to ensure that students using GenAI continue developing the basic skills required for rigorous scholarship, then this is likely to help prevent future declines in expertise for core academic practices and give more people the critical thinking skills required to resist the influence of persuasive AI. Similarly, if we can make good use of the opportunities AI brings to make education more personalized and holistic, this could improve the standard of education that institutions can offer their students by, for instance, providing deeper, more integrated programme-focused learning experiences.

However, we must not be naive about these opportunities. Decades of technological promise have often failed to meaningfully improve educational quality, or even to unseat long-held assumptions about teaching, learning, and educational practice. Companies often tout novel solutions as radical improvements, but when implemented, they turn out to provide only incremental benefits or frustrations. Many participants expressed concern about not getting too caught up in the promise of AI, especially when it is being developed without well-defined use cases by corporations with a track record of irresponsible practice. Reaping whatever benefits may come from AI, and avoiding its pitfalls, is likely to require far more clarity about how, and why, it is to be used, to ensure this actually aligns with the objectives of education.

Other important cross-cutting themes appear in many, or even all, of these issues. One of the most significant is the importance of continuing to ask deep questions about the nature and value of education, rather than letting its future be dictated solely by market forces. Most educators are constrained in their thinking about the traditional ways we design and deliver university education. However, the advent of GenAI is exposing deep fragilities in existing educational practices, with educators suddenly having to interrogate why their favoured approach to learning or assessment matters and how it can be adapted to changing circumstances. If educators have clearer ideas about what they are trying to achieve, and are supported in articulating their importance to students, administrators, and political leaders, then many of their fears about AI (such as a loss of student engagement or critical thinking) could be considerably mitigated. This underlines the need to resist pressure to fit education around AI by simply looking for opportunities to use this technology, and instead retain a humanistic approach that fits AI around educational problems and goals.

Another important theme is the social aspects of education, both in terms of the learning relationships that define current educational experience and the wider social functions of education. It is common to see discussions about the future of AI consider whether or not it will serve ‘human values’ or be under ‘human control’. However, from an educational perspective, it seems just as important to ask whether it will replace ‘human relationships’ or destabilize ‘human communities’. Amid significant interest in the personalization that AI might bring to education, and a widespread shift to more individualized consumer-driven models of education delivery, we should not lose sight of this important social context.

A third cross-cutting theme appearing in many issues is the political and economic context that surrounds the development and deployment of AI in education. Many developers and deployers of AI systems focus primarily on their technical capabilities and limitations. However, their impact on education will be just as much, if not more, determined by the institutions and practices surrounding their use. AI is entering a higher education sector that is already shaped - and degraded - by market-driven logics (at least in the Anglo-American sphere), where many institutions operate under intense financial pressure, while academic labour is casualized, precarious, and undervalued. Its impact on education will both shape and be shaped by this broader context, complicating all these already existing issues, while amplifying and causing a deeper erosion of institutional integrity.

5 Conclusion

Whether viewed as a risk, an opportunity, or some combination of the two, AI is undoubtedly having a significant impact on higher education. As such, it is very easy to get caught up in the immediate challenges facing academics and decision makers. In this study, we have used horizon scanning to try to get beyond this to understand the future issues most likely to impact the academy. While there is a significant degree of uncertainty about the importance and timescale of the issues participants identified, we believe our list provides a useful perspective on the range of possibilities facing the sector. We hope this can help colleagues across higher education, and beyond, to start thinking more clearly about the kinds of future they most want to see and the decisions they can take to move towards them.

References

Acemoglu, D., Autor, D., & Johnson, S. (2023). *Can We Have Pro-Worker AI?: Choosing a Path of Machines in Service of Minds*. Centre for Economic Policy Research.

Adeoye, M. A., & Otemuyiwa, B. I. (2024). *Navigating the Future: Strategies of EdTech Companies in Driving Educational Transformation*. Journal of Educational Research and Innovation Technology, 1(1), 43-50.

Alon, I., Haidar, H., Haidar, A., & Guimon, J. (2025). *The future of artificial intelligence: Insights from recent Delphi studies*. Futures, 165, 103514.

