Universities and AI: Developing New Models of Teaching and Learning in the Realm of Radical Uncertainty

A report by Dr James Ransom and Dr Richard Whittle.
Commissioned by the N-TUTORR National Digital Leadership Network

National Digital Leadership Network

Universities and AI: Developing New Models of Teaching and Learning in the Realm of Radical Uncertainty

A report by Dr James Ransom and Dr Richard Whittle

Commissioned by the N-TUTORR National Digital Leadership Network

ISBN 978-1-0683806-4-8

Author biography

Dr James Ransom

Dr James Ransom is a higher education specialist whose work looks at how universities can help adapt to and solve challenges facing society, including rapid technological change. He is an Honorary Senior Research Fellow at UCL Institute of Education, Head of Research at the National Centre for Entrepreneurship in Education (NCEE), and a Specialist Advisor on higher education to the European Bank for Reconstruction and Development. Previous work includes projects for the British Council, the Royal Society, and the British Academy, as well as jobs in policy at Universities UK, UNESCO, and the Association of Commonwealth Universities.

Dr Richard Whittle

Dr Richard Whittle is an expert in the economics of Artificial Intelligence. He researches Artificial Intelligence and Human Decision-Making and has published in world-leading journals such as Work Employment and Society, Public Administration, and the Cambridge Journal of Economics. Richard led the technical research for the Greater Manchester Independent Prosperity Review, and he is an academic advisor to the Manchester Digital Strategy. Richard has received research funding from numerous sources, including the Economic and Social Research Council (ESRC), UK Research and Innovation (UKRI), Research England, and the Money Advice Service, and he has been the recipient of a personal Capabilities in Academic Policy Engagement (CAPE) fellowship hosted at the Institute of Innovation and Public Purpose, UCL.

Abstract

This report investigates the implications of Artificial Intelligence (AI) on Learning, teaching and research. It presents a brief history of artificial intelligence and education to centre the reader, equipping them with insight into the future of these technologies. In this report we present a detailed examination of current AI use and considerations for its safe and ethical deployment. We conclude with horizon scanning and recommendations for educational establishments beginning to incorporate AI.

Introduction to the National Digital Leadership Network Report Series

The National Digital Leadership Network (NDLN) is a collaborative initiative designed to support digital transformation across Ireland's Technological Higher Education sector. Established under the N-TUTORR programme with funding provided through the EU's NextGenerationEU initiative, the network was officially launched in November 2024 to provide a national platform for digital leadership and complementary knowledge exchange and strategic collaboration. While the N-TUTORR programme has now concluded, our network continues its work under the guidance of a steering board composed of sector leaders and external experts.

Digital leadership in higher education extends far beyond technical expertise or the adoption of certain tools and platforms: it's about vision, strategy, and culture change. Effective digital leaders ensure that digital strategies and developments align with institutional and national priorities, not only enhancing teaching, learning, research, and administration functions but also upholding academic values, promoting equity, and driving business innovation. In this context, the NDLN fosters collaboration among higher education leaders, policymakers, and practitioners, providing opportunities to share insights, explore emerging challenges, and develop shared solutions.

As part of its work, the NDLN has commissioned a series of horizon-scanning reports authored by leading national and international scholars and practitioners. These reports explore key trends at the intersection of digital innovation, traditional leadership and strategic planning, providing actionable insights to support higher education institutions in aligning these trends and related opportunities with institutional and national priorities. Covering topics such as the evolving role of generative AI in academia, data-driven decision-making, academic integrity, new models of learning and teaching and new ways to plan for financial sustainability, this report series offers timely advice and direction for higher education leaders navigating the interrelated complexities of the digital and post-digital age.

We extend our gratitude to the N-TUTORR programme for its financial support, and to N-TUTORR Co-ordinator Dr Sharon Flynn for her direction and continued support of the network. Thank you also to members of our national steering board and to our external contributors, in particular Professor Lawrie Phipps.

A big personal thank you in addition to my colleagues in the Department of Technology Enhanced Learning (TEL) at MTU -- especially Darragh Coakley and Marta Guerra -- whose work has been vital to the preparation and publication of these reports. We are also very grateful to Dr. Catherine Cronin, our chief editor, and, of course, to all our authors whose insights, expertise, and dedication form the heart and foundation of this series.

We invite you to engage with these reports and join us in shaping the future of digital leadership in higher education.

Dr Gearóid Ó Súilleabháin
Department of Technology Enhanced Learning (TEL)
Munster Technological University (MTU)

Table of Contents

Executive Summary

While artificial intelligence (AI) has been intertwined with universities and education from its earliest days, recent developments in generative AI – over the past two years – have seen a profound shift in the accessibility of AI tools.

Large language models (LLMs) such as ChatGPT and Claude can now produce outputs that are realistically similar to what a human would produce in response to instructions known as prompts. It is now possible to ask an AI model to produce an essay on any subject, a film or image, or computer code. These models are widely available and exceptionally popular. Indeed, ChatGPT is the fastest-growing application in history, outpacing the adoption of Facebook, Twitter, and TikTok.

Despite the considerable excitement around these new technologies, they may not live up to the hype which surrounds them. They can get things wrong, convincingly impart false information, and produce incoherent work. There are also a large number of ethical concerns around generative AI, as will be raised in this report. Nevertheless, even current incarnations of these models are proving to be significantly disruptive in the academy. This is because their reliability has reached the level of practicality: although the answer they give may well contain inaccuracies – and will likely not be as good as the version provided by a skilled knowledgeable human – it will be an answer to the question given. Often this output will only be "good enough", with the user deciding when a near-instant and exceptionally cheap output is "good enough" when compared to a time consuming and costly one.

The journey from ChatGPT 3 to the latest models is one of increasing competence: LLM-generated outputs have significantly improved in quality and capability, and recent reasoning models (such as OpenAI o1) demonstrate new levels of possibility in this area.

The impact of AI is uncertain. Numerous questions remain unanswered and are likely currently unanswerable. Ultimately, however, the fundamental discussion is of value – how do we value skills and knowledge when AI can replicate facsimiles of these instantly?

The impact of AI is uncertain. Numerous questions remain unanswered and are likely currently unanswerable.

This report addresses the question of uncertainty by adopting Stuart Russell's definition of machine intelligence: "machines are intelligent to the extent their actions are expected to achieve their objectives" (Russell, 2021). This definition allows us to

address a key fact about AI – we are now at the point where, in an educational setting, we can reliably expect an AI system to accomplish the objectives we give it. Addressing this allows educational institutions to cut through some uncertainty and possibly some decision paralysis to take proactive action, rather than react to the latest AI developments in an endless cycle of AI advancement and institutional adaptation.[1]

This report provides the reader with a comprehensive overview of the current state of play concerning AI and universities, including a brief history of AI and a discussion of how it is being used within educational institutions. We end with a broad recommendation for institutions grappling with AI adoption. AI is best incorporated based on a wholesale understanding of what it means, not a simple to-do list.

AI has reached a level of practicality where, even with no further advancements, it will impact all areas of teaching and learning. Institutional decision-makers, academics, and students should be equipped with the skills to understand AI and make informed choices about its use. This report aims to support institutions on this journey.

> As AI continues to permeate higher education, universities stand at a crossroads of significant transformation. Navigating AI hype cycles strategically allows institutions to adopt technologies that genuinely enhance education. The evolving relationships between educators and students necessitate a re-evaluation roles and the intrinsic value placed on the learning journey.
>
> By focusing on critical thinking, ethical considerations, and the cultivation of uniquely human skills, universities can ensure that education remains a meaningful and enriching endeavour. Embracing these changes thoughtfully will prepare students to thrive in an AI-enhanced world, where the true value of knowledge lies not in its effortless showcasing but in its effective and responsible application. (Whittle,2024)

Getting Started: AI and Education

AI has long stood at the intersection of human ambition and technological possibility. Since its conceptual origins in the 1950s, AI has primarily existed as a specialised academic discipline, pursued by computer scientists, mathematicians, and cognitive psychologists seeking to understand and replicate human intelligence through computational means. AI has developed from early rule-based systems to machine learning algorithms that can master specific tasks.

[1]Generative artificial intelligence is also explored in two further reports in the National Digital Leadership Network report series; see Weller (2024) and Pratschke (2024).

However, the emergence of LLMs marks a pivotal transformation in AI's role and reach. These complex systems, trained on vast amounts of human-generated data, have dramatically widened access to AI capabilities, shifting AI from a narrow technical discipline to a universally relevant tool for teaching and learning. This transformation represents more than just a technological advancement; it signifies a fundamental change in how knowledge can be accessed, processed, and transmitted. Ultimately, AI has moved from being a specialist tool to one that is easily accessible. Higher education needs to adapt.

This is not to say that AI tools are infallible, simply that they are capable of producing human-realistic outputs can be generated almost instantly, at low cost and with little effort. The first mass market, broadly applicable AI tools are here.

The implications of this shift are profound. What was once the purview of specialised researchers working on specific problem domains has evolved into a technology that researchers, teachers, and students can leverage across virtually every field of study. From helping students understand complex concepts through interactive dialogue to assisting educators in creating personalised learning materials, AI has rapidly become part of the educational ecosystem. It has also created a significant amount of uncertainty: uncertainty over which skills to teach (and how to teach them), uncertainty over the reliability and validity of assessments, and more fundamental uncertainty over the changing nature of knowledge and the purpose of education.

The pace at which AI is being developed and adopted only increases this uncertainty. Educational institutions worldwide are participants in a mass experiment; the prize is robust, futureproofed, AI-enhanced learning. The costs of getting it wrong could be measured in generational terms.

LLMs may plateau and never live up to the hype of their key proponents. However, even their current capabilities are poised to disrupt teaching and learning models. First, we will look at where we are and how we got here.

Introduction to AI in Education

At its core, AI is a field in computer science focused on developing systems that can emulate human cognitive functions, including learning, reasoning, problem-solving, and pattern recognition. These systems utilise complex algorithms, neural networks, and advanced data-processing capabilities to analyse vast amounts of information, recognise intricate patterns, and make informed decisions, often improving their performance through experience and iterative learning. The fundamental goal of AI is

to create machines that can adapt to new situations, understand context, and perform tasks that traditionally required human intelligence.

In the educational context, AI is found in a broad spectrum of applications, from personalised learning platforms to comprehensive administrative tools. Adaptive learning systems like Duolingo (a popular language learning app) leverage AI to tailor educational content to individual student progress, automatically adjusting difficulty levels and reviewing concepts based on performance patterns. These systems continuously analyse student responses, identify knowledge gaps, and create customised learning pathways that optimise the learning experience. Beyond direct instruction, AI powers predictive analytics systems that help institutions optimise enrolment processes, predict student outcomes, and streamline administrative tasks. These tools can analyse historical data to identify at-risk students, suggest interventions, and help administrators make data-driven decisions about resource allocation and programme development.

Generative AI represents a particularly disruptive shift in educational technology, exemplified by broad tools like ChatGPT, Claude, Midjourney, and DALL-E. These systems can produce (a facsimile of) original content – whether text, images, or other media – in response to user prompts. One can instruct the model to produce (for example) a poem in the style of Shakespeare and it will. ChatGPT, Claude, and others can assist with writing, provide detailed explanations of complex concepts, and engage in interactive tutorials that adapt to student understanding. Midjourney creates sophisticated visual art and educational illustrations that can bring abstract concepts to life, while DALL-E generates customised images that can enhance educational materials and boost student engagement. These generative AI tools are challenging traditional ideas of content creation in education, allowing teachers to quickly develop customised learning materials and enabling students to explore new means of knowledge provision.

The integration of AI in education presents significant opportunities for enhancing the learning experience. Personalised learning platforms (as per the previous section) analyse individual student performance data to automatically adjust content difficulty, pacing, and teaching methods, creating tailored educational journeys that maximise engagement and retention. For instance, AI systems can identify when students struggle with specific concepts, and they can automatically provide additional practice materials or alternative explanations to ensure no student falls behind or becomes disengaged due to content that is too easy or difficult. Students themselves may use AI systems to transform their learning materials to their preferences[2]

[2]For instance, the proliferation of AI podcasting tools can easily change recommended reading into a convincing podcast discussion for the student to listen to.

AI's automation capabilities may reduce the administrative burden placed on teachers by handling time-consuming tasks like grading multiple-choice assessments, checking grammar in essays, and providing initial feedback on assignments. This automation allows educators to dedicate more time to crucial activities that require human insight, such as providing nuanced feedback on complex assignments, developing innovative teaching strategies, and offering emotional support to students. We must note the changing dynamics which come from such uses. Will the teacher–student relationship change if feedback is simply automated? Moreover, AI-powered assistive technologies may have uses for students with disabilities through tools like real-time speech-to-text conversion, automated sign language interpretation, and adaptive interfaces that adjust to individual needs.

However, the widespread adoption of AI in education also presents significant challenges that require careful consideration. Academic integrity has become increasingly complex with the emergence of sophisticated AI writing tools that can generate essays, solve maths problems, and complete coding assignments. Educational institutions must develop new assessment strategies and policies to ensure authentic learning while acknowledging AI's role as a learning tool. Furthermore, AI systems can inadvertently perpetuate societal biases present in their training data, potentially disadvantaging certain student groups through biased content recommendations or unfair assessment practices. These systems might favour specific learning styles, cultural perspectives, or language patterns, making it crucial to implement rigorous testing and monitoring protocols to ensure equitable educational outcomes.

