Centrum für Europäische Politik
FREIBURG | BERLIN

cepInput

No. 10 | 2024 13 August 2024

AI is Disrupting Education – For Better or Worse

Challenges and Strategies for Sustainable Learning and Institutional Resilience

Anselm Küsters

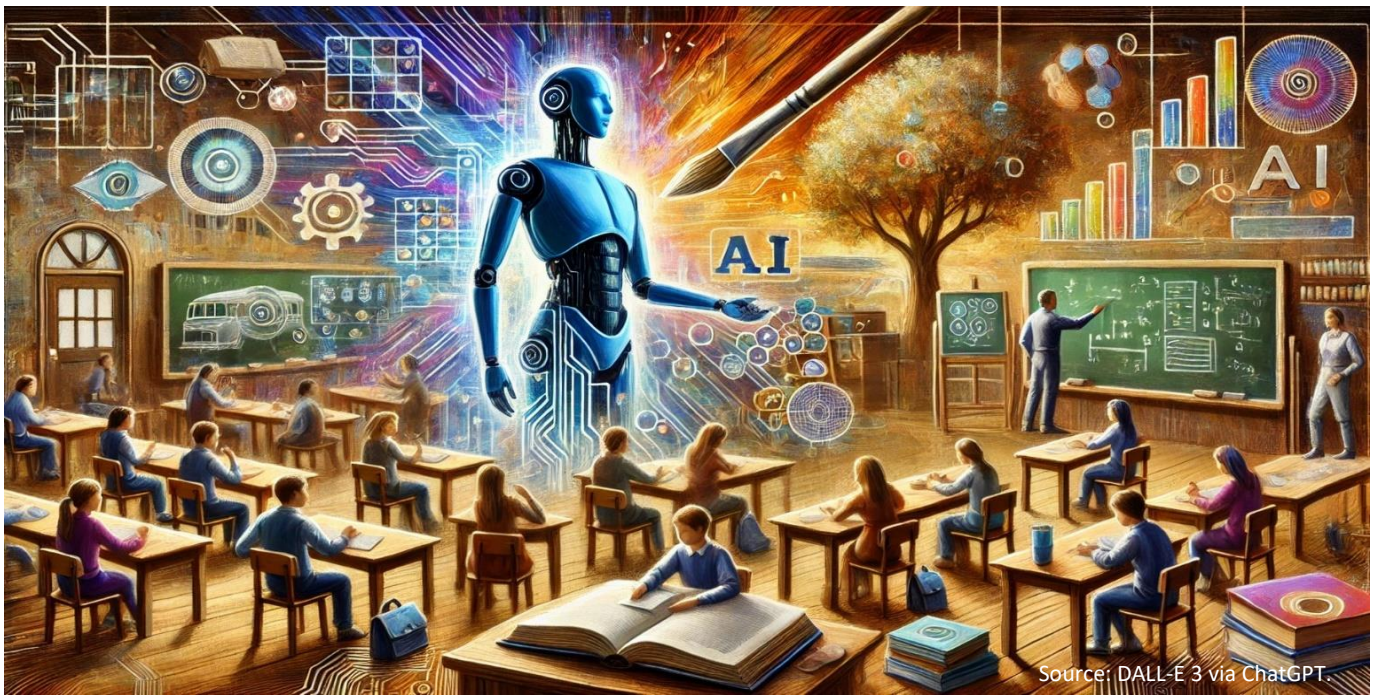

Source: DALL-E 3 via ChatGPT.

Since its release, ChatGPT has been used to write homework and scientific papers, although its use is rarely acknowledged. Large language models are changing our information practices more than the Covid pandemic. This cepInput explores the implications of this trend, ranging from unintentional plagiarism to compromised learning and peer review processes. It argues for an educational paradigm shift based on AI literacy curricula.

