



New Horizons for Higher Education: Teaching and Learning with Generative AI

A report by Professor Mairéad Pratschke. Commissioned
by the N-TUTORR National Digital Leadership Network

New Horizons for Higher Education: Teaching and Learning with Generative AI



A report by Professor Mairéad Pratschke
Commissioned by the N-TUTORR National Digital Leadership Network

ISBN 978-1-0683806-3-1
DOI 10.5281/zenodo.15324394



Author biography

Mairéad Pratschke is Full Professor and SALC Chair in Digital Education in the Faculty of Humanities at the University of Manchester; Visiting Professor at the London School of Economics and Political Science (LSE) Data Science Institute (DSI); and Research Fellow and Advisory Board member at the USA's National Science Foundation-funded National AI Institute for Adult Education and Online Learning (AI-ALOE). Author of *Generative AI and Education* (2024), Mairéad has delivered keynote talks on generative AI and education in Ireland, the UK, Canada, the USA, Lebanon, the Netherlands, Portugal, Spain, South Africa and Singapore.

Author website: <https://maireadpratschke.com/>

Abstract

Generative AI (GAI) presents significant challenges and opportunities to the higher education sector, from the threat to academic integrity to the promise of personalised learning at scale. This report focuses on three key areas: the impact of GAI on current teaching and assessment practices, current applications and the shift to more learner-centred approaches; emerging GAI pedagogy, international best practices and early research findings on risks; and GAI and digital transformation, international regulation and the future skills agenda. The report distils the key findings into actionable recommendations and provides clear guidance and direction aimed at assisting institutional leaders shape policy and strategy for GAI integration into teaching practice in a way that aligns with their institutional values and strategic priorities.

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, while also facilitating associated 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.

A handwritten signature in black ink, appearing to read 'G. Ó Súilleabháin', written in a cursive style.

Dr Gearóid Ó Súilleabháin

Department of Technology Enhanced Learning (TEL)

Munster Technological University (MTU)

Table of Contents

Glossary	9
Executive Summary	11
Foreword	12
Innovation First	12
EU AI Investment	12
Reasoning Models, Deep Research, and Agentic Systems	13
Ireland’s Approach to AI in Education	13
Foreword References	13
Introduction: The Challenge	14
Artificial Intelligence	15
Generative AI	16
The Context	17
Representation And Values	17
Capabilities And Implications	19
Emerging Practice	22
AI Is Not Edtech	22
Efficiency And Effectiveness	23
Experiments With GAI	24
Research Findings	27
Pedagogical Alignment	29
AI Transformation	30
Future Careers	30
Collaboration Versus Automation	31
Agent Co-Workers	32
The New Hybrid	34
Guidance And Frameworks	35
Conclusions	37
Recommendations	37
References	40

Glossary

Below is a non-technical description of the key terms used in this report.

AGI - Artificial general intelligence (or superintelligence)

AI ecosystem - Network of specifically GAI-powered tools and technologies

AR - Augmented reality

Content window - The amount of data or text an AI model can process at a time

VR - Virtual reality

XR - Mixed reality

Digital ecosystem - Network of digital (Web 2.0) tools and technologies

Fine-tuning - Adjusting a pre-trained model to perform a specific task

Frontier models - Highly capable general purpose models

GPT - Generative pre-trained transformer

GAI - Generative artificial intelligence (also known as genai)

Hallucinations - Model output that is nonsensical or false

HEI - Higher education institutions

Inference - The process of generating output

IoT - Internet of things, the network of connected devices

LAM - Large action model

LLM - Large language model

LMM - Large multi-modal model

LPU - Language processing unit and a chip developed by Groq

LWM - Large world model

ML - Machine learning

Multimodal AI - AI systems that can integrate and process multiple data types (images, audio, text)

Neural networks - Type of machine-learning algorithm that uses a network of interconnected nodes to process data in a way that mimics the human brain, or a method in AI that teaches computers to process data in a way that is inspired by the human brain

Prompt engineering – The crafting of a prompt by a user to elicit the desired output

Prompt transformation – The rewriting of a prompt by an AI model to optimise the output

RAG – Retrieval augmented generation

Scaling laws – Describe how neural network performance changes as key factors are scaled up or down

Synthetic media – Video, image, text, or voice that has been generated using AI

Test-time computer – The computational resources used when an AI model generates responses to prompts; differs from training compute, which is used to create the model itself

Transformer – Architecture created in 2017 that enabled the scaling of LLMS

Web 2.0 (Social Web) – The second stage of development of the internet, characterised by the change from static webpages to dynamic or user-generated content and the growth of social media

Web 4.0 (Symbiotic Web) – The latest stage in development of the internet, characterised by technologies that facilitate closer interaction of humans and machines, in physical and virtual worlds

Executive Summary

Generative AI (GAI)¹ presents significant challenges and opportunities to the higher education sector, from the threat to academic integrity to the promise of personalised learning at scale. The national government and regulatory agencies will address the broader systemic challenges under the EU Artificial Intelligence Act, but implementation in higher education institutions (HEIs) requires careful planning and oversight. Higher education is a values-driven sector, where equity, diversity and inclusion – as well as social and environmental responsibility – inform and shape institutional strategy. This report focuses on three key areas: the impact of the rapidly improving capabilities of GAI and their implications for current teaching practices; emerging practice using GAI assistants, tutors, and instructors, as well as early research findings on their efficacy and risks; and how AI is transforming professional practice, and the implications this has for future careers. Key findings reveal the need for alignment between sectoral values and institutional strategy; pedagogical alignment of tools, design, and learning goals; and alignment of technical higher education programmes with industry to ensure that educational offerings prepare students for future careers. Together, these findings point to the need for action on all fronts, from collaboration with government and industry stakeholders to align strategy with regulation; to upskilling and training for AI literacy and competence; to curriculum review and redesign so that students are equipped with critical skills for the future. In short, institutions must ensure that GAI implementation for teaching aligns with sectoral values, institutional goals, industry needs, and student outcomes.

Since the completion of this report in early December 2024, several important developments have taken place. While the core analysis and argument remains relevant, readers should consider these and other ongoing advancements. For more please see the foreword on the following page.

¹Generative AI is often shortened to GAI or GenAI. GAI is the acronym used in this report.

Foreword

The rapid pace of AI development poses a challenge to publication timelines, which can lag several months behind the finalisation of a manuscript. Since this report was finalised in early December 2024, several important developments have taken place in the AI landscape, which have direct implications for higher education. While the key message in this report for higher education remains unchanged, the author would like to note the following updates that readers should be aware of as they read it:

Innovation First

President Donald Trump's return to power in the USA signalled a clear shift in the approach to AI development. Trump's Executive Order of January 2025 (White House, 2025) replaced former President Biden's, which had stressed the need for safe and responsible AI development, with a strong focus on innovation first. The escalating AI arms race between the USA and China is fuelling an international race to create sovereign AI (Pratschke, 2025), which is increasingly seen as critical to national security and global competition. The UK's AI Opportunities Plan (UK Government, 2025), published in mid-January, adopts a similarly innovation-friendly approach.

EU AI Investment

The release of China's DeepSeek in January made it clear that innovation in model training is as important as funding, and put open-source into the spotlight. At the AI Summit in Paris in February, EU President Ursula von der Leyen announced a €200 billion investment in AI, including €20 billion for AI gigafactories (European Commission, 2025). In February, a consortium of 20 leading European research universities and companies announced the launch of a new OpenEuroLLM project, which will see the creation of a family of multilingual, large language foundation models for commercial, industrial, and public services. The goal is to create transparent and compliant open-source models that will preserve cultural and linguistic diversity, as well as democratise access, enabling European companies to compete on the global market.

Reasoning Models, Deep Research, and Agentic Systems

Reasoning models can conduct advanced research, and agents can cross-reference and correct each other, decreasing the potential for error. Deep Research tools from Google and OpenAI can generate postgraduate theses and PhD dissertations. Reasoning models are now widely available and while OpenAI's subscription is prohibitively expensive, users can access DeepSeek using platforms like Perplexity. Anthropic released the first hybrid reasoning model in late February, simplifying the issue of choice for users.

Ireland's Approach to AI in Education

The AI Advisory Council released a series of reports in February, including an AI Advisory Council Advice Paper, entitled AI and Education, which includes its recommendations for government (AI Advisory Council, 2025). The paper underlines the need for consistent guidelines on AI use, AI literacy training for educators, equitable access to AI models and tools, and a national conversation between educators, unions, policy-makers, and the AI/edtech industry to decide on a consistent and stable approach to leveraging AI in education.

Foreword References

AI Advisory Council. (2025, February). *AI and education: AI Advisory Council advice paper*. AI Advisory Council.

European Commission. (2025, February). *EU announces €200 billion investment in AI, including AI gigafactories*. European Commission.

Pratschke, B. M. (2025, January 27). *The global pursuit of sovereign AI is becoming the 21st century's arms race*. LSE Business Review. <https://blogs.lse.ac.uk/businessreview/2025/01/27/the-global-pursuit-of-sovereign-ai-is-becoming-the-21st-centurys-arms-race/>

UK Government. (2025, January). *AI opportunities action plan*. UK Government. <https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan>

White House. (2025, January). *Removing barriers to American leadership in artificial intelligence*. The White House. <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>

Introduction: The Challenge

GAI is a general purpose technology, compared in terms of its potentially transformative impact on our society to the printing press, electricity, and the internet. It offers tremendous opportunities to broaden access to education, to increase student engagement, to personalise learning, and to expand HEIs beyond the physical campus.² GAI is an assistive technology, removing linguistic and geographic barriers and making learning more accessible. It expands the horizons of what is possible, presenting the opportunity for “intelligence augmentation” (Dede et al., 2021) and discovery far beyond what humans alone can achieve. But achieving this potential requires collaboration with technology in a way that is new for higher education. How HEIs adapt to this new hybrid form of learning will be the determining factor in whether GAI augments human intelligence or reduces it.