These tools are not perfect, perhaps not even terribly good, but they are quick and cheap. Perhaps there are some occasions where free and "good enough" is acceptable?

A Brief History of AI and Education

The foundations of AI trace back to Alan Turing's1950 paper "Computing machinery and intelligence", which posed the question: "can machines think?" This philosophical inquiry not only laid the groundwork for AI development but also sparked decades of debate about the nature of intelligence and computation. The 1956 Dartmouth Conference, organised by John McCarthy and Marvin Minsky, formally established AI as a distinct academic discipline. This summer-long workshop brought together leading researchers to explore ways of simulating human intelligence through computational methods. The conference's focus on logical reasoning and problem-solving algorithms established the theoretical framework that would guide AI development for decades, emphasising symbolic manipulation and rule-based systems that attempted to codify human knowledge into machine-readable formats.

The practical application of AI in education began with the PLATO (Programmed Logic for Automatic Teaching Operations) system in the 1960s. Developed at the University of Illinois, PLATO represented a quantum leap in educational technology, connecting thousands of terminals worldwide through innovative features that we now consider fundamental to online learning. The system pioneered tools like instant messaging, online forums, and automated assessment systems, demonstrating the potential for technology to facilitate both individualised instruction and collaborative learning at scale. The 1970s witnessed the emergence of sophisticated expert systems, with MYCIN standing out as a particularly successful application in medical education. This system helped medical students diagnose bacterial infections by simulating the decision-making processes of experienced physicians, showcasing the potential of AI to support complex professional education. However, it was never used in practice owing to ethical (and legal) concerns regarding an incorrect diagnosis, and to the impracticability of the hardware in a time before personal computers. However, these early successes catalysed the development of Intelligent Tutoring Systems and Interactive Learning Environments, such as systems like AutoTutor, which could engage students in natural language dialogues, adapt to their responses, and provide personalised feedback based on their learning patterns.

The modern landscape of educational AI has evolved, with companies deploying far more sophisticated adaptive learning platforms that create highly personalised learning pathways. These systems utilise advanced algorithms to analyse student performance data in real-time, adjusting content difficulty, pacing, and teaching methods to optimise learning outcomes. The recent advent of LLMs, particularly generative pre-training transformer (GPT) technology, has potentially significantly disrupted educational accessibility and content creation. These models can generate contextually relevant educational materials, provide detailed explanations across diverse subjects, and offer immediate, personalised assistance to learners worldwide. Their ability to understand and generate human-like text has opened new possibilities for automated tutoring, content creation, and language learning, while raising important questions about the future role of traditional educational methods.

> The modern landscape of educational AI has evolved...

The history of AI development has been marked by significant fluctuations, known as "AI winters", when enthusiasm and funding waned due to unmet expectations and technological limitations. The first major winter occurred in the 1970s, when early excitement about machine translation and expert systems confronted practical limitations (see MYCIN, discussed earlier). Another significant downturn followed in the 1990s, when neural network research hit computational barriers. However, the current AI renaissance – powered by breakthroughs in deep learning, new computational resources, and vast datasets – has generated sustained optimism (a new hype?). Unlike previous cycles, modern AI has – this time – become widely used outside of narrow specialisms. It has the potential for significant disruption but correspondingly a pressure for rapid development and new capabilities to stave off a new, and much more public, AI winter. Indeed, such is public engagement with AI in its current boom that any following AI winter may be particularly cold.

Early attempts to integrate AI in education often struggled to address the complex realities of teaching and learning. The rigid, rule-based approaches of first-generation systems failed to account for the multifaceted nature of education, including social interactions, emotional factors, and diverse learning styles. Initial AI models emphasised knowledge transmission through predefined pathways, neglecting the development of critical thinking skills, creativity, and social-emotional learning. This fundamental misalignment between AI capabilities and educational objectives led to scepticism among educators and researchers, highlighting the need for more sophisticated approaches that better integrate pedagogical principles with technological capabilities. Many early systems also suffered from limited adaptability, and were incapable of adjusting to different cultural contexts, learning preferences, or educational philosophies. These challenges prompted a shift towards more holistic approaches that recognise the complexity of human learning and the importance of combining AI capabilities with sound pedagogical practices.

Generative AI and education

Generative AI is supporting educational practices through its ability to create personalised learning materials and provide instant, contextual support to students and staff alike. LLMs can generate practice problems, explain complex concepts in multiple ways, and even simulate dialogues in foreign languages, enabling a level of customisation previously impossible in traditional educational settings. These tools are particularly valuable for differentiated instruction, as they can adjust their explanations and examples to match each student's learning style, prior knowledge, and pace of understanding. For instance, when teaching a concept like photosynthesis, generative AI can create explanations ranging from simple analogies for younger students to detailed molecular processes for advanced learners.

However, at least at present, we need to ask a simple question: does generative AI provide correct answers? We call an instance of generative AI providing false information a hallucination. We need to be prepared for information being (often convincingly) wrong.

The technology's impact extends beyond content creation to assessment and feedback processes. Generative AI can provide detailed, formative feedback on student work, highlighting areas for improvement while suggesting specific strategies for enhancement. In writing instruction, for example, these systems can analyse essays not just for grammatical accuracy but also for logical flow, argument strength, and stylistic elements, providing specific suggestions for improvement. This capability allows teachers to focus on higher-order aspects of assessment while ensuring students receive timely, detailed feedback on their work.

However, the integration of generative AI in education raises significant pedagogical and ethical challenges that institutions must actively address. The ease with which these tools can generate content has sparked concerns about academic integrity and the development of critical thinking skills. Educators are grappling with questions about how to design assignments that promote genuine learning while acknowledging AI as a legitimate tool for research and composition.

The intersection of generative AI and pedagogical design is creating new possibilities for experiential learning and creative expression. Teachers are using AI to generate realistic scenarios for problem-based learning, create interactive storytelling experiences, and develop customised project prompts that align with individual student interests. In language learning, AI can generate contextually appropriate dialogues and cultural scenarios, while in STEM education it can create unique problem sets that build progressively on student understanding. These applications demonstrate how generative AI can enhance rather than replace traditional teaching methods, serving as a powerful tool for expanding educational possibilities.

Looking forward, the evolution of generative AI in education points towards increasingly sophisticated applications that blend AI with cognitive science. Current research explores how these systems can better understand and respond to student emotional states, learning preferences, and cognitive development stages. The goal is to create more empathetic and adaptive learning experiences that maintain the human element of education while leveraging AI's capabilities to provide unprecedented levels of personalisation and support. This development trajectory suggests a future where AI becomes an integral part of the educational ecosystem, enhancing both teaching and learning processes while preserving the essential role of human educators in guiding student development.

The following is an example of how an author uses generative AI in their teaching delivery. Here generative AI is not used simply as a recall tool, but rather various properties of generative AI models are used to create a practical understanding of a context.

Sample assignment: Exploring deceptive design with generative AI

Context: Deceptive design can be summarised as tricks built into websites and apps which manipulate the user into something they don't want to do (Brignull, 2011) such as discouraging them from cancelling a subscription through a long series of "are you sure?" checks (Mills et al., 2023)

Globally, regulators are moving to protect consumers in this space (Mills, 2024). The EU's Digital Services Act, the Federal Trade Commission's "Click-to-Cancel" rule and the UK's Digital Markets, Competition and Consumers Act 2024, all prevent firms from using deceptive design to manipulate customers.

Regulators will be auditing websites and apps to identify deceptive design, and firms will conduct internal audits to identify issues which may come to regulatory attention. Being able to conduct a behavioural audit to identify deceptive design is a key employability skill.

Generative AI can be a useful tool in completing a behavioural audit (e.g. Mills & Whittle, 2023). In particular, as persona generation capability allows an auditor to "see" the website or app from multiple perspectives, such as those of vulnerable customers.

For this assignment you will complete a behavioural audit of a firm's process of your choice.

You will need to:

1. Provide a detailed rationale of why generative AI can provide insight into your audit.
2. Demonstrate that your model can identify deceptive design.
3. Conduct behavioural audits for at least 3 relevant personas.

Notes: Part 1 requires you to investigate the potential of generative AI for behavioural science, noting the limitations. Part 2 requires you to build various deceptive design templates and record the results of a genAI model's interpretation of them (with your choice of genAI model). Part 3 requires you to submit your audit prompt pathways and responses. You will use these to build the various process maps and pathway plots required for the audit report.

This assignment requires students to use generative AI as a tool, akin to a skilled professional, rather than as a simple knowledge generator. Students are required to demonstrate they understand how a model works, how it can be validated, and then use it for an assessed process rather than a simple output. They are then required to use that output to demonstrate that they can use generative AI in a novel and highly skilled way. From the outset it requires the application of core behavioural economics knowledge, with better audits relying on the application of behavioural principles into model development. Additionally, the assignment is designed so that all the student's choices around the context the generative AI model will explore demonstrates their core behavioural knowledge.

Managing uncertainty

This section illustrates the uncertainty surrounding AI: uncertainty in its impact, its capabilities, and what may come next. Institutions will be justifiably uncertain over which aspects of the knowledge base imparted to students should be protected as "human only" – that is, which parts of a subject or field do students need to know without AI support. This may include what the student needs to truly understand to be the senior partner co-working with or supervising an AI tool. The Resolution Foundation views the impact of AI as radically uncertain (Bell et al., 2023), meaning that we cannot assign a likelihood of AI's impact, nor can we realistically envisage what it might be. Academics and institutions are unable to make effective proactive decisions in an environment of such uncertainty and frequently fall into the "resilience trap" (discussed in the final section) when best trying to manage this. In this report, we use a common general definition of AI and apply its principles to learning and teaching in order to provide a manageable (i.e. more certain) view of AI in learning and teaching: "machines are intelligent to the extent their actions are expected to achieve their objectives" (Russell, 2021). This definition allows us to address a key fact about AI: we are now in a time where, in an educational setting, we can reliably expect an AI system to meet the objectives we give it. Acknowledging this provides a basis for action in higher education. We propose the following as a definition of universities and AI:

"Artificial Intelligence has reached the point of practicality in education. That is, in many educational scenarios, we expect the Artificial Intelligence to reliably produce an output relevant to the objective it is given. This output will not necessarily be perfect but will frequently be of an acceptable minimum standard."

Taking this as a starting point, this report first provides an overview of current applications of AI in education before investigating the impact of AI on learning and teaching models, as well as beginning to explore the practical challenges of AI adoption.

The report concludes with a speculative investigation of the future of AI and education, as well as some broad recommendations.

> Artificial Intelligence has reached the point of practicality in education. That is, in many educational scenarios, we expect the Artificial Intelligence to reliably produce an output relevant to the objective it is given. This output will not necessarily be perfect but will frequently be of an acceptable minimum standard.

Current AI applications in education

Although generative AI has advanced with remarkable speed, any consideration of how it is used in teaching and learning needs to be framed in two ways. First, technology has a history of transforming classroom dynamics, and the university provision of education more broadly. As Ethan Mollick, an academic and commentator on AI, has noted, the introduction of the calculator could be seen as making us lazier, but should perhaps more accurately be considered as a strategic use of technology to take shortcuts (Colomé, 2024). He argues that we need to be similarly strategic with AI, especially given how open-ended, or generalised, it is as a technology. Other innovations – scheduling systems, software that aims to detect plagiarism, tools to improve spelling and grammar, video conferencing, and remote learning platforms – predate generative AI and have reshaped how teaching takes place, how students study, or the size of staff workload [3]. Although articles tend to describe these changes in terms of allowing "instructors to dispense their duties more effectively and efficiently" (Chen et al., 2020), the reality is probably more complex, with the management and administration of these tools taking up any free time that may have been created. In a similar vein, AI is likely to relieve some burdens while creating new ones.

Second, the speed of developments can simultaneously distract us from what is happening now, and result in our knowledge quickly becoming dated. In October 2024, a Professor of Learning with Digital Technology commented that

> If the AI community had ever asked teachers in higher education what we need from them, we would very definitely not have asked for a technology that makes it easy for students to deliver second-rate essays with a lot of factual errors (Laurillard, 2024).

Concerns about cheating notwithstanding (a theme we will return to), essays will not remain "second rate" and error riddled for long. Services such as Perplexity and ChatGPT Search already provide citations for claims [4], and the quality of AI models is expected to rapidly advance, at least in the near future (Azhar, 2024). At the same time, a preoccupation with

[3] Similarly, academic misconduct is not new – the introduction of the internet being perhaps the most notable example of a technology that has enabled students to cheat (Leo, 2023; Mollick, 2023)
[3] https://www.perplexity.ai/; https://openai.com/index/introducing-chatgpt-search/

where the technology is headed means we can lose sight of some of the novel applications of LLMs that are happening now, including by students.

"I think AI use will become the norm pretty soon and employers will expect us to be able to use it in a critical way. If we can do something much more efficiently, why wouldn't we?" – postgraduate student, University of Aberdeen (in Hack, 2023)

In the UK, where the authors of this report are based, the higher education sector has met AI with both excitement and concern (Freeman, 2024). Broadly speaking, however, the overall tone is one of trying to manage and control AI and minimise risks, rather than considering innovative ways of integrating it and improving teaching and learning (Illingworth, 2024). We would expect a similar pattern to emerge in other countries, including Ireland. Organisations such as Advance HE are supporting universities and encouraging institutions to look beyond a deficit model focused on academic misconduct, including work on AI and quality assurance, employability, and skills in an AI-dominated world of work. A first step for many institutions has been the release of guidance on the acceptable use of AI. Examples include guidelines and principles released by the Russell Group, which represents twenty-four universities (Russell Group, 2023). Jisc have also collected AI policies and guidance from UK universities to help share good practice (Barker, 2024b). A typical example of this is that of the University of the Highlands and Islands, which includes guidance for learning and teaching staff, support staff, and researchers to support the effective use of generative AI for activities associated with teaching and research, in ways that support and enhance practice. It includes approved generative AI tools, allows staff to decide how students use it for assessment, and gives examples of how staff themselves can use it appropriately (UHI, 2024a; 2024b).