Alsaiari, O., & Baghaei, N. (2024). *Persuasive AI Feedback: Enhancing Student Emotions and Engagement in Higher Education*. The Adjunct Proceedings of the 19th International Conference on Persuasive Technology.

American Association of University Professors (2025). *Artificial Intelligence and Academic Professions*. https://www.aaup.org/sites/default/files/2025-07/TREP-Artificial-Intelligence-and-Academic-Professions.pdf

Aoun, J. E. (2017). *Robot-proof: higher education in the age of artificial intelligence*. MIT Press.

Association of American Universities & Colleges (2025). *Survey of college and university leaders*. https://www.aacu.org/newsroom/higher-education-leaders-navigate-ai-disruption.

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2024). *Generative AI can harm learning*. The Wharton School Research Paper. http://dx.doi.org/10.2139/ssrn.4895486.

Baum, S. D. (2020). *Medium-term artificial intelligence and society*. Information, 11(6), 290.

Beard, S., Rees, M., Richards, C., & Rios Rojas, C. (2023). *The era of global risk: an introduction to existential risk studies*. Open Book Publishers.

Beard, S., & Connelly, M. (2025). B*ackground Briefing and Five Scenarios for Higher Education in the Age of Transformative AI*. SSRN: https://ssrn.com/abstract=5397691

Bengio, Y., Clare, S., Prunkl, C., Murray, M., Andriushchenko, M., Bucknall, B., … & Mindermann, S. (2026). I*nternational AI Safety Report 2026*. UK Department for Science, Innovation and Technology. https://internationalaisafetyreport.org

Bowen, J. A., & Watson, C. E. (2025). *Teaching with AI: A practical guide to a new era of human learning*. JHU Press.

Boyer, E. L. (1990). *Scholarship reconsidered: Priorities of the professoriate*. Princeton University Press.

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Digital Economy.

Bucknall, B. S., & Dori-Hacohen, S. (2022). *Current and near-term AI as a potential existential risk factor*. Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (pp. 119–129). Association for Computing Machinery. https://doi.org/10.1145/3514094.3534146

Cave, S., & Ó hÉigeartaigh, S. S. (2019). *Bridging near-and long-term concerns about AI*. Nature Machine Intelligence, 1(1), 5-6.

Cline, A., & Kaymak, B. (2025). *Are Young College Graduates Losing Their Edge in the Job Market?*. Federal Reserve Bank of Cleveland, Economic Commentary 2025-14. https://doi.org/10.26509/frbc-ec-202514.

Dalziel, M. V., Schaffer, K., & Martin, N. (2024). *Navigating AI in social work and beyond: A multidisciplinary review*. arXiv preprint https://arxiv.org/abs/2411.07245

Dinker, N. (2024). *Artificial intelligence and inequality: examining the social divides created by technological advancements*. International Journal of Innovations in Science, Engineering and Management, 228-236.

Eriksson, M., Purificato, E., Noroozian, A., Vinagre, J., Chaslot, G., Gomez, E., & Fernandez-Llorca, D. (2025). *Can we trust AI benchmarks? an interdisciplinary review of current issues in ai evaluation*. arXiv preprint https://arxiv.org/abs/2502.06559.

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., Gašević, D (2024). *Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance*. British Journal of Educational Technology, (56(2)), 489-530.

Freeman, J. (2025). *Student generative AI survey 2025*. Higher Education Policy Institute: London, UK.

Gallant, T. B., & Rettinger, D. A. (2025). *The Opposite of Cheating: Teaching for Integrity in the Age of AI*. University of Oklahoma Press.

Gerlich, M. (2025). *AI tools in society: Impacts on cognitive offloading and the future of critical thinking*. Societies, 15(1), 6.

Gokhale, A. (2025). *Countering flaws in algorithm design and applications: a Delphi study*. AI & SOCIETY, 40(5), 3721-3733.

Grace, K., Sandkühler, J. F., Stewart, H., Weinstein-Raun, B., Thomas, S., Stein-Perlman, Z., ... & Korzekwa, R. C. (2025). *Thousands of AI authors on the future of AI*. Journal of Artificial Intelligence Research, 84.