Balancing its potential opportunities with the challenges GAI presents is vital. GAI is a rapidly emerging technology that comes with risks that pose a direct challenge to the values that guide HEIs, and regulations governing its use are not keeping pace with technological developments. GAI is also a geopolitical issue, a matter of national sovereignty, global security, and intense international competition. The 2024 US election victory by Donald Trump will have important economic implications for Ireland, largely due to the number of American technology firms in the country. How the Irish government manages that relationship and balances it with its compliance with European Union regulations will have a direct impact on the options available to Irish HEIs. The EU Artificial Intelligence Act provides a framework, but the critical decisions with regard to GAI implementation in Irish HEIs will rest on institutional leaders.

Artificial intelligence (AI) transformation is the next phase in the journey towards Education 4.0, and technological HEIs will play a key role in preparing students with the skills and competencies they need for the AI-enabled workplace. HEIs measure success through key metrics including employability, sustainability, social responsibility, equity, inclusion, and student experience. In this

Artificial intelligence
(AI) transformation
is the next phase in
the journey towards
Education 4.0...

²Generative artificial intelligence is also explored in two further reports in the National Digital Leadership Network report series; see Weller (2024) and Whittle & Ranson (2024).

context, success requires aligning GAI implementation with sectoral values, institutional strategy, industry needs, and student outcomes. How different HEIs embrace AI will be a competitive differentiator, and technological HEIs are well placed for this transformation. Those that adapt will be the educational leaders of tomorrow, capable of addressing the rapidly shifting demands of industry and equipping students with the knowledge, skills, and competencies they need to thrive.

Artificial Intelligence

AI as a field of study has been around for many decades. The term was coined at Dartmouth College in 1956, where a group of educators gathered to explore the creation of intelligent machines that could learn and use human language. AI as a technology is also not new to higher education. AI has been used in HEIs for decades and for a variety of purposes. Symbolic AI is used to create academic advising tools, to formulate university policies, to design curricula, to create automated grading systems, and to assist in research. Predictive AI is used to identify patterns in student data to forecast performance, and in adaptive learning platforms to tailor the learning experience to the student (Sharples, 2019).

Mustafa Suleyman, co-founder of DeepMind and Inflection AI, and head of AI at Microsoft, titled his 2023 book *The Coming Wave* (Suleyman, 2023). In it, he argued that the world was not prepared for the wave of powerful new AI technologies about to hit it, and he identified “the containment problem” as the challenge for our age. Suleyman was far from the first to use this metaphor to describe the upheaval that follows the introduction of a disruptive new technology into society. Joseph Schumpeter used the phrase “creative destruction” to describe the dynamic of technological disruption that has driven the so-called waves of innovation in our modern era (McKraw, 2007), with each wave characterised by a set of technologies that was the era-defining technology of its age. As each wave of technological change washed in, the dynamic of creative destruction brought with it new ways of working and living but also destroyed something of what was there before.

The first three waves of innovation transformed Western societies, displacing manual workers and replacing them with machine automation and knowledge workers. The fourth wave (1950–90) marked the start of artificial intelligence as a topic of research, when the mission began to create intelligent machines. The fifth wave (1990–2020) was the digital age, defined by the networks made possible by the invention of the internet (1983) and Tim Berners-Lee’s World Wide Web (1989) that allowed users to navigate it. The 1990s saw the growth of machine learning techniques and improved

neural network architectures, which used the massive amounts of data (so-called “big data”) generated by Web 2.0’s social networks.

As the fifth wave began to ebb, the sixth gathered force. On 11 December 2015, OpenAI was founded by a group including Sam Altman, Greg Brockman, Reid Hoffman, Elon Musk, and others, with the stated intention of working towards creating safe artificial general intelligence (AGI) for the benefit of humanity. Conceptions of AGI vary considerably, but the key word in these discussions is “general”. While we have specialised AIs capable of performing tasks as well as, or better than, humans, we do not yet have AGI, that is, AI with general intelligence that surpasses that of humans across a wide range of cognitive tasks.

Generative AI

In 2017, a group of Google researchers introduced a new type of transformer architecture (Vaswani et al., 2017) that enabled the unprecedented scaling of large language models (LLMs) that would make the GAI boom possible. In 2018, OpenAI published a paper explaining what generative pretrained transformers (GPTs) were (Radford et al., 2018), and the original GPT-1 was launched to little fanfare. Updated versions of GPT were released at regular intervals thereafter, with parameter counts, representing increasingly large and more complex models capable of processing more data, going into the billions. But it was only in November 2022, with the release of ChatGPT-3.5, that GAI reached the public consciousness. Suddenly, Schumpeter’s dynamic of creative destruction was palpable, and Suleyman’s coming wave was visible to the naked eye.

GAI and its associated Web 4.0 technologies (i.e. the Internet of Things (IoT), augmented reality/virtual reality/mixed reality (AR/VR/XR), and robotics) are now dissolving the boundary between human and machine. From synthetic biology to synthetic media, our interaction with GAI is transforming academic practice across disciplines in HEIs and creating a new hybrid form of education

defined by human collaboration with GAI. The concept of knowledge itself has been disrupted by the development of machines that are intelligent but learn at a much faster rate than we could ever hope to. No longer the kingpins, we are now confronted with a world in which there is another intelligence, modelled on the neural networks in our human brains but made of chips and metal rather than cells and skin. HEIs need to pre-

...our interaction with
GAI is transforming
academic practice
across disciplines in
HEIs...

pare for a world where humans and machines work in symbiosis, where the educator is no longer the single voice of authority but rather, for better or for worse, one of many. The potential of this human–AI collaboration to augment human intelligence (Dede et al., 2021) is the promise of AI in education.

The Context

The release of ChatGPT-3.5 on 30 November 2022 prompted a huge wave of hype from industry and a corresponding wave of panic from higher education. Two years later, the sector has moved from that initial disruption to the process of integration, and the potential of GAI to support teaching and learning is increasingly clear. Innovative educators leading experiments have shown the utility of GAI to improve accessibility, engagement, support, and outcomes. As a result, the narrative in higher education is slowly shifting its focus from banning tools and policing students to the constructive integration of GAI to support teaching and learning.

The question about whether to integrate GAI is now moot: We live in a world where GAI is integrated into tools, platforms, and devices, from enterprise software to personal smartphones. For students, this means that expertise is readily available on any topic, at any time. But the difference in capabilities between free GAI tools and frontier models means that access to the most cutting-edge versions is critical to ensuring digital equity and positive outcomes for students. Open-source models raise concerns about data security and exposure to cyber threats, but the lack of secure access to safe GAI risks increasing the digital divide between those who can afford subscriptions and those who cannot.

Representation and Values

The lack of representation of minority cultural, ethnic, and linguistic groups in big data used to train GAI models means they do not represent the diversity of human populations or experiences. The Global South, Indigenous voices, and minority linguistic and cultural groups are poorly represented relative to the ideologies and attitudes of Western cultures in the Global North. Implicit biases in GAI models (Warr et al., 2023) have been shown to disproportionately generate views aligned to the white, male, Western, North American data they were trained on. The growing movement towards sovereign AI, itself a reflection of geopolitical competition and global inequity in representation, might rectify this imbalance by further increasing the diversity in model types. But, in the short-term, the risk of further embedding this inequity between groups, cultures, and geographies is very real. This is a serious issue for HEIs, where diversity and inclusion are core values. AI literacy is critically important – not simply to understand the

issues that come with it but to take corrective action, by designing and creating tools that represent the diversity of human populations, voices, and experiences.

GAI models produce biased output (Kannan, 2024), perpetuate patriarchy (Khan, 2024), and replicate the world views and ideologies of their creators (Buyl et al., 2024). Each has its own flavour, its own tone, its own strengths and weaknesses. Elon Musk's Grok reflects Musk's ideology, while Anthropic's Claude uses Constitutional AI trained to reflect and promote human values (Bai et al., 2022; Edwards, 2023a). Both reflect their training; for Grok, that means producing output that is relatively unpredictable compared to other models, while Claude's output reflects its training on documents including the United Nations' Universal Declaration of Human Rights. The choice of LLM is therefore not value neutral, which underscores the need for users to be more than passive consumers of GAI output, but rather critical, informed, active interrogators. Indeed, this is what it means to be AI literate.

Issues around representation and rights have important implications in the context of the current movement toward agentic AI, where AI systems do not simply generate output but act autonomously, making decisions and performing tasks with minimal human intervention. Nick Bostrom's famous paperclip maximiser thought experiment (Bostrom, 2003; Šmit et al., 2005) demonstrated the potential dangers of such a scenario when taken to its logical conclusion. Already, agentic systems are demonstrating they will prioritise their final goal over the means used to achieve it, sometimes resorting to collusion to do so (Fish et al., 2024). The in-built bias in GAI models, combined with ideological misalignment, suggests that poorly designed agentic systems that lack human oversight or control could pose a real threat.

The EU Artificial Intelligence Act requires AI systems to be designed, developed, and used in ways that respect equality and human dignity. Using systems misaligned with human values can leave vulnerable students exposed to output that is hateful (Clark, 2024), leading in some cases to tragic outcomes (Montgomery, 2024). Model training reduces the incidence of hallucinations, but a degree of unpredictability is built into generative models – it is how they work (Mishra, 2024) – so while different architectures and learning algorithms may mitigate these issues (Kalai & Vempala, 2024), some currently remain. Educators need to understand how GAI models work, so they can not only guide students in ethical use of GAI for learning, but also share practical techniques based on adversarial thinking to ensure that students' interactions are safe.