AI is currently being used by some higher education staff in many ways. In the classroom, AI can support curriculum development, content personalisation, and pedagogical methods through interactive learning environments (ILEs) and intelligent tutoring systems. These tools can facilitate personalised learning experiences and enable more effective instruction and are explored further in the next section. Some administrative tasks have been streamlined through AI automation, particularly in assessment and feedback processes – also covered later in this report. Staff can use AI-powered platforms for grading, reviewing student work, and providing timely feedback. The technology can also enhance communication between teachers and students through web-based platforms and virtual environments. Beyond traditional teaching, AI can support educational delivery through various technologies including virtual reality, robotics, and 3D visualisation tools. These innovations help create more engaging learning experiences while allowing staff to track student performance and refine their teaching methods (Chassignol et al., 2018; Chen et al., 2020).

"It seems lecturers and the university are confused with how to proceed with it. They don't know if they should promote or criticise it, which is slightly confusing" – student (in

To shift to the perspective of students, a few interesting surveys give us an insight into the use of generative AI for learning and coursework, as well as the pace of change. A survey of 1,000 undergraduate students in the US found 43 percent of respondents used ChatGPT or similar AI applications, 30 percent used AI for the majority of their assignment, and 17 percent acknowledged having submitted AI-generated coursework without edits (Welding, 2023). (This survey was conducted in March 2023, when ChatGPT was less than six months old).

Case study A

Maynooth University, Ireland

The "GenAI Guidelines for Teaching, Learning, and Assessment" is a collaborative project between the Centre for Teaching and Learning (CTL) and Student Skills and Success at Maynooth University. This initiative aims to support staff and students in using generative AI effectively in their educational practices (Maynooth University, 2024).

The project is developing comprehensive guidelines to help both staff and students navigate the use of generative AI in educational contexts. A key component is the expansion of the CTL GenAI Resource Hub, which will provide additional insight guides for staff on various related topics. To raise awareness about the guidelines and discuss issues related to assessment and academic integrity, the university plans to facilitate professional learning sessions featuring expert speakers.

Student-focused resources will be developed alongside staff guidelines to ensure learners understand how to use generative AI tools responsibly. The project will also collect and share case studies that demonstrate effective uses of generative AI across Maynooth University and other higher education institutions. Funded by the Strategic Alignment of Teaching and Learning Enhancement Fund in Higher Education (SATLE), this initiative ran from June to December 2024.

Maynooth University's proactive approach to integrating generative AI into higher education creates supportive frameworks that benefit both educators and students while addressing ethical considerations in academic AI use. Notably, the "Maynooth University student working guidelines for GenAI" were co-created by students from various academic disciplines and levels, working together with colleagues from Student Skills and Success, the CTL, and Critical Skills (Maynooth University, 2024).

In November 2023, the Higher Education Policy Institute (HEPI) surveyed 1,250 UK undergraduates on their attitudes towards generative AI tools (Freeman, 2024). The analysis mirrored the US study and found that generative AI has become widely adopted in higher education, with 53% of students incorporating it into their academic work. The technology primarily serves as an AI private tutor, with 36% of students using it to better understand concepts. Students also commonly use it for generating ideas and article summarisation. While most students use AI as a learning aid, some incorporate it directly into their assignments. Thirteen per cent of students used AI-generated text in their assessments, though most edit this content before submission; 5% of students submit AI-generated content without any personal modifications.

A third survey, this time by QS and conducted in late 2023, polled over 1,600 students and academics globally (Bizzozero, 2024). It found that AI is significantly influencing the academic landscape, with over one-third of students reporting that it affected their choice of course, university, or career path. Adoption is widespread – nearly 80% of students use tools like ChatGPT for their studies, and 92% of academics have some familiarity with generative AI. Computer science has now overtaken engineering as the top field of study for students, highlighting the growing attraction of AI-focused disciplines. While students value AI for improving their productivity and concentration, there are widespread concerns by academics about potential over-reliance. Academic institutions are particularly mindful of privacy issues and maintaining academic integrity, which has led to calls for clear ethical frameworks to guide AI integration into education (these are covered in more detail in the last section). A fourth survey in July 2024 illustrates the pace of change, finding that 86% of students globally regularly use AI in their studies (a doubling of the US sample from a year earlier) (Digital Education Council, 2024).

Table 1 summarises the findings of these surveys.

Table 1: Surveys on student AI use

Date	Organisation	Scope of survey	Key findings
March 2023	BestColleges (Welding, 2023)	1,000 US undergraduate students	• 43% used ChatGPT or similar AI applications • 30% used AI for majority of assignments • 17% submitted AI-generated work without edits
November 2023	HEPI (Freeman, 2024)	1,250 UK undergraduate students	• 53% incorporated AI into academic work • 36% used AI to better understand concepts • 13% used AI-generated text in assessments • 5% submitted AI content without modifications
Late 2023	QS (Bizzozero, 2024)	1,600+ students and academics globally	• Nearly 80% of students use tools like ChatGPT • 92% of academics are familiar with generative AI • Over 1/3 of students report AI influenced their choice of course/university/career
July 2024	Digital Education Council (2024)	3,800 students from 16 countries	• 86% regularly use AI in studies • 54% use AI on a weekly basis • 67% use AI to search for information • 58% feel they lack sufficient AI knowledge/skills • 48% don't feel prepared for an AI-enabled workplace • 80% say their university's AI integration doesn't meet expectations • 60% are worried about fairness of AI evaluations • Only 5% are fully aware of AI guidelines

It is important to add a final note before we explore examples of practice and implementation of AI within universities. Although there is burgeoning literature on AI and teaching and learning, it is nascent and there are plenty of areas in need of further research, which will hopefully be covered in the coming months and years. Bahroun et al. (2023) synthesised 207 papers on generative AI in education and recognised the need for longitudinal studies to track the long-term effects of generative AI integration in education, and deeper analysis of the sociocultural factors that influence the attitudes and perceptions of students, such as those covered above. Some academics are – quite rightly – asking whether we should be using AI for, for example, formative assessment in science, and how this could be done in an equitable manner, deepening and personalising rather than minimising teachers' interactions with students (Li et al., 2023) Others would say that the train has already left the station, as AI is already widely employed in formative assessment and outpacing the research community (Zhai & Nehm, 2023). The pace of change is indeed rapid, but this does not preclude universities from proactively shaping the use of this technology. The following sections dive more deeply into the implementation and impact of AI in teaching and learning.

> I have an open book exam in January, which I know I will struggle to pass. I have run typical questions through ChatGPT, and I know that would give me a good enough answer to pass – but what would be the point? I could get through and struggle on until the end of my degree but what then? ChatGPT and other AI tools will be on the market doing a better job than I could, so what am I being trained for?" – current undergraduate computer science student (in Hack, 2023)

Examples of AI Implementation

A few case studies of how universities are integrating AI into teaching and learning are scattered throughout this report. A common thread seen in the most effective examples is that they encourage students to learn about AI – to begin to understand and probe the mechanics and limitations of the technology – while they use it for a specific course-related task.

Students will be working with and alongside AI for the rest of their lives. Teaching students both how to use AI effectively and how to critically evaluate AI-generated content prepares them for this world. Most students expect to continue using AI after graduation, with 73% saying they expect to incorporate it into their future work to some extent (Freeman 2024). Translation emerged as the most anticipated application of AI, with 38% of students expecting to use AI for this purpose. Close behind, 37% expect to use AI to improve their writing, while 33% plan to use it for text summarisation. Fewer than one in five students anticipate using AI to generate text from scratch (Freeman, 2024). However, as models have developed in capability, this may change; AI hallucinations have decreased in volume and

severity, and correspondingly the increase in the humanisation (or indeed personalisation) of written content makes AI-generated text a more viable option (Lee & Low, 2024).

More research is needed on how the integration of AI impacts student learning outcomes (Wu & Yu, 2024), and a progressive yet critical approach to integrating AI encourages openness and transparency. This means institutions should actively promote open discussions about AI usage rather than treating it as a taboo subject.

This means institutions should actively promote open discussions about AI usage rather than treating it as a taboo subject

The statistics on AI usage among students suggest that, in the absence of these open discussions, students will use AI regardless and simply not disclose it. As Mollick puts it, there is a natural tendency to hide AI usage, including within workplaces: "if I use AI to do work, others will think I'm brilliant … you don't want people to know that you're not actually that brilliant … they're afraid that you'll realise their job is redundant, or that they'll be asked to do more work" (Colomé, 2024).

Some organisations are taking a progressive approach to disclosing AI usage. Jisc, a body supporting digital transformation in universities, has published an interesting and open breakdown of how they themselves use AI [5]. Some lecturers, such as Professor Francesc Pujol of the University of Navarra in Spain, are embracing AI fully: "in the semester's first class, I said: 'Welcome to this course. There are quite a few assignments, and you will have to use ChatGPT in all of them'" (Colomé, 2023). Part of the reasoning behind this thinking is to level the AI playing field between teacher and student – when students know a teacher is well-versed in AI tools, they will use AI to learn better and to support analytical thinking and critical reflection rather than to merely get faster answers. "Instead of using it to beat the system, the students use ChatGPT constructively," Pujol adds.

Some approaches are systemic, focused on overall approaches to learning. Leo (2023) suggests students receive mandatory training in each year of study to "foster fundamental values of academic integrity" and reduce misconduct (presumably through AI literacy training), and for staff to be encouraged to integrate experiential learning, giving the example of legal training simulators used at the University of Central Lancashire for dispute resolution and hostile negotiations. As learning becomes more experiential, staff and students will benefit from closer relationships between universities and employers. See case study A (presented earlier) for an example of an assessment which uses the properties of generative AI rather than simply its knowledge recall functionality.

[5] https://nationalcentreforai.jiscinvolve.org/wp/2024/11/21/navigating-the-terms-and-conditions-of-generative-ai/

Another – perhaps the most telling – means of understanding the use of AI in teaching and learning is to observe how it is used by students. We will examine how AI is used overall in more detail later, but commentators generally observe how student engagement with AI often evolves from initial attempts at automation to more sophisticated uses, such as employing AI as a sounding board for ideas and a tool for refining existing work. In turn, student experiences reveal both the potential and limitations of AI tools, and they have noted particular challenges with citation accuracy, technical terminology, and development of evidence-based arguments (Darwin et al., 2024; Hack, 2023). Newport (2024) provides a helpful summary of how one particular student experimented with ChatGPT and, in a process of iteration, developed a deeper understanding of the technology. "He was not outsourcing his exam to ChatGPT; he rarely made use of the new text or revisions that the chatbot provided," Newport summarised.

> "He also didn't seem to be streamlining or speeding up his writing process … if I had been Chris's professor, I would have wanted him to disclose his use of the tool, but I don't think I would have considered it cheating."

Institutional support frameworks will become increasingly important; the work of progressive professors may hint at what is possible, but as we learn more about the impact of AI, the risks and limitations, and how it can improve student outcomes, this needs to be captured and carefully rolled out. It also needs to be adjusted across disciplines and, as the following sections will discuss, adapted for different modes of learning and different student groups. AI interaction is already transforming student learning processes and revealing new patterns of knowledge construction, but the limited evidence so far suggests that, while AI serves as a collaborative tool rather than a replacement for traditional learning methods, we ignore AI to the detriment of teaching and learning.

The Impact of AI on Different Learning Modes

The integration of AI across different modes of learning – face-to-face, distance learning, and online provision – is reshaping how education is delivered and experienced. These changes are particularly significant given the shifts in educational delivery precipitated by the Covid-19 pandemic and the subsequent emergence of hybrid learning environments.

In face-to-face settings, AI can enhance traditional instruction by providing simulation-based learning experiences. These applications use various technologies, including virtual reality, to demonstrate concepts and provide students with practical, experiential learning opportunities (Chen et al., 2020). This integration of AI-powered simulations can help bridge the gap between theoretical understanding and practical application, making

complex concepts more accessible and engaging.

For group and remote learning scenarios, LLMs such as ChatGPT can act as valuable facilitators of collaborative learning. These models can structure group discussions, provide real-time feedback, and offer personalised guidance to enhance student participation and engagement. In collaborative writing projects, AI assists multiple students working together by offering style suggestions and editing support. For research activities, AI can help identify research questions and efficiently distribute topics among team members (Kasneci et al., 2023).

The role of AI in remote learning has become particularly significant (e.g. Sevnarayan & Potter, 2024). Chatbots have demonstrated effectiveness in maintaining student motivation and engagement, partly through their ability to promote self-directed learning and autonomy through tailored questions and assessments (Chen et al., 2020). This capability is especially valuable in addressing the challenges of student engagement that were exacerbated during the shift to remote learning during the Covid-19 pandemic. However, it is important to note that access to these tools is not universal – students without reliable internet access or computer equipment may be excluded from these benefits (this is covered in more detail in the next section) (Illingworth, 2023).