Gruetzemacher, R., & Whittlestone, J. (2022). *The transformative potential of artificial intelligence*. Futures, 135, 102884. https://doi.org/10.1016/j.futures.2021.102884.

Guo, Z., & Xu, Y. (2025). *AI-empowered integrated reform of teacher education: The experience of a Chinese university*. Education and Lifelong Development Research, 2(1), 30-38.

Habibulloh, M., Sholeh, M. I., & Idawati, K. (2024). *Exploring Technological Innovations and Approaches in Modern Education*. Journal of Studies in Academic, Humanities, Research, and Innovation, 1(1), 56-76.

Hackenburg, K., Tappin, B. M., Hewitt, L., Saunders, E., Black, S., Lin, H., ... & Summerfield, C. (2025). *The levers of political persuasion with conversational artificial intelligence*. Science, 390(6777), eaea3884.

Hadshar, R. (2023). *A review of the evidence for existential risk from AI via misaligned power-seeking*. arXiv preprint https://arxiv.org/abs/2310.18244

Hanea, A. M., McBride, M. F., Burgman, M. A., Wintle, B. C., Fidler, F., Flander, L., ... & Mascaro, S. (2017). *I nvestigate D discuss E stimate A ggregate for structured expert judgement*. International journal of forecasting, 33(1), 267-279.

Hanea, A. M., McBride, M. F., Burgman, M. A., & Wintle, B. C. (2018). *The value of performance weights and discussion in aggregated expert judgments*. Risk Analysis, 38(9), 1781-1794.

Harari, Y. N. (2024). *Nexus: A brief history of information networks from the Stone Age to AI*. Random House.

Harrad, R., Keasley, R., & Jefferies, L. (2024). *Academic integrity or academic misconduct? Conceptual difficulties in higher education and the potential contribution of student demographic factors*. Higher Education Research & Development, 43(7), 1556-1570.

Hemming, Victoria, Mark A. Burgman, Anca M. Hanea, Marissa F. McBride, and Bonnie C. Wintle. *A practical guide to structured expert elicitation using the IDEA protocol*. Methods in Ecology and Evolution 9, no. 1 (2018): 169-180.

Hendrick, C. (2025). *The Algorithmic Turn: The Emerging Evidence On AI Tutoring That's Hard to Ignore: Are We Approaching A Turing Test for Teaching?*. The Learning Dispatch.

Ismail, I. A., & Aloshi, J. M. R. (2025). *Data privacy in AI-driven education: An in-depth exploration into the data privacy concerns and potential solutions*. In AI Applications and Strategies in Teacher Education (pp. 223-252). IGI Global.

Jin, Y., & Sercu, L. (2025). *ChatGPT Interventions in Higher Education: A Systematic Review of Experimental Studies*. Journal of Computer Assisted Learning, 41(4), e70072. https://doi-org.ezproxy.cul.columbia.edu/10.1111/jcal.70072.

Jungherr, A. (2023). *Artificial intelligence and democracy: A conceptual framework*. Social media+ society, 9(3), 20563051231186353.

Kalaitzidis, T. (2025). *How generative AI fixes what higher education broke: reframing the "colonization" of knowing and learning*. AI & SOCIETY, 1-19.

Kemp, L., Adam, L., Boehm, C. R., Breitling, R., Casagrande, R., Dando, M., ... & Sutherland, W. J. (2020). *Bioengineering horizon scan 2020*. Elife, 9, e54489.

Kılıçkaya, F., & Kic-Drgas, J. (2025). *The future of tutoring: AI as an augmenting force*. AI & SOCIETY, 1-2.

Kim, Y. C., & Jung, J. H. (2019). *Conceptualizing shadow curriculum: Definition, features and the changing landscapes of learning cultures*. Journal of Curriculum Studies, 51(2), 141-161.

Kreijkes, P., Kewenig, V., Kuvalja, M., Lee, M., Hofman, J. M., Vitello, S., ... & Oates, T. (2025). *Effects of LLM use and note-taking on reading comprehension and memory: A randomised experiment in secondary schools*. Computers & Education, 105514.