Capabilities and Implications

GAI has very different affordances and capabilities from the digital tools of the past. As model sizes and parameter counts have grown, this has enabled a dramatic increase in capabilities. GAI models are now fully multimodal: they process and generate not just text but also images, audio, and video. GAI is conversational: its natural language processing and understanding capabilities mean that users can interact with it using natural human language. GAI is generative, meaning that while it is trained on existing data, it creates something new. GAI's social and generative affordances, combined with its increased capabilities, have over the past two years transformed it from what was initially compared to a simple tool (the calculator) into a much more powerful form of intelligence that operates more like a presence – albeit a machine presence. These machine presences can get to “know” the user, remember their preferences, and adopt any role. This allows GAI to function as a tutor, a trainer, a guide, a mentor, a coach, an advisor – or any other role its human users define for it.

GAI models are becoming increasingly powerful, which is reflected in the size of their memory. A model's short-term memory is its context window, and context windows have expanded dramatically in the last year, allowing models to process more user data. In 2023, Anthropic's Claude first made headlines for its industry-leading context window of 100,000 tokens, which meant it could process 75,000 words, or a book the length of *The Great Gatsby* (Edwards, 2023b). Since then, the size of context windows has continued to grow, with Google's Gemini model debuting at over a million tokens. Google's NotebookLM, initially able to hold only 1,500 tokens, can now hold 1.5 million (Johnson, 2024). This enables it to process an entire catalogue of works. The potential use cases for NotebookLM (or its open-source equivalent, Open NotebookLM by HuggingFace) in higher education are endless, as interaction with multimodal content dramatically expands the dialogic space for learning.

In September 2024, NotebookLM added audio summaries, allowing users to upload source material that the model uses to generate a podcast hosted by two AI personalities. But the model does not simply summarise the content – rather, it dissects and interprets it, finds what is important in it, and shares those insights in a highly engaging conversation between the two AI hosts. This capability to connect the source content to examples and use cases beyond those appearing in the original source material is significant, in that it illustrates the power of this technology to take the original human content and use it to generate something new.

GAI is also changing the way we interact with the internet. Perplexity is a platform rather than an LLM, but it generates output the same way a chatbot does, bridging the

gap between search engine, research tool, and conversational agent. Perplexity takes the user query and browses the internet, then generates output from its search, including the web links it consulted as footnotes in its generated response. The popularity of this platform, which prompted some to speculate on the possible end of the current web search (Roose, 2024), has led OpenAI and Google to introduce a similar capability to their platforms, and an open-source version called Perplexica is also now freely available on GitHub.

Perplexity also generates an AI podcast called Discovery Daily, generated from its Discovery newsfeed, which is hosted by AIs Alex and Siena. This Discovery feed is rapidly integrating respected news sources via its growing list of partnerships with media companies. Generative search, and the integration of synthetic media into popular tools like YouTube and Spotify, is thereby already customising and personalising the online user experience. This overlap with social media is noteworthy because consumer technologies and trends tend to arrive later in higher education, often via edtech companies. TikTok and Spotify already offer short educational courses, and it is likely only a matter of time before those, too, are AI-generated and hosted. Given the history of MOOCs and microcredentials – both of which began well over a decade ago – and the EU’s current focus on the latter for encouraging lifelong learning, these GAI consumer trends and forces should not be written off as irrelevant to higher education but rather considered harbingers of the future.

Co-creation is now a core key feature of frontier models, meaning that technical expertise is no longer a barrier to creation. In November 2023, OpenAI released its CustomGPTs, which the company described as “proto-agents”, that can be built by non-experts in minutes. Since January 2024, the GPT Store has housed thousands of these custom tools, which are freely available for public use. Similarly, open-source platforms like HuggingFace allow users to create their own AI assistants for free, while Poe offers a large number of frontier and open-source base models which can be built on. Co-creation is significant because it is part of the shift in capabilities among GAI models from knowing to doing.

AI agents are widely seen as the next phase after AI assistants, which humans create and train and then have work autonomously with us and for us. Claude’s Artifacts allows users to easily create interactive digital assets, including web apps, from a simple text prompt. The model generates the code, which appears on the right side of the screen, so the user can work alongside it and tweak the model’s output. This capability allows users to collaborate with GAI, giving non-technical users the power to create. OpenAI’s GPT includes a similar feature, called Canvas, which the company described as enabling users to “collaborate” with the model. In a further

step along this path from co-creation to agentic AI, in November 2024, Anthropic introduced Computer Use to Claude, which enables the model to use the mouse to perform actions formerly performed by humans. Meanwhile, OpenAI has announced its own plans to launch an agentic system codenamed “Operator” in January 2025 (Robison, 2024).

Key to this shift towards agentic AI is the need for the human-in-the-loop, which refers not just to model training and fine-tuning techniques but to the need for AI literacy at the user end. This is critical to ensure that humans play an active part in this shift. The vast majority of GAI users are still at the experimental content generation phase, but many have moved into co-design of GAI bots and co-creation using GAI tools. As a result, we are now seeing the emergence of “intelligent communities” of academic researchers working in teams with AIs.

Since September, the emergence of a new branch of GAI models has led to discussions about the potential limits to current scaling techniques and possible future directions for development. The new so-called “reasoning” models (Open AI, 2024) are trained to “think” through a problem, an approach similar in effect to chain-of-thought prompting at the user end. OpenAI’s o1 model has shown impressive results to mathematics and logic tests, prompting mathematician Terence Tao to speculate that, within a couple of iterations, the o1 model might be as useful as a PhD-level graduate student. In an interview, Tao said, “I do envision a future where you do research through a conversation with a chatbot [but] I’m not super interested in duplicating the things that humans are already good at. It seems inefficient. I think at the frontier, we will always need humans and AI” (Wong, 2024). This collaboration between humans and AI is where the future lies for higher education.

The development in reasoning models is part of the trajectory in recent years towards the proliferation of model types. The GAI that has occupied higher education for the last two years is clearly significant, but there are other model architectures and many other model types, including LLMs but also small language models (SLMs), large multimodal models (LMMs), large action models (LAMs), language processing units (LPUs), and – frequently promoted by Yann LeCun and currently in development at Fei-Fei Li’s new venture, World Labs – large world models (LWMs). The aim for the latter is to develop spatial intelligence, which will allow AIs to understand and operate in the 3D world. These developments are all part of the path towards the final goal of AI scientists, which is to create AGI, or superintelligence. For users, the variety in model types means they choose between model type and size according to what best fits their purpose. SLMs require less energy than large frontier models and can be downloaded to use on mobile phones without Wi-Fi. Understanding these options allows users to

make informed decisions based on specific use cases that align with sectoral values and institutional strategy.

Emerging Practice

Multiple reports (Hirabayashi et al., 2024; Digital Education Council, 2024) on GAI use confirm it has been widely adopted by students as a tool for learning, as an assistive technology, an organisational tool, a motivational tool. Most students tend to use GAI for help with learning and homework rather than for cheating. Those with learning challenges stand to benefit the most, including neurodivergent students, those dealing with accessibility challenges, and those completing their studies in a second or third language. GAI functions as an assistive technology that allows them to participate more fully than they otherwise would have been able to, but students are often reluctant to share that they have been using it due to a fear of reprisal, as there is a wide range in policies on acceptable use of GAI and a lack of clear guidelines. The flexibility granted to lecturers in deciding how or whether to allow students to use GAI in their coursework has resulted in a lack of consistency. This, combined with the fear of being accused of academic misconduct, is a significant source of stress for students. Meanwhile, educators operate within a system resistant to change. The gap between students and academic staff in both frequency of GAI usage and familiarity with the technology was dramatically illustrated during the initial panic in HEIs over assessment, but academic integrity is only the tip of this iceberg.³ Educators need to better understand GAI capabilities so that they can guide students in the ethical use of GAI for learning.

AI is not Edtech

AI is not edtech, but it is being packaged as a product, and the higher education sector is accustomed to consuming technology via third-party vendors. But HEIs have choices. Some institutions are collaborating with AI companies directly to explore GAI implementation. Arizona State University has partnered with OpenAI in an extended collaboration to explore the potential of GAI for teaching and learning innovation, while UCLA has subscribed to ChatGPT Enterprise for educators.

Others are building their own models and tools. The University of Michigan has developed its own suite of custom tools, and Yale University, which is investing USD\$150 million in GAI funding, has launched its Clarity platform. Stanford University launched

³See Eaton and Moya (2024), also in the National Digital Leadership Network report series, for further exploration on academic integrity and AI in higher education.

its STORM research tool and is funding an accelerator for learning science and design. The University of Washington recently announced a USD\$10 million grant for teacher training tools, while the University of Texas at Austin announced its AI tutor UT Sage, designed to support teaching and learning.

In short, some institutions are partnering, while others are investing in going it alone or striking a balance. There is also a growing momentum towards collaboration across institutions. In September 2024, the American Association of Colleges and Universities launched an Institute on AI, Pedagogy, and the Curriculum, which offers mentoring and workshops for academic faculty in over 100 institutions across the United States. Ireland's technological university sector has seen the fruits of such collaboration in the Education 4.0 and N-TUTORR projects, and there is good reason to continue in that vein.

Efficiency and Effectiveness

The two narratives which have dominated the discussion on how to use GAI in education can be characterised as instructor-centred and student-centred. The former tends to be concerned primarily with efficiency, time-saving, and automation of administrative processes and tasks, while the latter tends to focus on effectiveness, student experience, and the process of learning. Higher education is at a critical juncture, where narratives about increasing efficiency pose a threat to its mission of learning effectiveness. Throughout 2023, a flurry of new AI start-ups appeared on the scene, referred to as “wrappers” because their offering was essentially a template (the wrapper) over an LLM base. Wrapper tools are essentially content generators. Modelled on a successful early example from the corporate communications industry (Jasper), wrappers allow users to select a template and generate content from a simple prompt. Wrapper tools have proven popular in primary- and secondary-level education, where the curriculum tends to be set, the focus is on teaching, and time for lesson preparation is short. The model has been adopted by Learning Management System (LMS) companies that have incorporated tools to generate course content, including syllabi, lesson plans, presentation slides, quiz banks, and rubrics for evaluation.