- Beyond data analysis and writing assistance, generative AI can also stimulate creativity and make knowledge more accessible. However, simply giving every student a ChatGPT account will not automatically improve learning, as experiments show that careful design is essential to ensure effective integration of AI in education. To ensure responsible use and ongoing learning, AI literacy courses should be mandatory at all levels of education, covering technical foundations, ethics, practical applications, and general problem-solving skills.
- Already, AI-generated content is damaging academic processes, as evidenced by compromised peer review processes and systemic manipulation of journals. Moreover, a lack of AI literacy increasingly divides communication within scientific disciplines, hampering progress. Finally, structural concerns about AI in education include asymmetric access to cutting-edge generative AI models and services as well as significant development and implementation challenges for non-English LLMs. Policy recommendations include supporting the development of LLMs for underrepresented languages and mitigating data-related issues. Concrete internal guidelines are needed to oversee the ethical use of AI in educational settings and to ensure accountability.

Content

Figures

Tables

1 Introduction

In March 2024, Florida allocated $2 million for grants to use Artificial Intelligence (AI) in middle and high schools to lessen teachers' administrative tasks and enhance student tutoring and learning.[1] The initiative requires that the AI platforms used adhere to strict privacy standards and provide transparency around student interactions, with a mandatory performance benchmark set at or above the capabilities of GPT-4, the large language model behind the popular chatbot "ChatGPT". The move has positioned Khan Academy, an American non-profit educational organization with more than 150 million learners in some 50 languages, as a leading provider of educational AI tools. Its founder, Salman Khan, a second-generation immigrant who studied computer science at the Massachusetts Institute of Technology in Cambridge, recently published a book arguing that AI will transform education and "create a new golden age for humanity".[2] This is just one of many examples of a global, increasingly competitive push to incorporate such new technologies into the education system. Especially the rise of AI agents, such as Meta AI, and the release of powerful open-source language models, such as Llama 3, will greatly enhance the accessibility and application of AI in education and research, for instance through AI-assisted searches and interactions known as "conversational pedagogical agents".[3]

Fig. 1: AI in the education sector

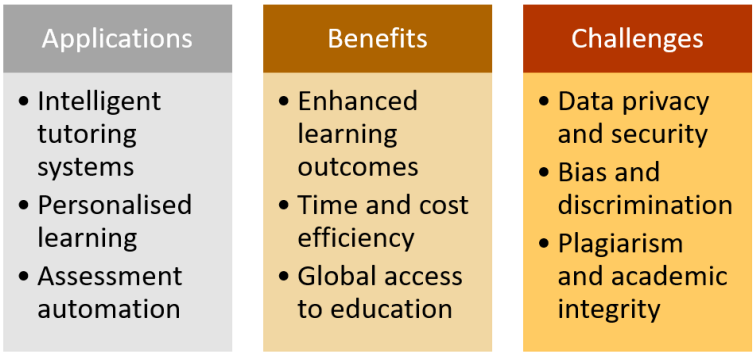

Source: Adapted from the literature review in: Abid et al. (2024).

The advent of generative AI in the fields of education, teaching, and research could amount to a significant turning point in the evolution of knowledge dissemination and acquisition. Potential benefits range from enhanced learning outcomes and global access to education to significant time and cost savings through personalised and intelligent tutoring system, while challenges exist with respect to data privacy and security, machine-inherent biases, as well as the threat of increased or involuntary plagiarism (Figure 1).[4] A recent literature review suggests that AI can revolutionise higher education by personalising learning methods, improving access to information beyond course materials, enhancing communication with professors, and applying motivational theories to increase student engagement.[5] Overall, the impact of these digital technologies to revolutionise writing and research cannot

1 Atterbury (2024), Florida takes its biggest leap into school AI programs, Politico.
2 Khan (2024), Brave New Words, Viking.
3 Hart (2024), Meta AI Faces Off Against Google, OpenAI With New Standalone Chatbot, Forbes.
4 Abid et al. (2024), A MAGDM approach for evaluating the impact of artificial intelligence on education, Front. Artif. Intell.
5 Triberti et al. (2024), "Better than my professor?" How to develop artificial intelligence tools for higher education, Front. Artif. Intell.

be overstated, given their suitability for enhancing so-called WINS ("Words, Images, Numbers, and Sounds") work.[6] As Wired magazine pointed out, AI cannot replace teaching, but it can make it better.[7]