The Open University has conducted research funded by the NCFE Assessment Innovation Fund on the implications of AI for remote assessment, examining seventeen different assessment types ranging from traditional formats to innovative approaches like infographics and role-plays (NCFE, 2024). Figure 1 visualises some of the common tells of AI-generated work, as identified by academic staff. Their findings revealed that role-play scenarios requiring practical application of learning and reflective tasks that draw on concrete examples were most resistant to AI generation. However, the study also found that AI-generated responses achieved passing marks in nearly all assessments (fifty-eight out of fifty-nine). While marker training improved AI detection rates, it also increased false positives, leading the project team to recommend prioritising assessment redesign over detection training. Assessment is explored in detail in the following section.

Figure 1: Common hallmarks of generative AI use, as identified by academics. Adapted from Hardie et al. (2024), with additional tells added by authors in the centre

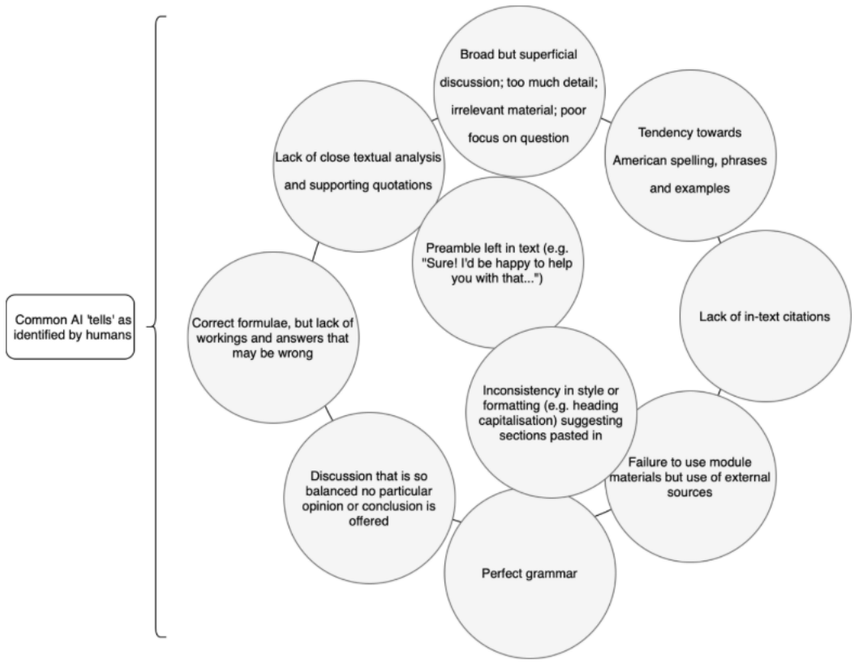

AI is also changing collaborative learning environments through intelligent tutoring systems and virtual learning assistants. These tools facilitate group discussions, provide guidance, and foster collaboration among students. The result is a dynamic learning environment that promotes active participation, critical thinking, and problem-solving skills, more closely reflecting real-world scenarios (Kaledio et al., 2024).

This transformation of different modes of learning through AI comes with implementation challenges, particularly around costs and infrastructure. The adoption of advanced technologies like Augmented and Virtual Reality requires significant investment (Varsik & Vosberg, 2024), and institutions need to carefully consider the balance between innovation, accessibility, and pedagogical effectiveness (while resisting the temptation to chase the shiniest new tech toys). As these technologies continue to evolve, ensuring equitable access while maintaining educational quality across different learning modes will be critical – as will understanding what works and why.

How AI Impacts Different Student Groups

While AI tools offer exciting possibilities for enhancing education, their integration raises important questions about equity and inclusion. There is currently a lack of research on the implications of AI tools for educational equity, including a scarcity of data and robust evaluations (Varsik & Vosberg, 2024). As with any major technological advancement or pedagogical innovation, the integration of generative AI tools in higher education presents both opportunities and challenges for different student groups.

Recent research suggests that a digital divide is beginning to emerge in AI usage patterns across different student demographics. This divide manifests in several ways, although it tends to exacerbate existing inequalities. In the UK, nearly three-fifths of students from the most privileged backgrounds (58%) use generative AI for assessments, compared with half (51%) from the least privileged backgrounds. There are also notable differences across ethnic groups, with Asian students indicating higher AI usage across a wider range of purposes, while white and Black students report lower usage rates (Freeman, 2024).

Gender and age also play significant roles in shaping AI engagement patterns (Stöhr et al., 2024). The overall proportion of male (57%) to female students (49 percent) using AI differs, and there are also clear differences as to how they use the technology. Male students tend to use AI text generators more frequently and are more likely to use AI for data analysis and coding, while female students more often employ AI for editing, translating text, and speech transcription. Age-related differences are also particularly striking when it comes to attitudes towards AI use: students aged twenty-five or over show significantly less tolerance for using AI for any purpose, and are more than twice as likely to consider no uses of AI acceptable (Freeman, 2024).

The implications of these disparities extend beyond current usage patterns to future opportunities. Students from more privileged backgrounds (i.e. POLAR quintile 5) are notably more likely to say they expect to use generative AI in their future careers (23%) compared to those from less privileged backgrounds (e.g. 15% for quintile 1 [6]). Perhaps more concerning is the fact that there is evidence suggesting that more privileged students may be better at using AI without detection, as they express less confidence in their institution's ability to spot AI use (Freeman, 2024).

To address these emerging inequities, Illingworth (2023) recommends that institutions take proactive steps. Universities should ensure equal access to digital resources such

[6]The POLAR (Participation of Local Areas) classification system divides the UK regions into five quintiles based on the proportion of eighteen- and nineteen-year-olds entering higher education, with quintile 1 representing the lowest participation rates and quintile 5 the highest. See https://www.officeforstudents.org.uk/data-and-analysis/young-participation-by-area/about-polar-and-adult-he/.

as generative AI by providing:

- Laptops students can borrow, and free or discounted internet access
- Partnerships with community organisations that can provide internet access or computer lending services
- Technology training to help students develop necessary skills and use AI tools effectively.

AI tools also present unique opportunities for enhancing accessibility for students with disabilities. LLMs can be combined with speech-to-text or text-to-speech solutions to assist people with visual impairment. These technologies can support inclusive learning strategies through adaptive writing, translation, and content highlighting in various formats. However, it is important that such implementations be guided by professionals such as speech therapists, educators, and specialists who can adapt the technology to specific disability-related needs (Kasneci et al., 2023). At the same time, students with disabilities may face unique challenges when using AI tools. For example, visually impaired students may face difficulties using chatbot interfaces, and support must be made available to help students with different needs access and effectively use new technologies (Illingworth, 2023).

The digital divide in AI access and usage threatens to exacerbate existing inequalities in educational opportunities and outcomes (Illingworth, 2024). To prevent this divide from widening, institutions should consider providing AI tools for students who cannot afford them where these tools have been identified as benefiting learning (Freeman, 2024). This

As AI becomes further integrated into higher education, universities need to maintain a careful focus on equity and inclusion. This means not viewing AI as a quick fix for educational challenges (Varsik & Vosberg, 2024), but rather as a tool that requires thoughtful implementation with consideration for the diverse needs of different student groups. Regular monitoring of usage patterns and outcomes across different demographic groups is critical for ensuring that AI integration enhances rather than hinders educational equity.

As AI becomes further integrated into higher education, universities need to maintain a careful focus on equity and inclusion

AI and Learning and Teaching Models

While concerns about academic integrity often dominate discussions of AI in education, LLMs such as ChatGPT, Claude, and other generative AI tools offer significant potential for enhancing student engagement and deepening learning outcomes. When used strategically rather than as a shortcut, AI can serve as a powerful tool for developing critical thinking and research capabilities. This section explores the relationship between AI and teaching and learning in detail, including how it can reshape the student–teacher relationship, the impact of personalisation, and what this means for the value of a university education.

For students, LLMs can support various aspects of the research and writing process, helping to develop critical thinking and problem-solving skills. These models can generate summaries and outlines of texts, enabling students to quickly grasp main points and organise their thoughts more effectively. They can assist in developing students' research skills by providing information and resources on particular topics, while also highlighting unexplored aspects and current research directions (Kasneci et al., 2023). Concerns around hallucinations can be partially addressed by models which seek to provide evidence, such as ChatGPT search, Perplexity, and specific academic tools such as Elicit [7].

However, there is an important distinction between using AI as a crutch versus using it as a 'co-intelligence' tool. As research has shown, students who rely on AI as a crutch risk diminishing the extent their learning by avoiding independent thinking. The key lies in using AI to enhance capabilities while keeping students actively engaged in the learning process (Mollick, 2024). We cover this further as part of a broader discussion on academic integrity in a later section, "Practical challenges, barriers, and ethics".

What tasks do students assign to chatbots? Recent research has identified three key practices that characterise productive interaction between students and generative AI interfaces: requesting, evaluating, and refining (see Table 2). Requesting involves making various types of queries to generate different forms of output. Evaluating requires students to actively analyse and judge the AI-generated content. Refining encompasses actions taken to improve the system's output to achieve specific learning goals (Pigg, 2024).

Table 2: A framework of interactions with generative AI in research writing

Theme category	Specific theme

[7] ChatGPT search and Perplexity are covered in the previous section "Current AI applications in education". Elicit is an example of an AI research assistant: https://elicit.com/.

Requesting	
Secondary research help	request for summary or key points of a secondary source request for references, citation, or reading ideas request for synthesis of multiple bodies of literature request for summary of available literature on a topic request to analyse the strengths and weaknesses of an aspect of a secondary source request to convert to a particular citation format request for simplified explanation of complex idea
Pre-writing help	request for outline request for scaffold, "recipe", or genre description request for quizzes, questions, or other ways to build content knowledge
Research report elements	request for abstract request for conclusion request for title request for implications or applications request for genre or research report section request for study contributions request to convert table into prose description request for code (i.e. LaTeX, Microsoft Visio)
Editing and revision help	request to shorten text request to improve writing for concision, active voice, etc. request to elaborate text request to proofread
Outreach help	request for outreach genre request to rewrite findings for nonexpert audience
Methodological or research design help	request for methodological advice request for analysis of a research gap request for a relevant theoretical approach request for potential new research questions related to prior readings request for new study objectives request for converting to tabular data request to analyse data for trends
Evaluating	for reference quality (or existence) for meeting genre expectations for coherence of meaning for alignment with writer's goals for specificity for plagiarism score for AI detection score for audience appropriateness
Refining	requesting more creative, detailed, or improved output requesting completion or continuation providing an example text providing content from the research project requesting output that better meets genre expectations repurposing prior output as new prompt

Reproduced from Pigg (2024)

For teaching staff, AI can support more efficient and effective research and writing tasks. At a basic level, it can help identify and correct typographical errors. More substantially, it can highlight potential grammatical inconsistencies and suggest personalised improvement strategies. The technology can also assist in identifying opportunities for topic-specific style improvement and generate summaries of challenging texts, enabling teachers to better support student understanding (Kasneci et al., 2023).

Research suggests that the most effective approaches to integrating AI in learning emphasise embodied practice and recognise the individual nature of how students engage with writing technologies. Importantly, studies have identified that experts in writing and research often demonstrate different uses of and attitudes towards these technologies compared to students who are still developing their research and writing skills (Pigg, 2024). This highlights the importance of providing clear guidance and support frameworks for AI integration in teaching and learning contexts.

AI and the Student–Teacher Relationship

The integration of AI into higher education will fundamentally reshape the relationship between students and teachers, particularly in the realm of assessment and feedback. This transformation presents both opportunities and challenges for the traditional teaching dynamic, with AI potentially serving as both a bridge and a barrier between educators and learners.

One of the most significant impacts of AI has been on assessment processes and feedback mechanisms. LLMs can semi-automate the process of grading student work by highlighting potential strengths and weaknesses in essays, research papers, and other writing assignments. This capability can help teachers identify areas where students are struggling, contributing to more accurate assessments of student learning development and challenges (Kasneci et al., 2023)

However, institutional responses to the growing availability of AI have varied considerably. A recent UK survey found that only 9% of students report their institution has changed assessment approaches "significantly", while 24% indicate approaches have remained unchanged. Nearly a quarter of students (23%) are uncertain about their institution's approach to AI in assessment, suggesting a need for clearer communication about institutional policies and practices (Freeman, 2024).

The automation of administrative tasks through AI aims to increase staff efficiency and save time, although a comprehensive survey would be required to see whether the marketing spiels match reality. Intelligent tutoring systems and AI-powered platforms provide functionalities for various administrative tasks, including grading and feedback provision (Chen et al., 2020). While this automation can enhance efficiency, some academics argue that we should question the premise that assessment needs to be faster, suggesting instead that formative assessment is integral to teaching rather than an activity to be outsourced (Li et al., 2023).

Research indicates that students are generally open to AI-supported assessment, with over 80% accepting it under certain conditions. However, this is conditional: students prefer continued human involvement in the assessment process, emphasising the importance of maintaining a "human-in-the-loop" approach (Braun et al., 2023). This aligns with calls from researchers that AI should be used to promote and support the interaction of teachers in the "sensemaking" of students, rather than merely codifying teacher decision-making (Li et al., 2023).