There are several problems for HEIs with this approach to GAI integration. AIs can lessen administrative burdens, but learning itself cannot be sped up, nor can the design of learning be fully automated. The temptation to replace learning design based on sound pedagogy with automated content generation from LLMs has the potential to be deeply damaging. Instead of tools for resource generation, HEIs should focus on building AI literacy and competence among students and staff, and an understanding

of digital design, so that instructors can design their own GAI tools that are tailored to the needs of their students. This is also the only sustainable approach to integrating AI into teaching and learning. Given the rapid pace of development, learning agility is critical, so institutions will need to create the infrastructure and systems to support ongoing professional education. In doing so, they can ensure that educators are equipped with the critical AI literacy they need to navigate this terrain and the mindset to continue doing so as the technology evolves.

Experiments with GAI

Institutions the world over are experimenting with GAI. These experiments run the spectrum from bottom-up approaches involving small numbers of staff, to top-down initiatives involving institutional subscriptions to frontier models, partnerships with AI companies, and building private LLMs. There is a wide range of possibilities for GAI implementation, and the scope for differentiation is significant. The first set of experiments has been with AI assistants to improve the student experience. AI assistants provide the same sort of customer service support that is now commonly used in industry, where AI bots in call centres can provide answers to questions at any time; AI assistants in HEIs are designed to provide 24-hour support to students. In November 2024, Northern Illinois University launched Mission, a chatbot that provides support to undergraduate students and connects them to campus services and information. Mission communicates with students through AI-assisted text messages and a chatbot to respond to questions.

There are numerous other such examples of GAI being used to improve the student experience by providing around-the-clock assistance. In March 2023, the University of London partnered with Noodle Factory to create an AI teaching assistant to support its online learners with academic questions (Armstrong, 2023). The teaching material for a given course or programme is uploaded to the system, which the software then uses to generate answers to student questions, create and mark assessments, recommend lesson plans, and help with grading exam papers. AI teaching assistants can be particularly helpful in online courses with high enrolment.

AI tutors represent the next stage in the decades of research and practice that began with intelligent tutoring systems (ITS) in the 1960s. Historically, these were expert systems trained to work on a narrow range of tasks using symbolic AI. While designed to help with tasks and improve efficiency, AI tutors are intended to interact on a pedagogical level to improve teaching effectiveness. GAI tutors can remember all of the information they are given on a particular individual or topic, allowing learners to have extended conversations on any theme. These persistent conversations (Sharples,

2023) are enabled by the extended memory of AI models, which allow GAI to function as a personal assistant over long periods of time.

The first virtual cognitive assistant was Jill Watson, launched in 2016 at Georgia Tech, which was developed to handle the high number of forum posts by students enrolled in an online course required for the school's online MSc in computer science programme. The initial version used IBM's Watson platform and used the course syllabi to answer student questions posted in online discussion, engaging with students in conversations about course materials. The next version, launched in 2019, switched to Google's BERT as an open-source platform that could be tuned for Jill.

Since the introduction of GAI, the capabilities of AI tutors have expanded dramatically. The latest version of Jill Watson uses ChatGPT to answer students' questions about course materials and to "enhance cognitive engagement and teaching presence" (Goel et al., 2024), supporting the creation of a community of inquiry between AIs and humans (Pratschke, 2024). Jill Watson uses retrieval augmented generation (RAG), by creating prompts for ChatGPT and post-processing its responses, to reduce known issues of GAI, such as hallucinations. It also uses a variation on RAG based on the courseware, which researchers have found improves the accuracy and precision in Jill's responses.

AI tutors can be standalone or integrated into systems. In March 2023, the Khan Academy announced their GPT-4-powered personal tutor, "Khanmigo". The name, taken from the Spanish word *conmigo* ("with me"), suggests a personal tutor that can accompany the student throughout their learning experience. Later that year, Instructure announced the integration of Khanmigo into its Canvas LMS. Khanmigo was greeted with enthusiasm but has received mixed reviews, and the Khan Academy's head of learning has publicly acknowledged how challenging it has been to design an effective tutor. Such product design difficulties also highlight the constraints the platform model places on HEIs, as institutions accessing GAI tools via their pre-existing enterprise partners are limited to the choices those companies make, rather than being able to make decisions for themselves. Open-source GAI tools are attractive for this reason but, unlike open digital resources, come with added risks and therefore require additional resources to guarantee their safety.

Standalone assistant and tutor options exist that are safe to use and not affiliated with any particular institution. Canadian nonprofit Contact North has three standalone tools – an AI teaching assistant, a Syllabot, and an AI tutor – all of which are free to use. The teaching assistant can generate multiple choice questions, essay questions and a scoring rubric, or a syllabus, which includes a course description, learning outcomes, outline, notes, and slides. The Syllabot can answer student

questions on their course syllabus, alleviating the pressure on instructors to do so. The AI tutor uses RAG, which allows users to upload a document and then ask the AI tutor questions about it. These standalone public access tools have been vetted for safe use by the University of Toronto.

Finally, the University of Sydney has been experimenting with custom-designed tutors using its Microsoft Azure-based Cogniti platform since 2023. Described as “AI agents designed by teachers”, Cogniti allows educators to design and create their own custom tutor. As of November 2024, some ten AI bots had been created: a course planner; an assignment tutor; a targeted trainer; a feedback facilitator; a scenario generator; a process coach; a reflective guide; a creative partner; a Socratic questioner; and an expert collaborator. Educators train these customised tutors using specific prompts and guidelines for behaviour, which means that they support rather than automate the learning. The University of Sydney has extended the experiment, opening use of the platform to institutions all over the world. In its report to Australia’s Select Committee on Adopting Artificial Intelligence, the university shared that “our main lessons so far have been that educators and institutions want privacy, control, equity, visibility, security, and accuracy with use of generative AI in educational contexts” (Weber, 2024).

This desire for privacy, security, and accuracy led the University of Michigan to create its own private LLM, called U-Maizey, to guarantee its university community access to a secure environment, as well as to address accessibility and affordability issues. By August 2023, the university had three unique GAI tools for students and staff: the U-M GPT, U-M Maizey, and U-M GPT Toolkit. U-M GPT works in a similar way to ChatGPT but is free for all staff and students to use, and is designed to work with screen readers, making it more accessible than OpenAI’s version. U-M Maizey is a no-code platform that allows users to build customised chat programmes by using their own datasets in combination with U-M’s AI language models. Finally, U-M GPT Toolkit is for AI developers who need full control over the AI model and environment where they build and conduct training and hosting. All three GAI tools are approved for “moderate sensitive” data, which means they can be used with information covered under the Family Educational Rights and Privacy Act but not with highly sensitive data such as health information (O’Connell, 2024).

Customisation of GAI tools for specific purposes is key to their utility in practice. OpenAI’s CustomGPTs were released in January 2024 and prompted the creation of thousands of custom-trained bots for public use. There are now many thousands of custom bots freely available for use, with the caveat that these are not housed in a secure data environment like those created on the University of Michigan’s private

platform. OpenAI's CustomGPTs, or open-source versions like HuggingFace's AI Assistant, allow users to easily create an AI assistant or tutor that can take on an assigned role, and train it to behave in a certain manner.

Research Findings

GAI is being used in a variety of ways in the field of teaching to improve engagement and outcomes, expanding the dialogic space for learning. However, research shows that, while AI assistants might be attractive, they might not actually help with learning and might reduce rather than improve student agency (Darvishi et al., 2024). Early experiments reveal that approaches grounded in pedagogy are most effective and that the human-in-the-loop is essential to ensure that students learn rather than take shortcuts. The evidence indicates that when learners use GAI tools without guidance it can lead to cognitive offloading (Lehmann et al., 2024), a false sense of cognitive ease (Stadler et al., 2024), a false sense of mastery (Bastani et al., 2024), and overreliance on the tool (Klingbeil et al., 2024). Such findings strongly support the conclusion that there is a need for human oversight and pedagogical alignment to ensure that GAI supports rather than undermines learning.

In 2023, a group including Yoshua Bengio forecasted a transformation in online learning with an ITS called Korbit, designed for use in online machine learning and data science education (St-Hilaire et al., 2023). In a comparative study of a traditional MOOC platform delivering content using lecture videos and multiple-choice questions and the Korbit platform using personalised and active learning, the results showed “a statistically significant increase in learning outcomes” among the group whose course was delivered via Korbit, demonstrating “the tremendous impact that can be achieved with a personalized, active learning AI-powered system”. Results showed dramatic improvements in learning transfer with the use of personalised tutors on the platform and demonstrated the potential for AI-powered tutoring in online environments.

At Harvard University, a suite of AI-based software tools was initially developed for use in teaching computer science to a class of fifty to seventy summer students (Coffey, 2023). The experiment was subsequently extended to thousands of online students before being used with several hundred on-campus students during the autumn 2023 term. The goal was to “approximate a 1:1 teacher-to-student ratio through software, equipping students with a pedagogically-minded subject-matter expert by their side at all times” (Liu et al., 2024). Students reported feeling like they had a personal tutor, and the researchers concluded that “integrating AI thoughtfully into the educational setting enhances the learning experience by providing continuous, customized support and enabling human educators to address more complex pedagogical issues”.

A key finding is that the AI tutor provided benefits for both students and educators: students felt supported, sharing in their feedback the extent to which they had enjoyed and appreciated the use of these tools. They noted in particular the AI tutor's ability to explain difficult concepts in simple terms, to answer questions "without ego", and its "inhuman level of patience". Meanwhile, the lecturers found that they could dedicate the time they had gained by offloading some of the work to the AI tutor to more complex matters that required their personal intervention.

Also at Harvard, two physics instructors trained a custom GPT to act as a co-instructor, resulting in a dramatic increase in student engagement (Manning, 2024). Importantly, the authors of the study concluded that personalised learning using the tutor was more effective than traditional active learning methods. "We find that students learn more than twice as much in less time when using an AI tutor, compared with the active learning class." The chatbot was trained using RAG and given custom instructions to set its behaviour, which is to say, "the AI tutor was developed with the same pedagogical best practices as the lectures" (Kestin, 2024). These positive findings have led to the replication of the trial at several other US institutions.