Lee, H. P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). *The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers*. Proceedings of the 2025 CHI conference on human factors in computing systems (pp. 1-22).

Levine, A., & Van Pelt, S. (2021). *The great upheaval: Higher education's past, present, and uncertain future*. JHU Press.

Niloy, A. C., Akter, S., Sultana, N., Sultana, J., & Rahman, S. I. U. (2024). *Is ChatGPT a menace for creative writing ability? An experiment*. Journal of computer assisted learning, 40(2), 919-930.

Oxford University Press. (2023). Researchers and AI: Survey findings. https://fdslive.oup.com/www.oup.com/academic/pdf/Researchers-and-AI-survey-findings.pdf.

Photopoulos, P., Tsonos, C., Stavrakas, I., & Triantis, D. (2022). *Remote and in-person learning: Utility versus social experience*. SN computer science, 4(2), 116.

Rahman, P., & Adaji, I. (2024). *Ethics in persuasive technologies: a systematic literature review*. Proceedings of the International Conference on Mobile and Ubiquitous Multimedia (pp. 106-118).

Reynolds, S. A., Beery, S., Burgess, N., Burgman, M., Butchart, S. H., Cooke, S. J., ... & Sutherland, W. J. (2025). *The potential for AI to revolutionize conservation: a horizon scan*. Trends in ecology & evolution, 40(2), 191-207.

Richter, F. (2025). *The ethical acceptability of personalization via intelligent systems in education*. AI & SOCIETY, 1-15.

Robey, N., Dunn, R., Haskins, N., & Dickter, C. (2022). *Disparities in student misconduct in higher education for students from underrepresented racial backgrounds or non-citizen visa status*. Journal of College and Character, 23(3), 238-257.

Ruediger, D., Blankstein, M. & Love, S. (2024). *Generative AI and postsecondary instructional practices*. Ithaka S+ R. https://www.jstor.org/stable/resrep61195.

Science and Security Board. (2026). *2026 Doomsday Clock statement*. Bulletin of the Atomic Scientists. https://thebulletin.org/doomsday-clock/2026-statement/.

Schleicher, A. (2018). *Educating learners for their future, not our past*. ECNU Review of Education, 1(1), 58-75.

Shukla, P., Bui, P., Levy, S. S., Kowalski, M., Baigelenov, A., & Parsons, P. (2025). *De-skilling, cognitive offloading, and misplaced responsibilities: potential ironies of AI-assisted design*. Proceedings of the CHI Conference on Human Factors in Computing Systems.

Sparrow, R., & Flenady, G. (2025). *Bullshit universities: the future of automated education*. AI & SOCIETY, 1-12.

Sposato, M. (2025). *Artificial intelligence in educational leadership: a comprehensive taxonomy and future directions*. International Journal of Educational Technology in Higher Education, 22(1), 20. https://doi.org/10.1186/s41239-025-00517-1.

Sutherland, W. J., Brotherton, P. N., Butterworth, H. M., Clarke, S. J., Davies, T. E., Doar, N., ... & Thornton, A. (2025). *A horizon scan of biological conservation issues for 2025*. Trends in Ecology & Evolution, 40(1), 80-89.

Tran, T. M., Bakajic, M., & Pullman, M. (2024). *Teacher's pet or rebel? Practitioners' perspectives on the impacts of ChatGPT on course design*. Higher Education, 1-24. https://doi.org/10.1007/s10734-024-01350-7.

Turner, R., Webb, O., & Pritchard, C. (2024). *'A freedom of students to choose': Student and staff perspectives on the future role of online learning in higher education*. Innovations in Education and Teaching International, 61(6), 1164-1183.

Turoff, M., & Linstone, H. A. (Eds.). (1975). *The Delphi Method: Techniques and Applications*. Addison-Wesley.

Tzirides, A. O. O., Zapata, G., Kastania, N. P., Saini, A. K., Castro, V., Ismael, S. A., ... & Kalantzis, M. (2024). *Combining human and artificial intelligence for enhanced AI literacy in higher education*. Computers & Education Open, 6, 100184.