In this context, this cepInput aims to assess the implications of integrating generative AI applications into processes of education and research.[8] This is crucial, as the EU's recent State of the Digital Decade report found that progress on digital skills is "alarmingly insufficient".[9] As policymakers and industry representative applaud the technology's potential, it is imperative to educate individuals through dedicated AI literacy curricula to ensure they have up-to-date methodological skills and to make them aware of certain pitfalls. Based on a meta-analysis of the current literature and the discussion of concrete use cases, this paper contrasts two approaches to generative AI in education. On the one hand, these AI applications could improve research, streamline the writing process, and foster a more inclusive academic environment (section 2). On the other hand, generative AI tools might still exacerbate socio-economic inequalities within academia, perpetuate biases, and lead to increased dependency of research on Big Tech (section 3). While these technologies offer potential for non-English-speaking scholars, enabling them to participate more fully in a global discourse, the accessibility of state-of-the-art models remains a formidable barrier. The financial implications of accessing and training cutting-edge AI systems risk widening the gap between well-funded private institutions and those with limited resources, typically public universities. Based on these pros and cons, the paper derives a number of policy recommendations (section 4). It concludes with a specific section on EU-related issues, which integrates some of the recommendations into the existing EU policy framework (section 5).

What are the main findings? First and foremost, it is crucial to explore strategies that can democratise access to open-source, transparent generative AI models to ensure that the trajectory of scientific progress and teaching is not dictated by economic prowess alone and to avoid perpetuating biases from the past. Achieving the optimistic scenario of integrating AI into research and education necessitates developing dedicated AI literacy curricula, including prompt design and tutorials about a conceptual understanding of large language models (LLMs). Finally, the rapid pace at which the field is advancing requires institutional updates to protect the learning path of students as well as the integrity of the peer review process, which is a cornerstone of academia. In short, educational providers and researchers must undergo a paradigm change that requires not only technical skills in AI but also the ability to evolve with AI technology and to critically assess its shortcomings.

2 Pro: Generative AI could improve research and learning

Generative AI, with its ability to analyse large datasets, produce content, and generate new ideas, promises to improve the research process across a range of disciplines. As Nathan Warren has noted, AI can act as both a "robust microscope, delving deep into vast data sets to unearth hidden patterns", and an "accelerator, speeding up our transformation of data into knowledge".[10] In other words, generative AI supports the basic research process through its ability to rapidly analyse and interpret large amounts of data. For instance, AI algorithms can analyse millions of papers, patents, and clinical trials

---

6 In general, see: Küsters (2024), Sprachtechnologie als Wettbewerbsvorteil der EU, cepInput No 7. For WINS work, see: Baier et al. (2023), Where Should Your Company Start with GenAI?, HBR.

7 Berdik (2024), AI Can't Replace Teaching, but It Can Make It Better, Wired.

8 This paper was initially written in the context of an online round table on the impact of AI on IR research and publishing, organized by the Istituto Affari Internazionali (IAI), which took place on June 11, 2024: https://youtu.be/cV_pXPM5BEM. A shortened version of this paper will be distributed as "forum article" in the *International Spectator*.

9 European Commission (2024), Communication, State of the Digital Decade 2024, COM/2024/260 final, p. 13.

10 Warren (2024), Unlocking the incredible with AI [Part 2] (exponentialview.co).

to identify barriers preventing scientists from translating their work into real-world applications, helping institutions to identify overlooked people and ideas.[11] This, in turn, might help to break down silos between many academic disciplines. Researchers are limited in the time they can devote to broad reading outside their narrow discipline and, due to cognitive limitations and behavioural heuristics, often fail to consistently link their decisions to best evidence.[12] In contrast, generative AI models can uncover interdisciplinary patterns at a speed and scale unattainable by human researchers alone.