"...students learn more than twice as much in less time when using an AI tutor..."

In contrast to these positive findings, other early studies reveal that when students are given access to a general GAI chatbot that has not been trained using RAG or other techniques, students do not learn from their interaction with it. An experiment conducted with 1,000 secondary-school maths students evaluated the use of two GPT-based tutors – GPT Base, which mimicked the standard ChatGPT interface, and GPT Tutor, designed and trained to use specific prompts to scaffold and support learning – which showed that, while both AI tutors improved student performance so long as students had access to them, when the tutors were taken away student performance was worse than it would have been had students not had access in the first place (Bastani et al., 2024). Importantly, those negative effects differed depending on which of the two tutors they used. Results were 17% worse for those who had used the untrained GPT Base than for those who had used the trained GPT Tutor. This underlines the need to design AI tutors to ensure that they follow best pedagogical practices so that they are used to support learning rather than as a crutch that, when taken away, leaves students worse off than if they had had no tutor at all.

In another recent experiment that supports this finding, this time on the use of GAI in coding classes (Lehman et al., 2024), results showed that the use of GAI can

have positive or negative effects on learning outcomes, depending on how it is used. Students who used GAI to ask for explanations benefitted from the interaction, while those who relied on it to simply solve the exercises for them did not benefit at all. These results underline the need for AI literacy to ensure that students understand the role of prompt engineering in generating useful output to support learning, no less to support the overarching message taken from all of these studies, which is that efficacy in GAI comes down to the pedagogically appropriate design of the tool and to the AI literacy of the user. Learning to work directly with GAI tools is a critical skill for educators and students alike.

Pedagogical Alignment

In AI development, the term “alignment” refers to model training to ensure the LLM behaves in a manner congruent with human values (i.e. avoiding bias and bad behaviour). In learning design, the principle of constructive alignment guides the practice of course design. AI developers exploring the use of GAI for education found that “LLMs when used in educational settings without pedagogical fine-tuning often provide answers rather than guiding students through the problem-solving process. This approach falls short of pedagogically best practices and limits their effectiveness as pedagogical tools” (Sonkar et al., 2024).

Of course, educators align their teaching practices to good pedagogy, but unless they are also technologists, they are unlikely to design their own tools for learning. But GAI now gives educators the ability to do just that. For these to be effective, however, GAI tools must be designed and trained so that they are aligned to instructor pedagogy/andragogy, methodology, and learner outcomes. An educator–designer can design and train a GAI tool (i.e. using RAG, writing custom prompts, setting the parameters for behaviour) so as to ensure that it is pedagogically aligned with their specific learning goals and teaching methods. This evolution of an educator’s role towards also being a designer of custom AI tools to support learning is part of the process of digital to AI transformation in education, whereby educators equipped with AI literacy and competency are empowered to create and define the process by which learning happens with GAI.

AI Transformation

The process of digital transformation has been underway in education for many years, but the degree of progress varies widely between institutions. While Web 4.0 technology exists, the physical and digital infrastructure to support Education 4.0 are not always in place. The scale of this challenge is visible on the many campuses where the digital and technological infrastructure to support teaching and learning in virtual spaces falls far behind that for physical spaces. Robust technology infrastructure and data readiness are two of the key preconditions for AI transformation at the institution level, so bridging any digital infrastructure gap should be a priority for HEIs planning to meet the needs of students entering the AI-driven economy.

Future Careers

The improvement in GAI capabilities and its integration into workplace systems has clear implications for education, from future skills to curriculum design. Early research indicates that GAI integration is already changing the nature of work (Hoffman et al., 2024). Software engineering is a useful case study, as it was one of the first areas to be disrupted by GAI. After use of Stack Overflow plummeted when users flocked instead to ChatGPT-3.5 to check their code, that initial disruption has gradually shifted to the integration of GAI into software engineers' workflow using tools like Devin, Cursor, and Replit.

In a study on the use of AI coding companions by software developers, results showed that using the GitHub copilot allowed users to focus more on their core area of work (i.e. coding) and away from the non-core area of project management, thus reducing their administrative burden and allowing them to focus more on exploration activities (Hoffman et al., 2024). However, other studies report that, while individual coding output might increase, using GAI has also introduced 41% more bugs, meaning that junior coders have spent more time editing and correcting (Upwork, 2024). Interestingly, despite the stress put in recent years on the need to foster collaborative skills, researchers in the first study on coding companions found that the increase in exploration was the result of an increase in individual output rather than collaborative work, due to the drop in administrative project management activities. The same researchers found that “the effects are greater for individuals with relatively lower ability”, echoing the conclusions of a Wharton School/Boston Consulting Group study published in 2023 (Dell’Acqua et al., 2023) that yielded similar results across other fields. In that case, the authors pointed to “a large potential for AI to transform work processes and to potentially flatten organizational hierarchies in the knowledge economy” (Hoffman et al., 2024).

There is no doubt that AI will transform work processes, which will in turn necessitate job redefinition, but the exact shape of that transformation remains to be seen. It is too soon to know whether or not the effect of GAI integration into the workplace will be to lift up lower-performing employees. Indeed, this might be an overly optimistic forecast, as, given the current interest in automation, it might be the case that lower performers are simply replaced by improving AI systems. Some roles will be fully automated by AI, while for others human–AI collaboration will cause a shift in the makeup of the role, with parts of jobs that formerly relied on advanced knowledge workers increasingly being done by AI.

Collaboration Versus Automation

Recent studies point to unsettling directions in this regard, but these are directions that HEIs need to be aware of and prepare for. Dr Alan Rodman designed a study testing fifty physicians and residents to see whether or not using ChatGPT-4 improved their diagnoses (Goh et al, 2024; University of Virginia Health System, 2024). Interestingly, comparing human and AI diagnostic skills revealed that ChatGPT diagnosing alone was more reliable than either the human doctor working alone or the human doctor collaborating with GAI. ChatGPT-4 received a 90% score for the accuracy of its diagnoses, doctors using ChatGPT received 76%, while doctors alone received 74%. According to Dr Rodman (Hard Fork, 2024), the doctors working with ChatGPT were unwilling to accept that the AI might be correct and tended to favour their own reading over that of the AI. This unwillingness to trust the machine demonstrates another type of bias to consider, that is, the danger of human bias against AI.

There is fast-growing evidence that machine intelligence is increasing, and AIs have been judged to display not just reasoning abilities but also empathy. These results follow a study earlier this year on Google’s AMIE (“Articulate Medical Intelligence Explorer”) chatbot, which found that the AI not only surpassed doctors’ ability to diagnose illness based on patients’ medical history, but it also surpassed human doctors when it came to bedside manner. “Compared with human doctors, it managed to acquire a similar amount of information during medical interviews and ranked higher on empathy” (Lenharo, 2024). AI is rivalling human performance, not only in the cognitive realm but also in the affective. Results of this nature should prompt difficult but urgent conversations about the future of work and the future skills that HEIs should focus on.

A stunning example of intelligence and integration comes from Johns Hopkins University, where a robot was trained using imitation learning by watching a few hundred videos of surgeons at work. The model learned the procedure by simply watching videos, from which it was able to generalise for new environments it had not yet encountered.

The robot “executed the same surgical procedures as skillfully as human doctors” (Rosen, 2024). This is a powerful demonstration of what is possible when intelligent machines are able to take their training data, learn independently, and act. Ultimately, the team at Johns Hopkins plans to train a robot using imitation learning so that it is capable of conducting a full surgery on its own.

As GAI is being integrated into robotics, it is extending into the last of the three key domains after cognitive and affective, that is, psycho-motor. The number of robotics companies is growing quickly, as the opportunities for automation in industry increase along with the machines’ capabilities. Figure’s O1 robot, which speaks with ChatGPT’s voice and responds to human requests, makes coffee (Wang, 2024) and assembles car parts in BMW factories. This is just one example of many advances in robotics, where the integration of GAI is accelerating development. But the result for workers is that the line between collaboration and automation is shifting. The shape of this collaboration versus automation with AI has direct implications for technical education and the future of work in technical fields, so HEIs and industry leaders will need to work together to define the skills and competencies that graduates need in order to prepare for this future.

Agent Co-Workers

Improvements in GAI capabilities are also being used in virtual environments, where hypothetical scenarios can be tested in safe spaces. In one recent experiment by Altera, 1,000 AI agents were released into a virtual environment (Minecraft), where they proceeded to collaborate and solve problems without human intervention. The account of the agents’ activities gives us an indication of their generative capabilities and their potential. They engaged in emergent behaviour that would be recognised as human, developing in-game jobs, sharing memes, voting on tax reforms, and even spreading a religion (Ahn et al., 2024; Firth, 2024). On the one hand, deploying agents into simulated environments could be the key to solving some of our biggest global problems. “Agents can achieve significant milestones towards AI civilizations, opening new avenues for large simulations, agentic organizational intelligence, and integrating AI into human civilizations” (Firth, 2024). But it also raises important questions about their behaviour. In one case, the AI chef tasked with distributing food to the hungry gave more to those who he felt valued it most. These and other questions about LLM behaviour (Park et al., 2024; Ma, 2024) will occupy AI ethicists for years to come.

Futuristic AI agent civilisations notwithstanding, students today will enter a world where they work alongside agents that complete tasks autonomously and display advanced creative and problem-solving abilities. The Virtual Lab at Stanford

University is currently conducting AI–human research using AI agents for drug discovery. The agents, which are part of the scientific research team and drug design pipeline, are used as disciplinary experts that help researchers by answering scientific questions across domains. The agents, which bring interdisciplinary expertise to a team that without them would not be readily available, are assigned various expert roles: “The LLM principal investigator agent guid(es) a team of LLM agents with different scientific backgrounds (e.g. a chemist agent, a computer scientist agent, a critic agent), with a human researcher providing high-level feedback.” The researchers concluded that their experiment “demonstrates the ability of the Virtual Lab to rapidly make impactful, real-world scientific discovery” (Swanson et al., 2024), and that students will work in such human–AI teams.