The timing of feedback also emerges as a critical consideration. Traditional teacher feedback on submitted assignments often comes days or weeks after submission, potentially diminishing its relevance for students. AI-generated feedback could provide immediate responses to draft work, though this is most effective when based on student-generated content rather than AI-generated submissions (Laurillard, 2024). Some innovative approaches involve using AI within peer review processes, where:

- Students provide feedback to AI-generated peer work using assessment rubrics
- Students prompt AI to give feedback on their own drafts
- Students revise their work based on both AI feedback and their experience of reviewing AI-generated work
- AI provides meta-feedback on the quality of student reviews (Laurillard, 2024)

However, the integration of AI fundamentally challenges the traditional teacher–student relationship in several ways: it can threaten the intellectual authority of lecturers, it disrupts conventional teaching spaces, it affects the emotional relationship between

teachers and students, and it necessitates strengthening ethical considerations in the teaching relationship (Cao, 2024). More pragmatically, one may wonder whether the use of AI tools by staff in assessments legitimises the deeper use of AI tools by students, encouraging an escalation in reliance on generative AI. escalation in reliance on generative AI.

A survey of 389 students and 36 staff across two universities revealed moderate usage of generative AI and consensus about which types of assessments are most impacted. While staff generally prefer adapted assessments that encourage critical thinking while assuming AI will be used, student reactions are mixed, partly due to concerns about losing creativity (Smolansky et al., 2023).

As AI continues to evolve, the focus should shift towards using it to deepen and personalise the interactions of teachers with students, rather than simply automating existing processes (Li et al., 2023). This might involve exploring various human-in-the-loop approaches where teachers and automated systems work collaboratively to enhance the assessment process in ways that combine the strengths of both (Li et al., 2023).

Evidence suggests that while AI can enhance certain aspects of the student–teacher relationship, particularly around feedback and assessment, it should not be seen as – and there is widespread determination that it not be seen as – a replacement for human interaction in education. Instead, the technology should be used to create more opportunities for meaningful engagement between teachers and students, while being mindful of the potential risks to the art and practice of teaching.

The Personalisation of Teaching and Learning

The ability of generative AI to offer personalised learning experiences holds significant promise for creating more inclusive and adaptable educational environments. This technology can enhance the student experience in several key ways, from providing real-time feedback to promoting deeper engagement with course materials (Illingworth, 2024). As the underlying learning models continue to improve and become more integrated with existing educational systems, they offer increasingly sophisticated ways to adapt to individual student needs (Hack, 2023).

One notable application is in language support and accessibility. Sophisticated translation technologies can adapt course content in real-time, enabling students from diverse linguistic backgrounds to interact with materials in their native or preferred language (Illingworth, 2024). This capability has particular relevance for international students and those studying in multilingual environments, potentially removing significant barriers to learning.

For teaching staff, LLMs offer valuable tools for creating personalised learning

experiences. These systems can analyse student writing and responses, providing tailored feedback and suggesting materials that align with specific learning needs. This support can significantly reduce the time required to create personalised materials and feedback, allowing teachers to focus on other aspects of teaching, such as developing engaging and interactive lessons (Kasneci et al., 2023).

The technology also shows promise in curriculum development and lesson planning. AI can assist in creating course syllabi and generating questions that encourage participation across different knowledge and ability levels. Importantly, these tools can help generate targeted practice problems and quizzes, allowing students to master the material at their own pace (Kasneci et al., 2023). This customisation of curriculum and content in line with the needs, abilities, and capabilities of learners represents a significant advancement in educational technology (Hack & Knight, 2023).

Perhaps one of the most immediate benefits for students is the availability of 24/7 support through AI-powered chatbots. These systems provide students with opportunities to ask questions they might otherwise feel uncomfortable raising in person (Hack, 2023). As these platforms analyse increasing amounts of data, they can identify patterns and offer personalised recommendations, potentially enhancing student engagement and motivation (Kaledio et al., 2024), especially if these insights are then married with existing educational infrastructure (the implications and risks of this data harvesting are covered in earlier sections).

AI's ability to provide immediate and constructive feedback represents another significant advantage. AI-powered assessment systems can evaluate assignments, quizzes, and examinations promptly, helping students understand their strengths and weaknesses in real-time (Kaledio et al., 2024). This rapid feedback loop facilitates self-reflection and enables students to make timely improvements in their learning approach.

However, it is worth noting that, while AI can create engaging learning experiences (Chen et al., 2020), questions remain about the long-term impact of such personalisation on learning outcomes. Further research is needed to understand how these tools affect different student groups and learning contexts, and to ensure that personalisation enhances rather than replaces meaningful human interaction in education.

Universities and AI

As AI becomes increasingly sophisticated, universities will need to adapt their teaching methods and learning environments to ensure that they continue to provide distinctive value in an AI-enhanced world – the (somewhat simplistic) question of "why go to university when you can be tutored by ChatGPT?" is unlikely to disappear anytime soon.

The emergence of AI tutoring and personalised learning systems will prompt a shift in how universities approach education. Rather than viewing AI as a threat to traditional teaching methods, progressive institutions will recognise it as an opportunity to transform their pedagogical approaches. Research shows current classroom dynamics often result in participation from only a small group of students, with others struggling to remain engaged, highlighting the need for more inclusive teaching methods. Lectures – a cornerstone of university teaching for centuries – may no longer represent the most effective learning format (Mollick in Colomé, 2024). Instead, there is growing evidence supporting a shift towards active learning environments where students engage directly with material and collaborate with their peers, and where the future of assessment may lie in combining traditional methods with AI-enhanced approaches, such as real-world case studies, internships, and project presentations.

This transformation extends beyond simply incorporating AI tools into existing frameworks. The "homework apocalypse" – referring to the disruption of assessment methods caused by the fact that many traditional homework assignments can easily be completed by AI – provides an opportunity to rethink how learning is structured and assessed (Colomé, 2024). Universities are beginning to explore models where classroom time is dedicated to problem-solving and collaborative work, while AI systems provide personalised support outside of class hours to fill knowledge gaps and provide additional guidance. And, as explored at the start of this section, "load sharing" with AI tools may become commonplace for complex assignments – with students being taught how to do so effectively. This still requires students to think carefully and write clearly, suggesting AI integration may reduce barriers without compromising learning outcomes (Newport, 2024).

Research into AI-powered learning technologies reveals both their potential and limitations. These systems can provide automated scaffolding for learning activities, including personalised reminders, real-time feedback on writing, and study recommendations. However, there remain questions about their impact on student agency and self-regulated learning. A study involving 1,625 students across multiple courses found that, while AI assistance was effective in improving immediate performance, students tended to rely on rather than learn from this assistance. When AI support was removed, students struggled to maintain the same level of performance, even when provided with self-monitoring tools (Darvishi et al., 2024).

> "Students tend to rely on AI assistance rather than learn from it."
>
> Darvishi et al. (2024)

This finding has important implications for how universities should integrate AI support systems. Rather than viewing AI as a replacement for traditional learning methods, institutions need to develop hybrid approaches that combine AI assistance with strategies to develop students' independent learning capabilities. The research suggests that the most effective approaches may be those that use AI to enhance rather than as a substitute for human-led teaching and learning (Darvishi et al., 2024).

As the previous section discussed, the challenge of maintaining academic integrity in an AI-enabled world is pushing universities to reconsider their assessment strategies. There is an emerging consensus that formative assessment should remain integral to teaching rather than being fully outsourced to AI systems. And, rather than engaging in what some describe as a "failed war on student academic misconduct", institutions are exploring more innovative approaches to evaluation (Leo, 2023). These can include:

- Limiting traditional essay-style assessments to one per semester per module
- Increasing emphasis on real-world case studies and practical applications
- Incorporating internships and project presentations into formal assessment
- Using time-limited, problem-based examinations
- Recognising excellence in extracurricular activities as part of assessment (Leo, 2023)

The role of AI in education extends beyond traditional degree programmes. LLMs and other generative AI tools are increasingly supporting research, writing, and problem-solving tasks across various domains (Kasneci et al., 2023). This has particular relevance for professional training and continuing education, where AI can provide domain-specific language skills and support for specific professional competencies.

Universities are well-positioned to leverage these capabilities in developing microcredentials and lifelong learning programmes. The integration of AI tools can help make continuing education more accessible and personalised, allowing universities to support learners throughout their careers rather than just during their initial degree studies.

The effective integration of AI into higher education requires careful consideration of ethical implications and strong interdisciplinary collaboration, covered further in the following section. Universities need to develop frameworks that balance the benefits of AI-enhanced learning with concerns about privacy, equity, and the development of critical thinking skills. This includes ensuring that AI tools are used in ways that support rather than undermine the development of deep subject knowledge and analytical capabilities.

The value of university education in an AI-enhanced world will increasingly lie in its ability to combine technological innovation with personalised learning experiences, while avoiding the technocratic temptation to focus on efficiency gains and cost savings. Although

AI can provide powerful tools for support, the unique value of university education remains the human element, and its ability to develop not just knowledge and skills but also critical thinking, ethical judgement, and the capacity for independent learning. As universities continue to adapt to technological change, maintaining this balance will be crucial to ensuring that they continue to provide meaningful value to students and society.

Case study B

University of Sheffield, UK

This case study was shared by a Senior Lecturer from the School of Biosciences at the University of Sheffield, with Jisc (Barker, 2024a). They describe the use of AI to create interactive sessions, and to encourage experimentation and deeper understanding.

"I teach students to use GenAI (Google Gemini) in my teaching sessions by providing them with customisable prompts that they can cut-n-paste from a shared GoogleDoc to save them typing the prompt in. I always include specific tasks that require students to critically evaluate the outputs."

In an example, they use Gemini to support student teams in brainstorming potential solutions to the UN sustainable development goals and critique the outcomes, with the following prompts:

"Act as a [CEO of a medium-sized biotech company specialising in enzymes]. How could your company contribute to tackling the UN SDGs?" Then, "Tell me more about [idea 2]." Then, "Act as a harsh critic. Criticise [idea 2]. Why might it fail?" And then, "Where can I find data on progress made on [SDG2]?"

For a second example, students use role-play to explore ethical issues around genetic screening.

'You are [Mary, a homemaker from Texas with four dyslexic sons who are struggling at school]. What is your opinion on combining IVF with preimplantation genetic testing for traits like [dyslexia]? Why?"

The lecturer adds that AI is used to explore and test ideas: "The outputs are not directly used in assessed work, so no refinements of outputs are required. I encourage students to continue on with the chat to explore thoughts and ideas – GenAI works best when used iteratively … The students really enjoy these sessions, and the large sessions can get pretty loud! It also gives students confidence when talking about challenging topics (e.g. ethics in the biosciences). These are deliberately designed as a safe, low-stakes environment where students can interact with GenAI without the fear of getting it 'wrong.'"

Practical Challenges, Barriers, and Ethics

Access, Resources, and Implementation Barriers

The integration of AI in higher education presents complex challenges around equity, access, and resource allocation. For many staff and students, the introduction of generative AI has generated considerable confusion and uncertainty, often overshadowing its potential benefits. This phenomenon has been described as the "Skynet effect" – a fear-driven response to new AI technologies that has fostered a cautious and frequently sceptical attitude within higher education (Illingworth, 2024). [8]

A more practical and immediate concern is the potential for AI to exacerbate existing digital inequities. Pre-pandemic data from the UK illustrates the scope of this challenge: an Ofcom survey found that 9% of households with children lacked access to a laptop, desktop, or tablet, with 4% having only smartphone access. This digital divide was particularly pronounced in lower-income households, where 21% of families whose main earner held a semi-skilled or unskilled occupation had no access to devices for their children's education (Illingworth, 2023) [9]. The societal implications of the "shift to online" driven by the Covid-19 pandemic are still being examined (Mills et al., 2021), but in the meantime regional and national governments have increased efforts to provide universal digital access (Peruzzo & Allan, 2024).

While tools like ChatGPT offer free basic access, the newest and premium features often reside behind paywalls. As AI technologies evolve and new generative tools emerge, universities face growing pressure to consider how best to respond to the proliferation of subscription-based tools. Institutions must develop strategies to ensure fair access, enabling both students and staff to utilise the AI resources necessary for their teaching and learning practices (Russell Group, 2023). This challenge extends beyond simple access to tools – it encompasses the need for comprehensive support systems and training.

The implementation of AI in education demands the development of new competencies and literacies. Teachers and learners need skills to understand both the technology's capabilities and its limitations, including the occasional "brittleness" of such systems, given we do not always fully understand how they work. A clear institutional strategy and pedagogical approach are essential, with particular emphasis on critical thinking and fact-checking strategies (Kasneci et al., 2023). This likely necessitates significant investment in professional development and support infrastructure.

[8]This does not preclude the need for serious examination of the broader impact AI will have on society, and the need for strong policies to mitigate the potential of serious harms it could cause (Russell, 2021; Suleyman & Bhaskar, 2023).
[9]We did not find comparable data for Ireland. However, just under 94% of households in Ireland had internet access in 2023 (Statista, 2024).

The cost implications of AI implementation can be substantial. There are three areas of concern: first, there are direct (albeit difficult to quantify) costs for both students and teachers. Students may become overly dependent on AI tools, which could potentially reduce engagement and learning outcomes, while teachers face an increased workload in training, using, or countering the effects of AI systems. Second, there is the monetary and time costs of working with AI education providers and consultants. Third, there is an opportunity cost, as the current focus on generative AI may reduce funding for research and development of more fundamental and potentially useful AI models, such as classic rule-based programmes for concept development and interactive learning (Laurillard, 2024).