Tzirides, A. O., Zapata, G. C., Bolger, P., Cope, B., Kalantzis, M., & Searsmith, D. (2025). *Assessing AI-generated feedback quality: Instructor perspectives on fine-tuned GenAI for higher education student writing*. In A. Rajala, A. Cortez, R. Hofmann, A. Jornet, H. Lotz-Sisitka, & L. Markauskaite (Eds.), ICLS 2025 Proceedings: Educating for world-making: Envisioning and enacting sustainable solutions to global crises (pp. 898-907). ISLS.

Walker, S. (2025). *Trends in assessment in higher education: considerations for policy and practice*. Jisc Report https://repository.jisc.ac.uk/9887/1/trends-in-assessment-report.pdf.

Wintle, Bonnie C., Mahlo NC Kennicutt II, and William J. Sutherland (2024). *Scanning Horizons in Research, Policy and Practice*. In SJ Beard and Tom Hobson ed. An Anthology of Global Risk. Open Book.

Younas, A., & Zeng, Y. (2025). *Moral grounding before algorithms: a cross-cultural critique of AI education in schools*. AI & SOCIETY, 1-9.

Zapata, G. C., Cope, B., Kalantzis, M., Tzirides, A. O, Saini, A. K., Searsmith, D., … & Abrantes da Silva, R. (2025). *AI and peer reviews in higher education: Students' multimodal views on benefits, differences and limitations*. Technology, Pedagogy and Education, 34(5), 583-601. https://doi.org/10.1080/1475939X.2025.2480807.

Zapata, G. C., Saini, A., Tzirides, A. O, Cope, B., & Kalantzis, M. (2024). *The role of AI feedback in university students' learning experiences: An exploration grounded in Activity Theory*. Ubiquitous Learning: An International Journal, 18(2). https://doi.org/10.18848/1835-9795/CGP/v18i02/1-30.

Appendix A

Detailed Methodology

The horizon scan involved a diverse group of 27 participants blending academic, industry, and policy expertise related to AI and/or education. 21 participants submitted short summaries of 2-5 issues relating to AI and higher education that they considered to be ‘on the horizon’: issues that had not yet become prominent in discussions but had the potential to substantially impact the sector. To ensure comparability between issues, participants were asked to frame them at similar levels of granularity. Very broad issues encompass a whole suite of more detailed issues and tend to be already well known and too vague to inform decision-making. On the other hand, very narrow issues, such as those only affecting a single academic programme, are unlikely to substantially impact higher education as a whole. To help ensure that issues were submitted with an appropriate level of granularity, the following example topic was circulated, framed at five different scales, with participants asked to submit issues between levels 2 and 4 in terms of granularity:

1. Regulation of generative AI,
2. Regulation of generative AI as a writing aid,
3. Establishing clear norms and practices for the use of generative AI as a writing aid,
4. Establishing clear norms and practices for the use of generative AI through university writing courses,
5. Establishing a generative AI module as part of the university writing course at Columbia University.

The exercise facilitators also produced a briefing document that summarized some of the research and uncertainties around future AI trajectories and their potential to impact higher education, as well as developing a set of five possible scenarios to stimulate creative reflection (Beard & Connelly, 2025).

A total of 72 issue summaries were collected, anonymized, and circulated to all participants. Participants were assigned a subset of 10 issues to consider in more detail and asked to evaluate each according to its likelihood, timescale, and novelty, as well as to raise points of reflection and discussion. We used principal components analysis to study the relationship between the likelihood, timescale, and novelty of issues in order to identify those that were on the horizon and those that were already well known. Each of the listed issues had 2–3 investigators assigned to it, who were not the person who submitted the issue and were not necessarily experts in that particular topic. This meant that workshop discussion could include the person who submitted the topic, others who were already knowledgeable in the area, and the people who had previously been assigned the topic, allowing for a more informed discussion.