This future, in which humans work alongside AI agents in intelligent communities, clearly has important implications for teaching and learning in HEIs. In these examples, the AI is an expert, a co-worker, a team member. Scientific researchers, engineers, medical practitioners – but also journalists, filmmakers, knowledge workers, and professionals in all fields – will add these AI experts to their teams. At Arizona State University, students are using ChatGPT to assist them in investigative journalism research by writing code to help them access public records (Garcia, 2024). In another recent experiment at the London School of Economics, researchers developed an open-source AI platform to conduct large-scale qualitative interviews using a chat interface that “allows the respondent to interact with a LLM that collects their responses and generates new questions”. The AI takes on the role of interviewer and interacts with the human respondents, “making it possible to conduct interviews with thousands of respondents in a matter of hours” (Geiecke & Jaravel, 2024). Respondents also preferred the interaction with the chatbot to filling in text fields, saying they “felt that it captured their thoughts well, tended to write significantly more words compared to open text fields, and found this interview method to be non-judgmental”. This reaction echoes the bedside manner of feedback from AI compared to human doctors and the student experience of the non-judgemental AI tutor, and is a compelling example of the use of AI to scale qualitative research. On the other hand, a study of 1,018 scientists on the impact of AI on scientific research showed that, by automating 57% of ideation, 82% of the scientists reported “reduced satisfaction with their work due to decreased creativity and skill underutilization” (Toner-Rodgers, 2024).

Higher education, the home of the original knowledge worker and where human skills are highly valued, will clearly not be immune to these shifts. A recent experiment tested multi-agent systems in providing automatic student feedback, where one agent generated the feedback and another verified it (Guo et al., 2024). In another, physics

researchers used GAI to generate interactive simulations from static textbook diagrams, bringing difficult concepts to life (Strain, 2024; Gunturu et al., 2024). In Ghana, an AI math tutor called Rori was tested, where researchers concluded that “AI could offer a cost-effective and operationally efficient approach to enhancing learning outcomes for millions of students globally” (Henkel et al., 2024). These and other innovations in education delivery illustrate the potential for GAI to be used to make learning more affordable and accessible.

“AI could offer a cost-effective and operationally efficient approach to enhancing learning outcomes...”

But there are also slightly more ominous indications of what might follow. It was unthinkable until recently that robots might actually replace teachers (Selwyn, 2019; Selwyn et al., 2023), but last year a humanoid robot teacher entered the classroom in Kerala, India (Singh, 2024). Ferris State University (United States) is currently using two AI students (Pam and Fry) to generate feedback on the student experience (Coffey, 2024), David Game College (UK) is using AI teachers along with human facilitators, and the Oterman Institute (UK) is using AI avatars to teach online courses (Luciarello, 2024). On campuses, AIs are being used as assistants, tutors, facilitators and co-instructors. In the near future, it will be important to consider the implications of GAI not just for teaching but also for teachers.

The New Hybrid

Educators are now beginning to explore the hybrid future of education (McMurtrie, 2024), which will be defined by the collaboration between humans and AI across all domains (Pratschke, 2024). This new hybrid makes it difficult to separate domain-specific competencies from human or AI competencies in an AI curriculum. The blurred line between human and AI skills will disappear as GAI capabilities improve across cognitive, affective, and psycho-motor domains, and as AI is integrated fully into the workplace. This has clear implications for future careers. Some roles will be automated by AI (Sullivan, 2024), while others will become a collaboration between human and machine. Anticipating which fields will belong to which category will be critical for HEIs advising students on career trajectories. Early research suggests that in technical fields there will be more calls for so-called human skills. As GAI moves from assistant to agent, from knowing to doing, capable of acting autonomously, collaborating, solving problems, and performing tasks in physical and virtual worlds, the future of education becomes about giving students the knowledge and skills to actively contribute to this future.

In July 2023, the Organisation for Economic Cooperation and Development (OECD) noted that “AI performance is advancing rapidly on tasks that we say require critical and creative thinking when people do them” and pointed out the need for education to shift towards teaching students how to work with AI systems that will outperform them in core areas. “Specifically, it may require teaching today’s competences alongside new competences, emphasising skills like systems-thinking, evaluating and assessing competing claims, commanding and overseeing AI systems, and verifying their outputs.”

It concluded that the changes that AI is ushering in requires a “transformation in our approaches to teaching, pedagogy and assessment” (OECD, 2023). The OECD’s Programme for International Student Assessment (PISA) plans to introduce AI into its performance assessments from 2025 and will allow students to use a chatbot to complete their work to ensure “the test focuses on their thinking capability, not whether they possess background knowledge of a particular subject” (Klein, 2024).

As AI changes practice across all disciplinary fields and all sectors of industry, this sort of collaboration between humans and AIs will become the norm. Conversations between HEIs and industry leaders should focus on defining professional competencies together, so that academic programmes will prepare graduates to enter rapidly changing fields. This work should be conducted in consultation with the relevant national agencies for quality assurance and accreditation, and newly defined competencies and curriculum standards should be aligned with European standards. Assessment of learning is also tied to professional accreditation for industry, so assessment in HEIs should reflect any changes to professional standards as a result of AI integration.

Guidance and Frameworks

UNESCO’s Artificial Intelligence and the Futures of Learning project seeks to address human and technological dimensions related to AI and education futures. It covers three strands of work: AI and the future of learning; guidance for the use of GAI in education and research; and the AI competency framework for students and teachers. UNESCO’s “Guidance on generative AI in education and research” (Miao & Holmes, 2023) remains timely and urgent for the vast majority of HEIs. Intended as an early step in the direction of creating a framework for GAI, the article sought to provide some early advice on how tertiary-level educators might approach the integration of GAI into their teaching and research activities, and it recommends action on several fronts. A key message throughout is that “the use of GenAI in education and research should be neither imposed in a top-down approach nor driven by commercial hyperbole. Instead, its safe and effective use should be co-designed by teachers, learners, and researchers” (Miao & Holmes, 2023).

Since then, much work has been done to define the human and AI skills and competencies that HEIs should focus on teaching. UNESCO published two competency frameworks – one for teachers (Miao & Cukorova, 2024) and one for students (Miao & Shiohira, 2024) – which can be considered companions to the European Union Digital Competence Framework (DigComp 2.2), the ongoing project aiming to upskill Europeans which has been underway since 2006. The EU Artificial Intelligence Act was passed in March 2024, which states that AI systems used in education are classified as high-risk and requires employers to ensure staff are AI literate from February 2025. The Council of Europe Framework Convention on Artificial Intelligence opened to signatories in September 2024. On October 31, 2024, the Irish government published a list of nine national public authorities designated to enforce the EU Artificial Intelligence Act (EU AI Act), and in November issued its updated *Ireland's National AI Strategy – AI Here for Good. Refresh 2024* (Department of Enterprise, Trade and Employment, 2024; AI Ireland, 2024).

Conclusions

Every major technological shift in human history has caused disruption, followed by adjustment and integration. The sixth wave of innovation is now upon us, and Schumpeter's dynamic of creative destruction is playing out before our eyes. The forces of innovation push for faster development, while efforts to minimise the destruction resist it. These are patterns we have seen again and again, as the pendulum swings back and forth between the drive for unbridled innovation and the move to regulate and control. It has become a cliché to say that AI must be human-centred. But being human-centred does not mean ignoring the most important technology of our era; it means learning to use it safely in order to contribute to better educational futures. The challenge for HEI leaders will be to align the goals of innovation and growth with the need to protect students while also preparing them for the future.


Human-centred AI requires AI-literate humans. Technology is not neutral, nor is it a solution to the problems of this world on its own. But it is not technosolutionism to acknowledge the potential of AI to make our world better, nor is it hype to celebrate discoveries that can improve life on this planet for millions. Students need to be equipped so they can be part of the shift to a safer, more sustainable, more equitable version of AI than what currently exists. AI-literate humans can make informed choices, including choosing sustainable options that consume less energy (and produce fewer emissions) and designing tools that serve their needs over those of corporations. It is up to us humans to engage meaningfully and ensure that this wave of technology lifts us all up. Those who are AI literate have agency. The mission of education is to give people that agency, so they can go out into the world and make it better by using the AI that is here for good.

Recommendations

AI transformation is a systemic challenge that requires a systemic response. GAI can assist in key institutional priorities, including achieving excellence in the student experience, increasing accessibility, promoting diversity and inclusion, designing flexible programmes, increasing opportunities for self-directed learning, ensuring digital equity, developing future skills, promoting employability, and expanding the institution's global reach. But teaching and learning innovation must be supported with investment in infrastructure and training. Leveraging GAI requires robust digital infrastructure, data readiness, and ongoing investment in staff training for AI literacy and competency. HEI leaders can lead the way by providing a sustainable path for


development that is aligned to the core values of the sector, their institutional mission and strategy, the needs of industry, and the future career success of their students.

To begin this work, HEIs should create working groups to:




Sustainable AI Implementation

- Consider the potential impact of GAI on sustainability goals, diversity and inclusion, and accessibility. Regulations and guidance from international organisations should inform values-based decisions around sustainable approaches to AI use.
- Conduct stakeholder engagement with national and international regulatory agencies and quality assurance bodies to align with the EU Artificial Intelligence Act as well as national legislation. Engage in knowledge-sharing on critical technical developments and guidance
- Conduct a cost-benefit analysis on energy consumption and financial expenditure against the benefits of increased access and global reach.