Additional challenges include the risk of inaccurate or misinterpreted information, as data within generative AI tools is gathered from diverse sources, including poorly referenced or incorrect materials. There are also ongoing concerns about plagiarism and copyright infringement, as AI tools are trained on information developed by others (Russell Group, 2023). Furthermore, there is a risk that students could become over-reliant on AI technologies, po-

To address these challenges, staff should actively raise awareness of digital equity issues. This includes helping students understand that access and privileges in digital settings are not uniformly distributed. Through class discussions and activities that encourage critical thinking about digital equity and social justice, students can develop empathy and awareness of digital poverty (Illingworth, 2023). This approach aligns with broader efforts to ensure that AI adoption in education supports, rather than undermines, equitable and inclusive learning environments (Varsik & Vosberg, 2024).

To address these challenges, staff should actively raise awareness of digital equity issues.

The implementation of AI in education requires striking a careful balance between leveraging technological benefits and addressing ethical considerations. A comprehensive approach must consider privacy concerns, enhance cultural responsiveness, manage techno-ableism, and provide continuing professional development opportunities. Another point, perhaps sometimes overlooked, is that institutions must also maintain educational integrity amid growing influence of commercial entities in the higher education sector (Varsik & Vosberg, 2024). Large technology companies wield enormous power and have interests that do not always necessarily align with the users of their tools, including universities (Unwin, 2017). This necessitates ongoing research into the implications of AI tools for equity and inclusion, ensuring that technological advancement serves to reduce rather than amplify

existing educational disparities. The following sections dive into the topics of data privacy, AI bias, and algorithmic transparency, as well as the ethical use of AI, in more detail.

Data Privacy, Security, and AI Harms

As universities begin to integrate generative AI tools into teaching and learning, concerns have emerged regarding the balance between innovation on one hand and data privacy, security, and potential harms on the other. These issues require careful consideration and proactive management to ensure responsible, effective, and sustainable implementation.

A primary concern focuses on the handling and protection of student data. As AI systems rely on collecting and analysing personal information to provide personalised learning experiences, safeguarding student privacy has become essential for educators and policymakers (Kaledio et al., 2024). Further research is needed into how to develop robust data protection measures to ensure responsible and secure handling of student information (Bahroun et al., 2023).

The commercialisation of educational AI tools presents additional challenges. As private companies become more involved in developing and providing these tools, there is growing concern that commercial interests might overshadow educational objectives. This trend raises critical questions about the primary focus of educational tools, with the risk that profit motives could take precedence over educational outcomes. The access these commercial entities have to vast amounts of sensitive student data amplifies the potential for misuse, suggesting a need for policymakers and stakeholders to explore options for maintaining educational integrity amidst growing commercial influence (Varsik & Vosberg, 2024). Given the advance of open-source models (such as Meta's LLAMA) that are broadly competitive with the leading commercial models, we may see entrepreneurial universities fine-tune and deploy LLMs themselves in the future [10].

Privacy considerations extend to the everyday use of AI tools in teaching and assessment. For instance, the widespread practice of academics "testing" ChatGPT with exam questions and coursework inadvertently supplies more training data to these systems (Hack & Knight, 2023) – unless they have opted out of data sharing. Similarly, the use of AI detection tools raises both legal and ethical questions. Uploading student work to such tools may breach data protection legislation or institutional policies (Hack & Knight, 2023). Moreover, depending on whether a generative AI tool is designed to learn directly from its users' inputs or not, there are inherent risks to privacy and intellectual property associated with the

[10]This would not be a straightforward endeavour and would still likely require paying for cloud compute from the likes of Amazon Web Services, Microsoft, or Google. There is also disagreement about how open models such as LLAMA (https://www.llama.com/) truly are (Maffulli, 2023).

information that students and staff may enter (Russell Group, 2023).

The effectiveness and reliability of AI detection tools also present challenges. False positive rates for AI detection tools reportedly vary between 2% and 20%, and these tools have yet to undergo independent verification. Additionally, users can evade detection relatively easily by adjusting parameters within AI tools themselves, such as increasing randomness and deploying different writing styles. Given the intrinsic learning capabilities of AI systems, detection tools may always lag behind the technology they aim to monitor (Hack & Knight, 2023).

Misunderstandings also abound. A professor in the Agricultural Sciences Department at Texas A&M University initially failed their entire class, believing all students had submitted AI-generated work. The professor had attempted to verify this by asking ChatGPT itself if it had written the assignments. After it became clear that this detection method was not reliable, the professor reversed the failing grades (Newport, 2024).

These challenges suggest that higher education providers need to think beyond detection and consider fundamental questions about teaching methods and assessment approaches (covered later in the report). There is a growing recognition that authentic assessment should provide students with opportunities to demonstrate skills in working with AI responsibly. This then would necessitate developing students' ability to critique AI-generated content and understand its underlying algorithms and datasets, including potential biases

The implementation of AI tools in education also requires careful consideration of user agreements and data collection practices. Jisc have helpfully scrutinised user agreements of various generative AI tools, particularly regarding age restrictions and whether providers collect user inputs for training their models [11]. This examination extends across different types of AI services, including text, image, and audio tools, though the rapidly evolving nature of this technology necessitates regular review of such policies (Shepperd, 2024).

The implementation of AI tools in education also requires careful consideration of user agreements and data collection practices

AI Bias and Algorithmic Transparency

A significant challenge in the educational application of AI lies in understanding and addressing issues of bias (which is reproduced from training data) and transparency (of how

[11] See https://nationalcentreforai.jiscinvolve.org/wp/2024/10/31/navigating-the-terms-and-conditions-of-generative-ai/. The list is updated regularly.

LLMs work under the hood). Recent surveys indicate concerning gaps in the awareness students have of AI limitations: more than a third of students who have used generative AI (35%) report being unaware of how frequently it produces fabricated facts, statistics, or citations (i.e. hallucinations). This lack of awareness may suggest either that students have limited experience with the technology or, of perhaps greater concern, that they are not making efforts to verify AI-generated information before incorporating it into their work (Freeman, 2024).

The challenge of bias in AI systems presents particular concerns in educational contexts. AI models like ChatGPT learn from existing data and can consequently inherit and perpetuate societal biases present in their training data (Gichoya et al., 2023; Illingworth, 2024). There have been a few high-profile examples of "racist" LLMs as a result; tech companies have worked to contain this, but the issue persists (see, e.g., Blodgett & Talat, 2024). This has significant implications for curriculum development and delivery, and could potentially undermine efforts to create inclusive educational experiences. Without careful monitoring and calibration, these tools risk reinforcing existing stereotypes and overlooking important cultural and contextual nuances in educational content (Illingworth, 2024).

The implications of these biases may disproportionately affect certain student groups. There are particular concerns about how students from non-dominant cultural and linguistic backgrounds will be impacted (Li et al., 2023). Some researchers argue that AI systems can overemphasise certain cultural perspectives, in particular by taking "the culture, language, and representations of White people as the standard against which all answers ought to be seen, heard, and measured" (Cheuk, 2021). This bias is compounded by the fact that the cultural and linguistic practices of students from minority backgrounds may be underrepresented in the training datasets used to develop these AI systems (Yao et al., 2020).

A related challenge is the lack of transparency in how AI systems operate (sometimes described as "interpretability", i.e. how we understand or trace the reasoning behind specific outputs). The "black box" nature of many AI models makes it difficult to understand the reasoning behind their outputs (Kasneci et al., 2023), with even the heads of AI companies unable to fully explain how they work (Tangermann, 2024). This opacity presents challenges for both staff and students in evaluating the reliability and appropriateness of AI-generated content. Future research needs to focus on enhancing the transparency of AI models, as understanding how outputs are generated is crucial for building trust and acceptance in university settings (Bahroun et al., 2023).

To address these challenges, several areas require attention:

1. Monitoring and correcting biases: establishing robust processes for identifying and mitigating biases in generative AI models, in order to maintain curriculum

2. Enhancing transparency: developing methods to provide clear explanations of how AI-generated outputs are produced, making the technology more accountable and trustworthy (Bahroun et al., 2023). For now, this will likely involve acknowledging what we do not yet understand
3. Promoting equity: addressing the specific needs and perspectives of students from diverse backgrounds, ensuring AI tools do not perpetuate existing educational inequalities (Li et al., 2023). The human-in-the-loop processes described for assessment in the previous section could be useful here
4. Verifying processes: similarly, implementing systems to help students and staff verify AI-generated information and identify potential hallucinations or biases (Russell Group, 2023).

The processes for mitigating bias in training datasets are often not publicly available (Cheuk, 2021), which only highlights the need for greater scrutiny and transparency in how educational AI tools are developed and deployed. As these technologies become more embedded in educational settings, addressing challenges of bias and transparency is crucial for ensuring they serve all students effectively and equitably.

The Ethical use of AI in Academia

There are numerous ethical counterweights that accompany the considerably promising potential of AI in higher education, from threats to academic integrity to its significant environmental impact. We have covered a few of these in previous sections, and will summarise the landscape in this section. Recent surveys have suggested that, while institutions are making progress in establishing clear guidelines, significant challenges remain in ensuring ethical and equitable use of AI technologies. Most higher education institutions have begun developing policies around AI use, with varying degrees of success in communicating these to students. Research indicates that a majority of students (63%) believe their institution has a clear policy on AI use, and a minority (12%) find these policies unclear. Fourteen per cent of students report being unaware of how clear their institution's policy is, suggesting ongoing challenges in communication (Freeman, 2024).

Students' views on acceptable AI use are nuanced. While a majority consider it appropriate to use generative AI for explaining concepts (66%), suggesting research ideas (54%), and summarising articles (53%), only 3% believe it acceptable to use AI-generated text in assessments without editing (Freeman, 2024). This alignment between institutional policies and student attitudes provides a foundation for ethical implementation, though the focus must remain on promoting ethical usage and literacy, including educating both students

and staff about AI's capabilities, limitations, and broader societal implications (Illingworth, 2024).

The relationship between AI and academic integrity presents complex challenges that extend beyond simple questions of cheating. As Newport (2024) notes, "the binary question 'is it cheating?' hides the possibility that something new and inventive might be going on here". This complexity is particularly evident in specific disciplines – studies have shown how the fields of computer science and engineering education, for example, face concerns about the erosion of core competencies through over-reliance on AI tools (Bahroun et al., 2023).

Higher education institutions have long struggled with academic misconduct, and their approach to AI-assisted work follows similar patterns (Mollick, 2023). Many institutions have adopted AI writing detection software, but this approach faces several challenges. As generative AI models advance in producing human-seeming content, the reliability of detection tools becomes increasingly questionable (Leo, 2023). More concerning is the potential for these tools to "exacerbate systemic biases against non-native authors" by misidentifying their original work as AI-generated (Liang et al., 2023).

The fundamental issue may lie in how academic achievement is measured. As astrophysicist Neil deGrasse Tyson has observed, "we have created a system where we value grades more than learning". This focus on grades as "the currency of learning" creates incentives that can work against genuine educational outcomes (Tyson in Leo, 2023).

More broadly, the environmental impact of AI usage in education represents an often-overlooked ethical consideration. The energy requirements for AI systems are substantial – training GPT-3 (the precursor to the current GPT-4 family of models developed by OpenAI), for example, consumed 1.3 million kWh, equivalent to of the yearly energy consumption of sixteen people in the United States. The operational costs are equally significant: while a Google search consumes approximately 0.0003 kWh, a single ChatGPT prompt requires 0.1 kWh – 330 times more energy (Laurillard, 2024).

The environmental footprint of AI extends beyond energy consumption. Research indicates that global AI demand may require between 4.2 and 6.6 billion cubic meters of water withdrawal by 2027, exceeding the total annual water withdrawal of half of the UK. Even seemingly simple tasks like AI image generation can consume significant resources, with some models using approximately half a smartphone charge per image (Barker, 2024c). Amazon, Meta, and Microsoft all have plans to use nuclear power to supply their data centres, but they have encountered numerous roadblocks (Chant, 2024); a reliance on nuclear power is also unlikely to evade ethical scrutiny.

There are also concerns about how AI tools are developed. Some developers have outsourced data labelling to low-wage workers in poor conditions (Perrigo, 2023), raising questions about the ethical implications of using these tools in educational settings. Furthermore, while ethics codes are embedded in some AI tools, their implementation is not universal, and users may find it difficult to verify whether or not they are incorporated (Russell Group, 2023).

The path forward requires a more strategic approach than was taken with previous technological innovations. As Mollick (in Colomé, 2024) notes, drawing parallels with social media:

> "AI does so many things that we need to set guardrails on what we don't want to give up. It's a very weird, general-purpose technology, which means it will affect all kinds of things, and we'll have to adjust socially. We did a very bad job with the last major social adjustment, social media. This time we need to be more deliberate."

Students are seeking more support and tools from their institutions, with 30% agreeing that their institution should provide AI tools, while only 9% say they currently do so. However, satisfaction with current AI support remains low, with only 22% of students expressing satisfaction (Freeman, 2024). These statistics suggest a need for institutions to not only develop clear policies but also to provide practical support and resources for ethical AI use.

Case study C

Cornell University, United States

To inspire staff to be creative in the classroom, Cornell recognised the work of five faculty members using generative AI in early 2024 (Cornell University, 2024). The university created "How to Implement This in Your Class" summaries of each, which are available for download as PDFs on their website.

In the first example, students used AI image generators like DALL-E and Midjourney to create eighty images of "weird things", exploring how different prompts affected the outputs. They critically analysed the AI's limitations, biases, and image sources through class discussions. Students then selected one of their prompts and corresponding images to develop into a student-crafted 3D model, alternating between AI and traditional methods to understand the capabilities and constraints of the tools.