We held two separate discussion workshops, both conducted remotely over Zoom in September 2025. In the first workshop, 10 participants discussed the issues under 11 overarching themes, in order to explore the connections between them and the general challenges and opportunities they posed to the sector. In the second workshop, 16 participants discussed specific issues that had been identified using their survey responses as falling into one of three groups:

1. issues that were especially speculative, long-term, or novel and that participants might therefore be less familiar with,
2. issues where there was the most disagreement, indicating different perceptions or judgements about them, and
3. issues that could be removed, merged, or edited due to being close duplicates, not truly ‘on the horizon’, or poorly framed.

Participants were prompted to consider whether they had enough information to evaluate these issues, what was novel about them, and what particular challenges and opportunities they presented. Feedback from the survey and workshop discussions was used to reduce the list of issues down to a shortlist of 40 that were most clearly significant horizon issues, and to redraft these to ensure their most important aspects were well presented.

Finally, participants individually and confidentially scored each of these 40 issues according to their importance and the timescale over which they would become most important. The first 11 responses were used to further narrow down the shortlist to the 31 highest importance issues, to reduce the burden being placed on the remaining participants. These importance scores were converted to standardised Z scores; that is, the mean and standard deviation of each participant's set of scores were first calculated, then each score in the set was standardized by subtracting the mean and dividing by the standard deviation. The resulting Z scores retain information about the distribution of scores and can be meaningfully aggregated across participants who have provided sets of scores with different distributions. At the end of the Aggregation process, the 14 top-scoring issues (based on these Z scores) were presented to participants for any final comments and feedback, together with the rest of this paper.

Appendix B

List of Participants

Name	Institutional Affiliation	Speciality	Completed Identification Survey	Completed Discussion survey	Attended Discussion Workshop	Completed Evaluation Survey
James Beadle	Cambridge University Press & Assessment	Private Sector	Yes	Yes	Yes	Yes
Mehdi Bilgrami	Pearson	Private Sector	Yes	Yes	Yes	Yes
Gareth A. Brinkworth	University of Cambridge	Education Research	Yes	Yes	Yes	Yes
Clare Carroll	Pearson	Private Sector	Yes	Yes		
Lydia Chilton	Columbia University	AI Research	Yes	Yes	Yes	
Martin Compton	University of East London	Education Research	Yes	Yes	Yes	Yes

Kate Crawford	University of Southern California	AI Research	Yes	Yes	Yes	Yes
Judith Donath	Harvard / MIT	AI Research	Yes		Yes	
Megan Ennion	University of Cambridge	Education Research	Yes		Yes	Yes
Lance Gharavi	Arizona State University	Education Research	Yes	Yes	Yes	Yes
Claire Gordon	London School of Economics	Education Research	Yes			
Velislava Hillman	Goldsmiths, University of London	Education Research	Yes		Yes	Yes
Amy Hungerford	Columbia University	Education Leadership		Yes	Yes	Yes
David Madigan	Northeastern University	Education Leadership	Yes	Yes	Yes	Yes
Alexandru Marcoci	University of Cambridge	AI Research	Yes	Yes		Yes
Alice E.Marwick	Data & Society	AI Research	Yes		Yes	
Alondra Nelson	Institute for Advanced Study	Education Research	Yes		Yes	
Sebastián Otero	Columbia University	Education Research		Yes	Yes	Yes
Matthew Pittinsky	Arizona State University	Private Sector	Yes	Yes	Yes	Yes
Vincent Ponzo	Amazon	Private Sector	Yes	Yes	Yes	
Bruce Schneier	Harvard	AI Research			Yes	
Miguel Urquiola	University of Columbia	Education Leadership		Yes	Yes	Yes
Simon Walker	London School of Economics	Education Research	Yes	Yes	Yes	Yes
Dan Wang	Columbia Business School	Private Sector	Yes		Yes	Yes
Meg Young	Data & Society Institute	AI Research		Yes	Yes	
Gabriela C. Zapata	University of Nottingham	Education Research	Yes	Yes		Yes

Supplementary material

- Longlist of 72 issues from the identification stage
- Anonymized summary of results from the discussion stage surveys
- Anonymized summary of results from the evaluation stage survey