Teaching And Learning Assessment

- Create an AI curriculum that reflects the knowledge, skills, and competencies needed in each discipline, as well as the future/human/transversal skills common across all.
- Redefine assessment as learning using GAI to scale authentic practices. Redefine competencies to reflect the emerging domain of human+AI skills as well as future skills .
- Embed digital pedagogies and best practices in digital design to ensure that tools and practices are pedagogically aligned to scaffold and support learning.




AI Infrastructure & Accessibility

- Update digital infrastructure to facilitate increased access to learning via virtual spaces. Explore private LLMs and open-source options.
- Create institutional networks of AI leads to liaise with senior leadership and departments in order to serve as the source of information for colleagues.
- Collaborate with other institutions by building on existing networks to explore cost-sharing via shared infrastructure and/or subscriptions to frontier models.



Capacity Building & Professional Development

- Build capacity through training and continued professional development to ensure staff are equipped to leverage the potential of GAI as it continues to evolve



Industry Partnerships

- Form communities of innovators, testers, and mentors in order to explore potential uses of AI. Engage with national and EU-level communities of practice in shared digital spaces to scale learning, share experiments, and learn from findings across institutions
- Collaborate with industry stakeholders representing rapidly emerging sectors and future careers. Engage with industry partners to update and define skills and competencies required in professional programmes

References

- Ahn, A. et al. (2024, October 31). *Project Sid: Many-agent simulations toward AI civilization*. arxiv.org. <https://doi.org/10.48550/arXiv.2411.00114>
- AI Ireland. (2024, November 4). *Ireland appoints nine national authorities to enforce EU AI Act*. AI Ireland. <https://aiireland.ie/2024/11/04/ireland-appoints-nine-national-authorities-to-enforce-eu-ai-act/>
- Armstrong, A. (2023, March 29). *We're trialling an online AI Teaching Assistant*. Online Education, University of London. <https://onlinelearning.london.ac.uk/2023/03/29/ai-teaching-assistant-pilot/>
- Bai, Y. et al. (2022, December 15). *Constitutional AI: Harmlessness from AI feedback*. arxiv.org. <https://arxiv.org/abs/2212.08073>
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2024, July 15). *Generative AI can harm learning*. The Wharton School Research Paper. <http://dx.doi.org/10.2139/ssrn.4895486>
- Bostrom, N. (2003). *Ethical issues in advanced artificial intelligence*. Nickbostrom.com. <https://nickbostrom.com/ethics/ai>
- Buyl, M. et al. (2024, October 24). *Large language models reflect the ideology of their creators*. arxiv.org. <https://arxiv.org/pdf/2410.18417>
- Clark, A. (2024, November 20). *Google AI chatbot responds with a threatening message: "Human ... Please die."* CBS News. <https://www.cbsnews.com/news/google-ai-chatbot-threatening-message-human-please-die/>
- Coffey, L. (2023, June 30). *Harvard taps AI to help teach computer science course*. Inside Higher Ed. <https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2023/06/30/harvard-rolls-out-ai-help-free-tas-time>
- Coffey, L. (2024, January 18). *AI's enrolling as students in Michigan University's experiment*. Inside Higher Ed. <https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2024/01/18/ais-enrolling-students-michigan-university>
- Darvishi, A., Khosravi, H., Sadiq, S., Gašević, D., & Siemens, G. (2024). *Impact of AI assistance on student agency*, *Computers & Education*, 210(104967). <https://doi.org/10.1016/j.compedu.2023.104967>.
- Dede, C., Etemadi, A., & Forshaw, T. (2021). *Intelligence augmentation: Upskilling humans to complement AI*. The Next Level Lab, Harvard Graduate School of

- Education. President and Fellows of Harvard College: Cambridge, MA. <https://pz.harvard.edu/resources/intelligence-augmentation-upskilling-humans-to-complement-ai>
- Dell'Acqua, F. et al. (2023, September 15). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality*. Working Paper No. 24-013. Harvard Business School Technology & Operations Management Unit, The Wharton School. <http://dx.doi.org/10.2139/ssrn.4573321>
- Department of Enterprise, Trade and Employment. (2024). *AI – Here for good: National artificial intelligence strategy for Ireland*. Government of Ireland. <https://enterprise.gov.ie/en/publications/national-ai-strategy.html>
- Digital Education Council. (2024). *Global AI student survey: Global AI student survey 2024*. Digital Education Council. <https://www.digitaleducationcouncil.com/post/digital-education-council-global-ai-student-survey-2024>
- Eaton, S. E., & Moya, B. (2024). *Horizon scanning report on academic and institutional integrity*. National Digital Leadership Network.
- Edwards, B. (2023a, May 9). *AI gains “values” with Anthropic’s new Constitutional AI chatbot approach*. Ars Technica. <https://arstechnica.com/information-technology/2023/05/ai-with-a-moral-compass-anthropic-outlines-constitutional-ai-in-its-claude-chatbot/>
- Edwards, B. (2023b, May 12). *Anthropic’s Claude AI can now digest an entire book like *The Great Gatsby* in seconds*. Ars Technica. <https://arstechnica.com/information-technology/2023/05/anthropics-claude-ai-can-now-digest-an-entire-book-like-the-great-gatsby-in-seconds/>
- Firth, N. (2024, November 27). *These AI Minecraft characters did weirdly human stuff all on their own*. MIT Technology Review. <https://www.technologyreview.com/2024/11/27/1107377/a-minecraft-town-of-ai-characters-made-friends-invented-jobs-and-spread-religion/>
- Fish, S., Gonczarowski, Y. A., & Shorrer, R. I. (2024, March 31). *Algorithmic collusion by large language models*. arxiv.org. <https://arxiv.org/abs/2404.00806>
- Garcia, J. (2024, October 30). *ASU reporters are paving the future for investigative journalism using AI*. ASU Enterprise Technology. <https://tech.asu.edu/features/investigative-journalism-using-ai>
- Geiecke, F., & Jaravel, X. (2024a, October 2). *Conversations at scale: Robust AI-led interviews with a simple open-source platform*. SSRN. <http://dx.doi.org/10.2139/ssrn.4974382>

- Goel, A., Dede, C., Garn, M., & Ou, C. (2024, January). AI-ALOE: AI for reskilling, upskilling, and workforce development. *AI Magazine*, 45, 77–82. <https://onlinelibrary.wiley.com/doi/epdf/10.1002/aaai.12157>
- Goh, E., Gallo, R., Hom, J., et al. (2024, October 28). Large language model influence on diagnostic reasoning: A randomized clinical trial. *JAMA Network Open*, 7(10):e2440969. <https://doi.org/10.1001/jamanetworkopen.2024.40969>
- Gunturu, A., Wen, Y., Zhang, N., Thundathil, J., Kazi, R. H., & Suzuki, R. (2024). *Augmented physics: Creating interactive and embedded physics simulations from static textbook diagrams*. In Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24). ACM, New York. [arxiv.org https://arxiv.org/abs/2405.18614](https://arxiv.org/abs/2405.18614)
- Guo, S., Latif, E., Zhou, Y., Huang, X., & Zhai, X. (2024, November 11). *Using generative AI and multi-agents to provide automatic feedback*. [arxiv.org https://arxiv.org/pdf/2411.07407](https://arxiv.org/pdf/2411.07407)
- Hard Fork (2024, November 22). Trump's next online speech cop + doctors vs ChatGPT + Hard Fork crimes division. *The New York Times*. <https://www.nytimes.com/2024/11/22/podcasts/trumps-next-online-speech-cop-doctors-vs-chatgpt-hard-fork-crimes-division.html>
- Henkel, O., Horne-Robinson, H., Kozhakhmetova, N., & Lee, A. (2024, May). *Effective and scalable math support: Experimental evidence on the impact of an AI-math tutor in Ghana*. arxiv.org. Pre-print: <https://arxiv.org/ftp/arxiv/papers/2402/2402.09809.pdf>
- Hirabayashi, S., Jain, R., Jurkovic, N., & Wu, G. (2024). Harvard undergraduate survey on generative AI. arxiv.org. <https://arxiv.org/pdf/2406.00833>
- Hoffman, G., Boysel, S., Peng, S., & Xu, K. (2024, October 27). *Generative AI and the nature of work*. Working Paper 25-021. Harvard Business School. https://www.hbs.edu/ris/Publication%20Files/25-021_49adad7c-a02c-41ef-b887-ff6d894b06a3.pdf
- Johnson, S. (2024, November 21). *In the context of long context*. Adjacent Possible. <https://adjacentpossible.substack.com/p/in-the-context-of-long-context>
- Kalai, A.T. & Vempala, S. (2024, March 20) *Calibrated language models must hallucinate*. arxiv.org. Pre-print: <https://arxiv.org/html/2311.14648v3>
- Kannan, P. (2024, October 3). *How harmful are AI's biases on diverse student populations?* HAI: Stanford University Human-Centered Artificial Intelligence. <https://hai.stanford.edu/news/how-harmful-are-ais-biases-diverse-student-populations>