Second, students explored the integration of AI-generated texts into academic writing throughout the semester. The instructor, Tracy Carrick, approached AI tools like ChatGPT as an extension of existing writing aids such as QuillBot and Grammarly. By participating in AI-free zones and generative AI integration activities, students learned to use AI tools strategically and ethically to enhance their writing skills. Carrick told students: "figure out ways to let tools and resources do some of the heavy lifting for you … when used ethically, responsibly, and strategically, they can help you learn to be a stronger writer."

In the third example, students used generative AI to create high-quality images illustrating fibre-science innovations. Through weekly blogs, presentations, and a final poster project, they documented their prompts and refined their image generation skills. The instructor, Juan Hinestroza, noted that AI tools democratised the creation of high-quality visuals for group projects. "The ability to generate high-quality images via generative AI was indeed transformative, as it allowed anyone in the group to contribute to the imagery of the poster," he said.

"Having this kind of interaction with AI encouraged us to become more familiar with it, but also to challenge the information it provided." – Student

Fourth, students studying controversies in American politics used LLMs like ChatGPT to examine counterarguments and definitions of key concepts. Through guided steps, they iteratively worked with AI to co-produce written argumentation while maintaining their own analytical input. The process helped students engage critically with complex terms like "democracy" and "authoritarianism", combining AI and non-AI approaches to build critical thinking skills.

The fifth example takes a different approach. Instructors implemented a colour-coding strategy for written argumentation that helped circumvent potential AI use. Students colour-coded their claims and supporting evidence using matching colours, which helped them visualise their argument structure and identify gaps in their reasoning. This approach not only limited AI use but also provided unexpected benefits for both students and faculty in terms of the organisation of their writing and feedback efficiency.

The Future of AI in Learning and Teaching

As this report has explored, even the near future of artificial intelligence in education is particularly hard to predict. The further into the future we go – perhaps looking forward five to ten years – the murkier our view. Where institutions may land is dependent on a myriad of factors, outside the non-linear path of AI itself. Universities globally are being challenged

by changing geopolitics, domestic economics, competition from new types of providers, as well as demographic changes. Aside from uncertainty generated by increasing AI capability, AI adoption in educational settings requires considerable human interaction, and this in itself is unpredictable. Skills and capabilities may differ across and within institutions, and an institution's ability to transform as required is highly variable. The rapid increase in AI models specifically designed for education further impacts our ability to look ahead. Universities are being forced to run at far quicker timetables than their normal decision-making processes allow for.

AI progression frameworks can help us understand the potential evolution of AI. We may move from simple chatbots through to more complex reasoning systems and agentic AI. This is an AI that will independently complete tasks and make decisions. However, university policy primarily focuses only on the first rung of this ladder. Each stage presents different implications for education, however the timeline and likelihood of these transitions remains highly contested. Anticipated AI developments might never arrive. Universities, however, have the potential to develop policy that is more far reaching than one simply focused on chatbots or the first stages of this new AI.

We can put together a framework that helps us understand where we may be going. First, we need to question how (or whether) AI capabilities will expand beyond current limitations. Second, we need to query to what depth AI integrate into the educational process. Third, and only possible after considering 1 and 2, we should consider what governance mechanisms will effectively control educational AI systems? Intertwined with all this, we must consider how human learning patterns will adapt to these changes.

These questions interact in complex ways, their answers are highly variable, and they create a tapestry of possible futures that institutions must learn to navigate.

The near-term changes are likely to include significant shifts in teaching delivery and assessment practise. It is likely that university administration and operations will also adapt with AI. Universities must consider their infrastructure, policies, and staff and student experience while maintaining educational quality and accessibility. In the short term, we expect to see the emergence of new models for the student–teacher relationship as AI systems increasingly reinforce educational content and support students in their learning. This report considers there to be tremendous potential for AI to augment human capabilities in an educational setting, we caution against AI being used for simply substitutive means or in a cost-cutting way. New systems of learning and skill development are likely to emerge, requiring the very skills that institutions may be looking to outsource to AI. Institutions will have to consider how they validate knowledge and expertise in this era.

Educational institutions' interactions with AI in these initial stages are coinciding with a flurry of different advice, of evolving regulatory frameworks and new ethical guidelines. Universities must navigate these changes while maintaining their fundamental role in knowledge creation, dissemination, and education. Our institutions will change, they must not lose sight of their core functions or identities when they do so. Institutions will succeed when they embrace the uncertainty around AI and manage it proactively and with understanding, while maintaining clear principles for educational quality and student development in this new world.

Students are also facing the uncertainty: the uncertainty of what to study and how to study it. Institutions must help students beyond graduation. This AI-driven world may well be one in which graduates face requirements in rapid succession to continually upskill (Brown & Whittle, 2020).

However we navigate the future, one possibility is that this current AI cycle crashes, and then institutions which have reshaped themselves for a world that revolves around AI may find themselves having to rapidly shift back.

Successful integration of AI technologies depends on systematic evaluation and testing. However, academic institutions tend to run on a cycle where continual testing and change is particularly difficult. Institutions may benefit from starting with small-scale pilot programmes, then targeting specific educational or administrative challenges. These pilots may provide valuable insights into both the feasibility and practicality of implementation, while minimising institutional risks. The results from these experiments may inform broader implementation strategies and help identify potential challenges before significant resources are committed. An evidence-based approach will help institutions avoid common pitfalls such as investing in technologies based on speculative potential rather than their actual demonstrated value. Generative AI promises a lot to universities, and superficially it may seem like a panacea against low levels of funding and rising costs. However, rushed adoption (particularly when it is financially motivated – see case study D) of these technologies could be institutionally damaging.

The financial implications of AI adoption are significant, and, as educational institutions are generally publicly funded, this demands careful planning of resource allocation. A balanced investment approach includes adequate funding for pilot programmes, evaluation processes, and ongoing assessment. However, even these initial costs may be too much for some institutions.

Case study D

AI integration

A financial crisis is looming for UK universities (Richmond & Regan, 2024), artificial intelligence (AI) is emerging as both an opportunity and a temptation – one that, if mishandled, could deepen rather than alleviate the sector's problems.

AI holds significant promise for enhancing education. It can potentially streamline administrative processes, personalise learning, and support innovative teaching methods, as well as support staff and address workload issues. However, the current financial crisis facing UK universities threatens to distort these potential benefits. Faced with tightening budgets, universities may adopt AI reactively, prioritising cutting costs over meaningful improvement. Automation of tasks such as grading, student support, and even content delivery might reduce costs in the short term, but this risks compromising the quality and richness of the student experience.

Furthermore, rushed implementation often overlooks critical safeguards, including equity considerations, ethical concerns, and the ability of existing infrastructure to cope. The pressure to integrate AI quickly may lead to reliance on commercial tools without adequate scrutiny of their data practices, bias mitigation, or long-term implications. Staff may lack proper training in using AI, and students may experience uneven access to these technologies, exacerbating existing inequalities.

To avoid this pitfall, universities must approach AI integration strategically. AI should be seen not as a replacement for human expertise but as a tool to augment it, supporting innovation and improving student outcomes. This requires careful planning, transparent communication, and investments in training and infrastructure. If the sector succumbs to the allure of AI as a quick fix, it risks undermining its fundamental mission: to provide equitable, high-quality education and knowledge generation in an era of significant change.

As institutions develop AI policies, they must balance innovation with equity, ensuring that technological advancement doesn't exacerbate existing educational disparities. They should advocate for developing frameworks that promote responsible AI development, while protecting student privacy, academic integrity, and equal access to educational opportunities.

The rapid evolution of AI technologies demands a fundamental shift in how universities prepare students for future employment. Rather than focusing solely on technical skills,

institutions must develop interdisciplinary programmes that combine technological literacy with critical thinking, ethical reasoning, and adaptable problem-solving abilities. This preparation requires universities to commit to shaping wider ethical AI use through research, policy, advocacy, and educational leadership. As our institutions develop their AI policies, they must balance innovation with equity, ensuring that technological advancement doesn't exacerbate existing educational disparities. Institutions should focus on developing frameworks that promote responsible AI use while protecting student privacy, academic integrity, and equal access to educational opportunities. University policy may well be place-based, supporting regional initiatives and localised skill needs.

The rapid evolution of AI technologies requires a fundamental shift in how universities prepare students for future employment. Institutions may look towards developing interdisciplinary programmes that combine technological literacy with critical thinking, reasoning, and adaptable and transferable problem-solving abilities. Universities should ensure they keep pace with changing economic environments and any new skill requirements. It is a possibility in this rapidly changing world that higher education award structures, content, and approach may change mid-cycle.

Educational institutions must actively address any digital divide and ensure equitable access to AI resources. They will need to develop further infrastructure support systems, provide additional and increased levels of training, and ensure that AI implementation doesn't disadvantage students from under-resourced backgrounds.

The Next Step: Avoiding the Resilience Trap

The running theme through this report is of AI creating an environment of profound uncertainty for education. This uncertainty, where future outcomes and events are so difficult to foresee, is best described as radical. Decision-making in a state of such deep uncertainty is difficult.

In such circumstances, many institutions rely not then on anticipating and preparing for future outcomes but on building their resilience. They will do this by training staff in AI literacy, implementing skills training, and developing new policy.

This is useful – indeed additional investment in staff AI understanding and skills is to be welcomed – however, resilience is not an end point. With regards to AI, we term this "the resilience trap". This is where institutions attempt to build defences against an unknowable future rather than wholly integrate with and shape that future. AI may impact many aspects of our world; educational institutions should be leading voices in the debate rather than passive resistors.

Resilience – in the form of AI literacy, new policies around assessment and engagement, and increased awareness – is only the first step. The key to this transformation lies not in general preparedness, but in developing a more fundamental understanding of AI for learning and teaching. Universities must play an active role in defining how AI is to be used and the impacts it will have.

This report has proposed a simple definition of AI and education to help universities navigate this future. By adopting the perspective that AI has reached the point of practicality in educational settings, higher education institutions can begin to narrow the scope of what they need to consider and prepare for. This definition effectively removes several highly uncertain aspects of AI (such as artificial general intelligence) from immediate consideration, replacing them with a more practical focus on systems that are designed to achieve specific objectives.

This reframing shifts the conversation from nebulous uncertainties to more manageable risks. Instead of attempting to prepare for every conceivable future, institutions can focus on understanding and responding to specific systems with defined objectives. This approach brings several advantages:

1. It allows institutions to make more concrete decisions based on observable system capabilities and limitations
2. It creates clearer pathways for research and inquiry focused on specific aspects of AI systems
3. It enables more effective resource allocation by focusing on practical rather than theoretical concerns

The current institutional emphasis on resilience reflects a broader tendency to try to make decisions under uncertainty rather than working to reduce that uncertainty. While some degree of uncertainty will always exist with emerging technologies, the goal should be to minimise it by achieving a deep multidisciplinary understanding of AI.

Recommendations:

It is not possible for this report to provide a step-by-step checklist for an institution to successfully adapt to the changes brought about by recent developments in AI. For an institution to succeed here, it will likely be due to a combination of numerous factors, including developing a full understanding of the implications of AI in its own specific context.

This understanding will be different across the range of educational institutions, depending on existing strengths, local and national economies, and an institution's civic role.

If institutions develop resilience as the first step of a strategy, rather than relying simply on acquired resilience to fortify them against the rapid change of AI, they can reduce the uncertainty around AI, allowing a proactive strategy to be developed.

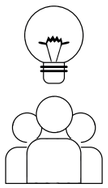

Upskill existing staff in AI understanding:
AI literacy is a part of this, but a deeper multidisciplinary approach is often required where staff learn not only what AI is, but what it means.

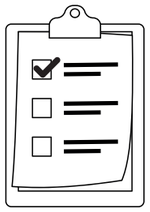

Assess institutional AI strengths and weaknesses:
Equipped with this understanding, assess institutional AI strengths and weaknesses, recognising that strengths in understanding this new mass market approach to AI may be found in unexpected places.

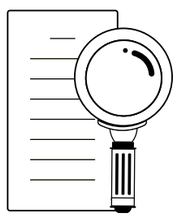

Work with local business and regional policy:
Work with local business and regional policy to develop a place-based understanding of AI – part of this role may be supporting business and policy in a temporal understanding of AI to avoid decision-making looking only to the short-term.

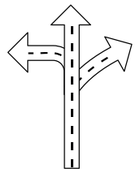

Challenge institutional norms:
Be prepared to challenge institutional norms around timescales, strategy, and policy. Rapid change may require equally rapid decisions. Institutions must be willing to have "difficult conversations".

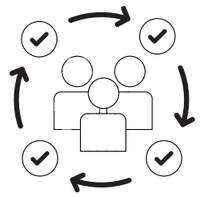

Test, Learn, and Adapt:
Develop a "Test, Learn, and Adapt" approach internally, trail new approaches with sophisticated evaluation procedures, and rapidly incorporate findings. Remember, what works today may not tomorrow.

Lastly, institutions must be prepared for failure. The radical uncertainty AI has ushered cannot be changed into a low-risk environment. Correct, informed decisions will sometimes go wrong; institutional policies need to be robust enough to accommodate the coming transition to a higher-risk environment.

Conclusion:

This report examines the impact of AI on models of learning and teaching. The potential impact of this AI boom goes beyond altering existing models and threatens to disrupt the entire educational model. Institutions will need to adapt, but the uncertainty around how AI will develop and how society and the economy will respond makes effective institutional response difficult. Wider factors around the role of educational institutions, domestic politics, and, for some institutions, low levels of funding create an environment where institutions are unlikely to experiment to determine an effective response. Financially, many institutions simply cannot afford to take the appropriate level of risk required to successfully adapt their existing models to AI. Given the uncertainty around AI, educational institutions are defaulting to low-risk resilience strategies, such as equipping their staff with a level of AI literacy sufficient to allow for some incorporation of AI into existing models.