Fig. 2: AI has surpassed human performance on several benchmarks

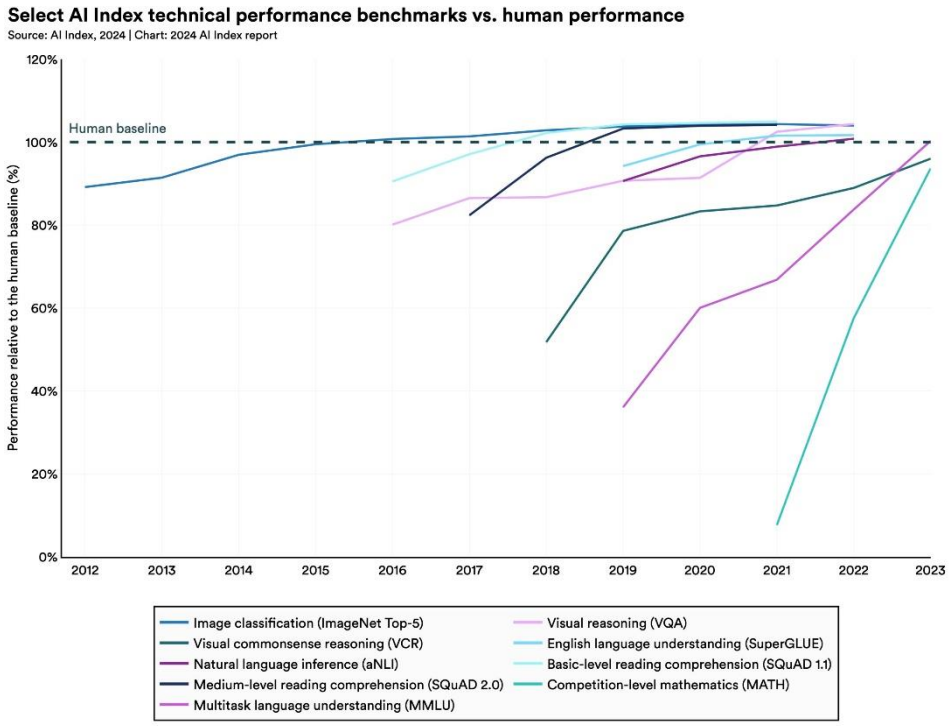

Source: Taken from: AI Index Steering Committee (2024), "The AI Index 2024 Annual Report.

The recent Artificial Intelligence Index Report 2024 from Stanford University's Institute for Human-Centered Artificial Intelligence underscores how broad the range of data is that generative AI models such as ChatGPT can nowadays analyse for researchers. When doing so, **these models often match or exceed human performance in areas as diverse as reading comprehension, image recognition, and advanced mathematics** (Figure 2).[13] This evolution requires novel assessment methods that address more complex cognitive tasks such as abstraction and reasoning. The report highlights the introduction of sophisticated testing frameworks, including the Graduate-Level Google-Proof Q&A Benchmark (GPQA), which challenges both AI and human PhD students with rigorous web-based multiple-choice questions.[14] Current results thus show that AI performs comparably to specialised human scholars. Still, it is worth bearing in mind that these models are ultimately probabilistic and thus continue to suffer from errors known as "hallucinations". It has also been speculated that many benchmarks are no longer valid as their questions and answers have been used as training data, thus "contaminating"

[11] Wang (2024), How I'm using AI tools to help universities maximize research impacts (nature.com).
[12] Morris (2018), Human Cognitive Limitations. Broad, Consistent, Clinical Application of Physiological Principles Will Require Decision Support. Ann Am Thorac Soc. 2018 Feb;15(Suppl 1):S53-S56.
[13] AI Index Steering Committee (2024), "The AI Index 2024 Annual Report," Stanford University, HAI_AI-Index-Report-2024.pdf (stanford.edu).
[14] Rein et al. (2023), Preprint at arXiv, https://doi.org/10.48550/arXiv.2311.12022.

the results. This means that there is a continued need for improved evaluation mechanisms to assess the competencies of AI systems in complex settings like education.