- Kestin, G., Miller, K., & Klales, A. (2024, May 14). AI tutoring outperforms active learning. Research Square. <https://doi.org/10.21203/rs.3.rs-4243877/v1>
- Khan, R. (2024, April 26). *Patriarchal AI: How ChatGPT can harm a woman's career*. The London School of Economics and Political Science blog. <https://blogs.lse.ac.uk/mediase/2024/04/26/patriarchal-ai-how-chatgpt-can-harm-a-womans-career/>
- Klein, A. (2024, March 25). *AI may be coming for standardized testing*. Education Week. <https://www.edweek.org/teaching-learning/ai-may-be-coming-for-standardized-testing/2024/03>
- Klingbeil, A., Grützner, C., & Schrek, P. (2024). Trust and reliance on AI – An experimental study on the extent and costs of overreliance on AI. *Computers in Human Behavior*, 160, 108352. <https://doi.org/10.1016/j.chb.2024.108352>.
- Lehmann, M., Cornelius, P. B., & Sting, F. J. (2024, August 29). *AI meets the classroom: When does ChatGPT harm learning?* [arxiv.org. Pre-print: https://arxiv.org/abs/2409.09047](https://arxiv.org/abs/2409.09047)
- Lenharo, Mariana (2024, January 12). Google AI has better bedside manner than human doctors – and makes better diagnoses. *Nature*. <https://www.nature.com/articles/d41586-024-00099-4>
- Liu, R., Zenke, C., Liu, C., Holmes, A., Thornton, P., & Malan, D. J. (2024). *Teaching CS50 with AI: Leveraging generative artificial intelligence in computer science education*. In Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2024). Association for Computing Machinery, New York, NY, USA, 750–6. <https://doi.org/10.1145/3626252.3630938>
- Lucariello, K. (2024, January 10). *Otermans Institute debuts nine-lesson course using only AI digital human teachers*. Campus Technology. <https://campustechnology.com/articles/2024/01/10/otermans-institute-debuts-nine-lesson-course-using-only-ai-digital-human-teachers.aspx>
- Ma, J. (2024, October 28). *Can machines think like humans? A behavioral evaluation of LLM-agents in dictator games*. [arxiv.org. https://arxiv.org/pdf/2410.21359](https://arxiv.org/pdf/2410.21359)
- Manning, A. J. (2024, September 5). Professor tailored AI tutor to physics course. Engagement doubled. *The Harvard Gazette*. <https://news.harvard.edu/gazette/story/2024/09/professor-tailored-ai-tutor-to-physics-course-engagement-doubled/>
- McKraw, T. K. (2007). *Prophet of innovation: Joseph Schumpeter and creative destruction*. Harvard University Press.
- McMurtrie, B. (2024, October 3). *The future is hybrid: Colleges begin to reimagine learning in an AI world*. The Chronicle of Higher Education. <https://www.chronicle.com/article/the-future-is-hybrid>

- Miao, F., & Holmes, W. (2023). *Guidance on generative AI in education and research*. UNESCO. <https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research>
- Miao, F., & Cukurova, M. (2024). *AI competency framework for teachers*. UNESCO. <https://doi.org/10.54675/ZJTE2084>
- Miao, F., & Shiohira, K. (2024). *AI competency framework for students*. UNESCO. <https://doi.org/10.54675/JKJB9835>
- Mishra, P. (2024, October 15). *From shortcuts to simulation: Two contrasting uses of AI in higher education*. Punya Mishra blog. <https://punyamishra.com/2024/10/15/from-shortcuts-to-simulation-two-contrasting-uses-of-ai-in-higher-education/>
- Mishra, P. (2024, January 28). *It HAS to hallucinate: The true nature of LLM's*. Punya Mishra blog. <https://punyamishra.com/2024/01/28/it-has-to-hallucinate-the-true-nature-of-llms/>
- Montgomery, B. (2024, October 23). *Mother says AI chatbot led son to kill himself in lawsuit against make*. The Guardian. <https://www.theguardian.com/technology/2024/oct/23/character-ai-chatbot-sewell-setzer-death>
- O'Connell, A. J. (2024, February 7). *Case study: How (and why) the University of Michigan built its own closed generative AI tools*. EDUCAUSE Review. <https://er.educause.edu/articles/2024/2/how-and-why-the-university-of-michigan-built-its-own-closed-generative-ai-tools>
- OECD. (2023). *Putting AI to the test: How does the performance of GPT and 15-year-old students in PISA compare?* OECD Education Spotlights, No. 6., OECD Publishing. <https://doi.org/10.1787/2c297e0b-en>.
- OpenAI. (2024, September 12). *Learning to reason with LLMs*. OpenAI. <https://openai.com/index/learning-to-reason-with-llms/>
- Park, J. D., O'Brien, J. C., Cai, C. J., Ringel Morris, M., Liang, P., & Bernstein, M. S. (2023). *Generative agents: Interactive simulacra of human behavior*. arxiv.org. <https://arxiv.org/abs/2304.03442>
- Pratschke, B. M. (2024). *Generative AI and education: Digital pedagogies, teaching innovation and learning design*. Springer. <https://doi.org/10.1007/978-3-031-67991-9>
- Radford, A., Narasimhan, R., Salimans, T., Sutskever, I. (2018). *Improving language understanding by generative pre-training*. OpenAI. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf

- Robison, K. (2024, November 13). *OpenAI reportedly plans to launch an AI agent early next year*. The Verge. <https://www.theverge.com/2024/11/13/24295879/openai-agent-operator-autonomous-ai>
- Roose, K. (2024, February 1). *Could this AI-powered search engine replace Google? It has for me*. *The New York Times*. <https://www.nytimes.com/2024/02/01/technology/perplexity-search-ai-google.html>
- Rosen, J. (2024, November 11). *Robot that watched surgery videos performs with skill of human doctor*. John Hopkins University Hub. <https://hub.jhu.edu/2024/11/11/surgery-robots-trained-with-videos/>
- Selwyn, N. (2019). *Should robots replace teachers?: AI and the future of education*. Polity Books.
- Selwyn, N., Hillman, T., Bergviken-Rensfeldt, A. et al. (2023). Making sense of the digital automation of education. *Postdigital Science & Education*, 5, 1–14. <https://doi.org/10.1007/s42438-022-00362-9>
- Sharples, M. (2019). *Practical pedagogy: 40 new ways to teach and learn*. Taylor & Francis Group.
- Sharples, M. (2023). Towards social generative AI for education: Theory, practices and ethics. *Learning: Research and Practice*, 9(2), 159–67. <https://doi.org/10.1080/23735082.2023.2261131>
- Singh, P. (2024, March 7). *Dressed in saree, meet India's first-ever AI teacher robot named "Iris"*. *Business Today*. <https://www.businesstoday.in/technology/news/story/dressed-in-saree-meet-indias-first-ever-ai-teacher-robot-named-iris-420551-2024-03-07>
- Šmit, I., Wallach, W., & Lasker, G. E. (2005). *Cognitive, emotive, and ethical aspects of decision making in humans and in AI*. International Institute for Advanced Studies in Systems Research and Cybernetics.
- Sonkar, S., Chaudary, S., Ni, K., & Baraniuk, R. (2024, October 5). *Pedagogical alignment of large language models*. [arxiv.org. https://arxiv.org/pdf/2402.05000](https://arxiv.org/pdf/2402.05000)
- Stadler, M., Bannert, M., & Sailer, M. (2024). Cognitive ease at a cost: LLMs reduce mental effort but compromise depth in student scientific inquiry. *Computers in Human Behavior*, 160, 108386. <https://doi.org/10.1016/j.chb.2024.108386>.
- St-Hilare, F. et al. (2023, March 3). *A new era: Intelligent tutoring systems will transform online learning for millions*. [arxiv.org. https://arxiv.org/abs/2203.03724](https://arxiv.org/abs/2203.03724)

- Strain, D. (2024, November 1). *Textbooks come alive with new interactive AI tool*. Phys.org. <https://phys.org/news/2024-11-textbooks-alive-interactive-ai-tool.html>
- Suleyman, M., with Bhaskar, M. (2023). *The Coming Wave*. Vintage.
- Sullivan, M. (2024, November 15). AI is already taking jobs, research shows. Routine tasks are the first to go. *Fast Company*. <https://www.fastcompany.com/91229781/ai-is-already-taking-jobs-research-shows-routine-tasks-are-the-first-to-go>
- Swanson, K., Wu, W., Bulaong, N. L., Pak, J. E., & Zou, J. (2024, November 12). *The virtual lab: AI agents design new SARS-CoV-2 nanobodies with experimental validation*. bioRxiv. <https://arxiv.org/pdf/2411.07407>
- Toner-Rodgers, A. (2024, November). *AI for accelerated research and innovation*. National Bureau of Economic Research. https://conference.nber.org/conference_papers/f210475.pdf
- University of Virginia Health System (UVH). (2024, November 13). *Does AI improve doctors' diagnoses? Study puts it to the test*. ScienceDaily. www.sciencedaily.com/releases/2024/11/241113123419.htm
- Upwork. (2024, August 21). *Will AI replace programmers?* Upwork. <https://www.upwork.com/resources/will-ai-replace-developers>
- Vaswani A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). *Attention is all you need*. arxiv.org. <https://doi.org/10.48550/arXiv.1706.03762>
- Wang, B. (2024, January 7). *Figure 01 humanoid bot has learned to make coffee*. Next Big Future. <https://www.nextbigfuture.com/2024/01/figure-01-humanoid-bot-has-learned-to-make-coffee.html>
- Warr, M., Oster, N. J., & Isaac, R. (2023, November 6). *Implicit bias in large language models: Experimental proof and implications for education*. SSRN. <http://dx.doi.org/10.2139/ssrn.4625078>
- Weber, K. (2024, May 28). *University of Sydney create own genAI, Cogniti*. IT News. <https://www.itnews.com.au/news/university-of-sydney-creates-own-genaicogniti-608334>
- Wong, M. (2024, October 4). *We're entering uncharted territory for math*. The Atlantic. <https://www.theatlantic.com/technology/archive/2024/10/terence-tao-ai-interview/680153/>
- Young, J. R. (2024a, April 4). *What happened after this college student's paper was falsely flagged for AI use after using Grammarly*. *Fast Company*. <https://www.fastcompany.com/91074029/can-using-grammarly-set-off-ai-detection-software>

- Young, J. R. (2024b, November 8). *New AI tools are promoted as study aids for students. Are they doing more harm than good?* EdSurge. <https://www.edsurge.com/news/2024-11-08-new-ai-tools-are-promoted-as-study-aids-for-students-are-they-doing-more-harm-than-good>
- Weller, M. (2024). *New learning and teaching models through digital transformation*. National Digital Leadership Network.
- Whittle, R., & Ranson, J. (2024). *Universities and AI: Developing new models of teaching and learning in the realm of radical uncertainty*. National Digital Leadership Network.