This will not be enough. The rewards of AI are potentially large, but low-risk strategies will not realise them. This report advocates that institutions take risk appropriate to a high-reward outcome.

References

Azhar, A. (2024). What to expect when you're expecting GPT-5. Exponential view. https://www.exponentialview.co/p/expecting-gpt5

Bahroun, Z., Anane, C., Ahmed, V., & Zacca, A. (2023). Transforming education: A comprehensive review of generative artificial intelligence in educational settings through bibliometric and content analysis. Sustainability, 15(17), 12983.

Barker, C. (2024a). Innovative insights: AI success stories from the community. National Centre for AI. https://nationalcentreforai.jiscinvolve.org/wp/2024/04/15/innovative-insights-ai-success-stories-from-the-community-2/

Barker, C. (2024b). Navigating the future: Higher education policies and guidance on generative AI. National Centre for AI. http://nationalcentreforai.jiscinvolve.org/wp/2024/07/31/navigating-the-future-higher-education-policies-and-guidance-on-generative-ai/

Barker, C. (2024c). Artificial intelligence and the environment: Taking a responsible approach. National Centre for AI. https://nationalcentreforai.jiscinvolve.org/wp/2024/09/18/artificial-intelligence-and-the-environment-taking-a-responsible-approach/

Bell, T., Clark, T., Fry, E., Kelly, G., & Thwaites, G. (2023). Ending stagnation: A new economic strategy for Britain. Centre for Economic Performance.

Bizzozero, S. (2024). Generative AI in higher education: Navigating policy, ethics, and skills development. HEPI. https://www.hepi.ac.uk/2024/11/08/generative-ai-in-higher-education-navigating-policy-ethics-and-skills-development/

Braun, D., Rogetzer, P., Stoica, E., & Kurzhals, H. (2023). Students' Perspective on AI-Supported Assessment of Open-Ended Questions in Higher Education. 15th International Conference on Computer Supported Education, 73–79.

Brignull, H. (2011). Dark patterns: Deception vs. honesty in UI design. Interaction Design, Usability, 338, 2–4.

Blodgett, S. L., & Talat, Z. (2024). LLMs produce racist output when prompted in African American English. Nature, 633(8028), 40–1. https://doi.org/10.1038/d41586-024-02527-x

Brown, G., & Whittle, R. (2020). Algorithms, blockchain & cryptocurrency: Implications for the future of the workplace. Emerald Group Publishing.

Cao, M. (2024). Teacher–student relationships in the context of artificial intelligence. Advances in Educational Technology and Psychology, 8(4), 75–80. https://doi.org/10.23977/aetp.2024.080412

Chassignol, M., Khoroshavin, A., Klimova, A., & Bilyatdinova, A. (2018). Artificial intelligence trends in education: A narrative overview. Procedia Computer Science, 136, 16–24. https://doi.org/10.1016/j.procs.2018.08.233

Chant, T. D. (2024). Regulators deliver successive blows to Amazon and Meta's nuclear power ambitions. TechCrunch. https://techcrunch.com/2024/11/04/regulators-deliver-successive-blows-to-amazon-and-metas-nuclear-power-ambitions/

Chen, L., Chen, P., & Lin, Z. (2020). Artificial intelligence in education: A review. IEEE Access, 8, 75264–78.

Cheuk, T. (2021). Can AI be racist? Color-evasiveness in the application of machine learning to science assessments. Science Education, 105(5), 825–36. https://doi.org/10.1002/sce.21671

Colomé, J. P. (2023). AI in the classroom: "I require my students to use ChatGPT." El País English. https://english.elpais.com/science-tech/2023-02-24/ai-in-the-classroom-i-require-my-students-to-use-chatgpt.html

Colomé, J. P. (2024). Ethan Mollick, analyst: "Students who use AI as a crutch don't learn anything." El País English. https://english.elpais.com/technology/2024-10-03/ethan-mollick-analyst-students-who-use-ai-as-a-crutch-dont-learn-anything.html

Cornell University (2024). Teaching Innovation Case Studies. Center for Teaching Innovation. https://teaching.cornell.edu/teaching-spotlight/teaching-innovation-case-studies

Darvishi, A., Khosravi, H., Sadiq, S., Gašević, D., & Siemens, G. (2024). Impact of AI assistance on student agency. Computers & Education, 210, 104967.

Darwin, Rusdin, D., Mukminatien, N., Suryati, N., Laksmi, E. D., & Marzuki. (2024). Critical thinking in the AI era: An exploration of EFL students' perceptions, benefits, and limitations. Cogent Education, 11(1), 2290342.

Digital Education Council (2024). What Students Want: Key Results from DEC Global AI Student Survey 2024. https://www.digitaleducationcouncil.com/post/what-students-want-key-results-from-dec-global-ai-student-survey-2024

Freeman, J. (2024). Provide or Punish? Students' Views on Generative AI in Higher Education. HEPI Policy Note 51. Higher Education Policy Institute.

Gichoya, J. W., Thomas, K., Celi, L. A., Safdar, N., Banerjee, I., Banja, J. D.,…Purkayastha, S. (2023). AI pitfalls and what not to do: mitigating bias in AI. The British Journal of Radiology, 96(1150), 20230023.

Hack, K. (2023). "Excited, concerned and curious": Student perspectives on learning and working in the era of AI. https://www.advance-he.ac.uk/news-and-views/excited-concerned-and-curious-student-perspectives-learning-and-working-era-ai

Hack, K., & Knight, C. (2023). Higher education in the era of AI. https://www.advance-he.ac.uk/news-and-views/higher-education-era-ai.

Hardie, L., Lowe, J., Pride, M., Waugh, K., Hauck, M., Ryan, F., Maguire, C. (2024). Developing robust assessment in the light of Generative AI developments. The Open University. https://law-school.open.ac.uk/research/developing-robust-assessment-light-generative-ai-developments

Illingworth, S. (2023). If AI is to become a key tool in education, access has to be equal. The Conversation. http://theconversation.com/if-ai-is-to-become-a-key-tool-in-education-access-has-to-be-equal-204156

Illingworth, S. (2024). Generative AI in academia: Friend or foe? BERA. https://www.bera.ac.uk/blog/generative-ai-in-academia-friend-or-foe

Kaledio, P., Robert, A., & Frank, L. (2024). The impact of artificial intelligence on students' learning experience. Available at SSRN 4716747. http://dx.doi.org/10.2139/ssrn.4716747

Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F. … Hüllermeier, E. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. Learning and Individual Differences, 103, 102274.

Laurillard, D. (2024). How AI-generated feedback could assist student learning. HEPI. https://www.hepi.ac.uk/2024/10/22/how-ai-generated-feedback-could-assist-student-learning/.

Lee, C. C., & Low, M. Y. H. (2024). Using genAI in education: the case for critical thinking. Frontiers in Artificial Intelligence, 7, 1452131.

Leo, U. (2023). Generative AI should mark the end of a failed war on student academic misconduct. LSE Impact Blog. https://blogs.lse.ac.uk/impactofsocialsciences/2023/07/21/generative-ai-should-mark-the-end-of-a-failed-war-on-student-academic-misconduct/

Li, T., Reigh, E., He, P., & Adah Miller, E. (2023). Can we and should we use artificial intelligence for formative assessment in science. Journal of Research in Science Teaching, 60(6), 1385–9.

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. Patterns, 4(7). https://doi.org/10.1016/j.patter.2023.100779

Maffulli, S. (2023). Meta's LLaMa 2 license is not open source. Open Source Initiative. https://opensource.org/blog/metas-llama-2-license-is-not-open-source

Maynooth University. (2024). Maynooth University student working guidelines for GenAI. https://www.maynoothuniversity.ie/sites/default/files/assets/document/MU%20Student%20Working%20Guidelines%20for%20GenAI%202024.pdf

Mills, S. (2024). Deceptive choice architecture and behavioral audits: A principles-based approach. Regulation & Governance, 18(4), 1426–41

Mills, S., & Whittle, R. (2023). Detecting dark patterns using generative AI: Some preliminary results. SSRN. http://dx.doi.org/10.2139/ssrn.4614907

Mills, S., Whittle, R., & Brown, G. (2021). The implications of a crisis-driven societal shift to online consumption. In P. McCann, & T. Vorley (Eds.), Productivity and the pandemic: Challenges and insights from COVID-19 (pp. 88–103). Edward Elgar Publishing.

Mills, S., Whittle, R., Ahmed, R., Walsh, T., & Wessel, M. (2023). Dark patterns and sludge audits: an integrated approach. Behavioural Public Policy, 1–27. doi:10.1017/bpp.2023.24

Mollick, E. (2024). Co-Intelligence: Living and working with AI. Pengiun Random House.

NCFE. (2024). The Open University: Developing robust assessment in the light of Generative AI developments. Our Assessment Innovation Fund pilots: The Open University. https://www.ncfe.org.uk/help-shape-the-future-of-learning-and-assessment/aif-pilots/the-open-university/

Newport, C. (2024). What kind of writer Is ChatGPT? The New Yorker. https://www.newyorker.com/culture/annals-of-inquiry/what-kind-of-writer-is-chatgpt

Perrigo, B. (2023). Exclusive: The $2 per hour workers who made ChatGPT safer. TIME Magazine. https://time.com/6247678/openai-chatgpt-kenya-workers/

Peruzzo, F., & Allan, J. (2024). Rethinking inclusive (digital) education: Lessons from the pandemic to reconceptualise inclusion through convivial technologies. Learning, Media and Technology, 49(2), 244–58.

Pigg, S. (2024). Research writing with ChatGPT: A descriptive embodied practice framework. Computers and Composition, 71, 102830. https://doi.org/10.1016/j.compcom.2024.102830

Richmond, T., & Regan, E. (2024, July). Augar reviewed: Why post-18 education in England is still broken, and how to fix it. Tom-richmond.com. https://www.tom-richmond.com/wp-content/uploads/2024/07/EDSK-Augar-Reviewed.pdf

Russell, S. (2021). Human-compatible: AI and the problem of control. Penguin.

Russell Group. (2023). Russell Group principles on the use of generative AI tools in education. Russell Group. https://russellgroup.ac.uk/media/6137/rg_ai_principles-final. pdf

Sevnarayan, K., & Potter, M.-A. (2024). Generative artificial intelligence in distance education: Transformations, challenges, and impact on academic integrity and student voice. Journal of Applied Learning and Teaching, 7(1), 1–11.

Shepperd, P. (2024). Navigating the terms and conditions of generative AI. National Centre for AI. https://nationalcentreforai.jiscinvolve.org/wp/2024/10/31/navigating-the-terms-and-conditions-of-generative-ai/

Smolansky, A., Cram, A., Raduescu, C., Zeivots, S., Huber, E., & Kizilcec, R. F. (2023). Educator and student perspectives on the impact of generative AI on assessments in higher education. Proceedings of the Tenth ACM Conference on Learning @ Scale, 378–82. https://doi.org/10.1145/3573051.3596191

Statista. (2024). Ireland: Household internet access 2009–2023. Statista. https://www.statista.com/statistics/377692/household-internet-access-in-ireland/

Stöhr, C., Ou, A. W., & Malmström, H. (2024). 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, 100259.

Suleyman, M., & Bhaskar, M. (2023). The coming wave: Technology, power, and the twenty-first century's greatest dilemma. Crown.

Tangermann, V. (2024). Sam Altman admits that OpenAI doesn't actually understand how its AI works. Futurism. https://futurism.com/sam-altman-admits-openai-understand-ai

UHI. (2024a). Generative artificial intelligence (GAI) policy. UHI. https://www.uhi.ac.uk/en/t4-media/one-web/university/about-uhi/governance/policies-and-regulations/policies/generative-ai/generative-ai-policy.pdf

UHI. (2024b). UHI principles on the use of generative AI tools. UHI. https://www.uhi.ac.uk/en/t4-media/one-web/university/about-uhi/governance/policies-and-regulations/policies/generative-ai/uhi-principles-on-the-use-of-generative-ai-tools.pdf

Unwin, T. (2017). Reclaiming information and communication technologies for development. Oxford University Press.

Varsik, S., & Vosberg, L. (2024). The potential impact of Artificial Intelligence on equity and inclusion in education. OECD Artificial Intelligence Papers, No. 23, OECD Publishing.

Welding, L. (2023). Half of college students say using AI is cheating. BestColleges. https://www.bestcolleges.com/research/college-students-ai-tools-survey/

Wu, R., & Yu, Z. (2024). Do AI chatbots improve students learning outcomes? Evidence from a meta-analysis. British Journal of Educational Technology, 55(1), 10–33.

Yao, L., Cahill, A., & McCaffrey, D. (2020). The impact of training data quality on automated content scoring performance. Semantic Scholar. https://www.semanticscholar.org/paper/The-Impact-of-Training-Data-Quality-on-Automated-Yao-Cahill/8a7a0a47dd4fd989fb4431ffed2a1b3ec4b02556

Zhai, X., & Nehm, R. H. (2023). AI and formative assessment: The train has left the station. Journal of Research in Science Teaching, 60(6), 1390–8.

National Digital Leadership Network

ISBN 978-1-0683806-7-9