Fig. 3: Percentage of AI-engaged papers by field, 1985–2023 by field

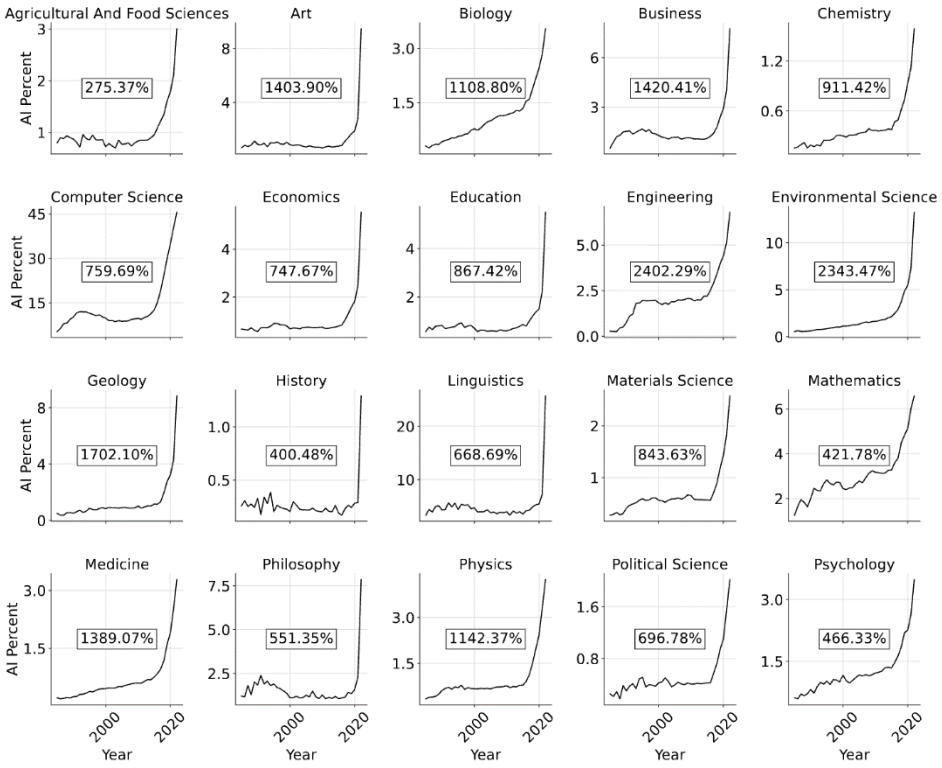

Source: Taken from: Duede et al. (2024), "Oil & Water? Diffusion of AI Within and Across Scientific Fields".

AI methods are increasingly being used in research projects both large and small, which has led to hopes that this technology will further accelerate research in various fields.[15] Google DeepMind, for instance, has developed "Graph Networks for Materials Exploration" (GNoME), which aims to help chemists discover new materials. This project uses advanced graph network technologies to analyse material properties and predict nearly 400,000 stable substances, potentially speeding up the discovery process in materials science labs.[16] The project "GraphCast" uses similar graph network techniques to provide fast and accurate weather forecasts, improving researchers' ability to respond to meteorological conditions more effectively.[17] Beyond such flagship projects of cutting-edge AI, AI is nowadays widely used in various STEM fields, particularly in medicine, materials science, robotics, agriculture, genetics, and computer science, as highlighted in a Royal Society report.[18] An empirical study of around 80 million research publications across 20 different scientific fields found an exponential growth in scientific engagement with AI from 1985 to 2022, with the number of AI-related publications increasing by around thirteen times across all fields (Figure 3).[19] However, this rapid uptake has also led to some

15 See: UN AI Advisory Body (2023), Governing AI for Humanity, interim_report.pdf (un.org), p. 7.
16 Peplow (2023), Google AI and robots join forces to build new materials, Nature.
17 Wong (2023), DeepMind AI accurately forecasts weather — on a desktop computer, Nature.
18 The Royal Society (2024), Science in the Age of AI, Report.
19 Duede et al. (2024), Oil & Water? Diffusion of AI Within and Across Scientific Fields, Arxiv.

problems, which are discussed below: The spate of so-called reproducibility failures suggests a certain immaturity of ML-based science, and the need for ongoing work on methods and best practices.[20]

Over the past months, the focus has been on generative AI systems, which promise to help streamline the writing process. There is indeed **strong evidence that ChatGPT has been used to write plenty of scientific papers since its release in November 2022, although it is not usually listed as an author or mentioned in the methodological appendix**.[21] Of the more than 1,600 scientists who responded to a Nature survey in 2023, nearly 30% said they had used generative AI to write a paper, and around 15% said they had used it for their own literature reviews or to write grant applications.[22] According to one estimate, around 10% of all research abstracts in 2024 will be written by ChatGPT.[23] The first systematic large-scale analysis of 950,965 academic papers published between January 2020 and February 2024 estimates a steady increase in the use of LLMs, with the largest and fastest growth observed in computer science papers (up to 17.5%), while mathematics papers and the Nature portfolio showed the least LLM change (up to 6.3%).[24] Overall, these results suggest that **LLMs are widely used in scholarly writing and are already changing our information and knowledge practices**. Based on the frequency of certain style words, **the advent of LLM-based writing assistants has had an even greater impact on the scientific literature than the Covid pandemic**.[25]

Fig. 4: Comparing LLMs and humans on the Divergent Association Task (DAT)

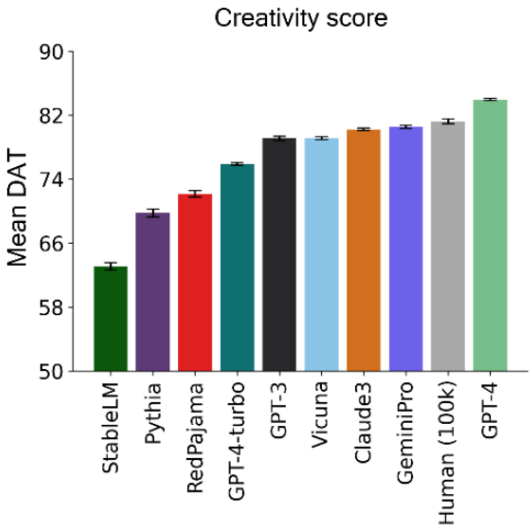

Notes: Mean DAT score and 95% confidence intervals. DAT = Divergent Association Task. Source: Taken from: Bellemare-Pepin et al. (2024), "Divergent Creativity in Humans and Large Language Models", p. 6.

Beyond data analysis and writing assistance, generative AI could stimulate creativity and help improve the *quality* of research, not just the *quantity*. In general, a systematic evaluation found evidence that LLMs can outperform humans in specific creative tasks such as divergent association and creative writing, based on a comparison between state-of-the-art LLMs and a large dataset of 100,000 humans

[20] Kapoor and Narayanan (2023), Leakage and the reproducibility crisis in machine-learning-based science, Patterns, 4/9.
[21] Conroy (2023), Scientific sleuths spot dishonest ChatGPT use in papers, Nature; Kannan (2024), How Much Research Is Being Written by Large Language Models?, Stanford.
[22] Van Noorden and Perkel (2023), AI and science: what 1,600 researchers think, Nature.
[23] Kobak et al. (2024), Delving into ChatGPT usage in academic writing through excess vocabulary, Arxiv.
[24] Liang et al. (2024), Mapping the Increasing Use of LLMs in Scientific Papers, Arxiv.
[25] Kobak et al. (2024), Delving into ChatGPT usage in academic writing through excess vocabulary, Arxiv.

(Figure 4).[26] AI models, particularly those designed for idea generation, can suggest research questions, hypotheses, and experimental designs that may not be immediately obvious to human researchers and thus help to reframe or even solve the problem.[27] Recent studies found that the average quality of AI ideas is quite high, and that the diversity of AI-generated ideas can be significantly improved through prompt engineering, approaching that achieved by larger groups of human subjects.[28] Prompt engineering is typically used as a shorthand for several different methods of manually improving and crafting the input (known as "prompt") to a generative AI model in order to improve the output. Other research has found that facilitating conversational learning, which involves a dynamic exchange of questions and answers between an AI-driven chatbot and a researcher, stimulates critical thinking.[29]

Consider, for example, LEGALSTORIES, a novel application of LLMs in legal education, where generated stories and questions helped to facilitate the learning and assessment of legal concepts.[30] To achieve this, the researchers first created a dataset comprising legal doctrines, LLM-generated stories, and expert-reviewed questions. By using a human-in-the-loop pipeline that integrates LLM output with expert input, high-quality educational content was then produced that enhances both comprehension and critical thinking (Figure 5). This innovative method not only improves legal literacy, but also highlights the wider potential of LLMs in education. All in all, **the integration of AI into research and teaching has the potential to be a source of qualitatively new insights**.

Fig. 5: Expert-in-the-loop pipeline for producing LLM-enhanced learning materials

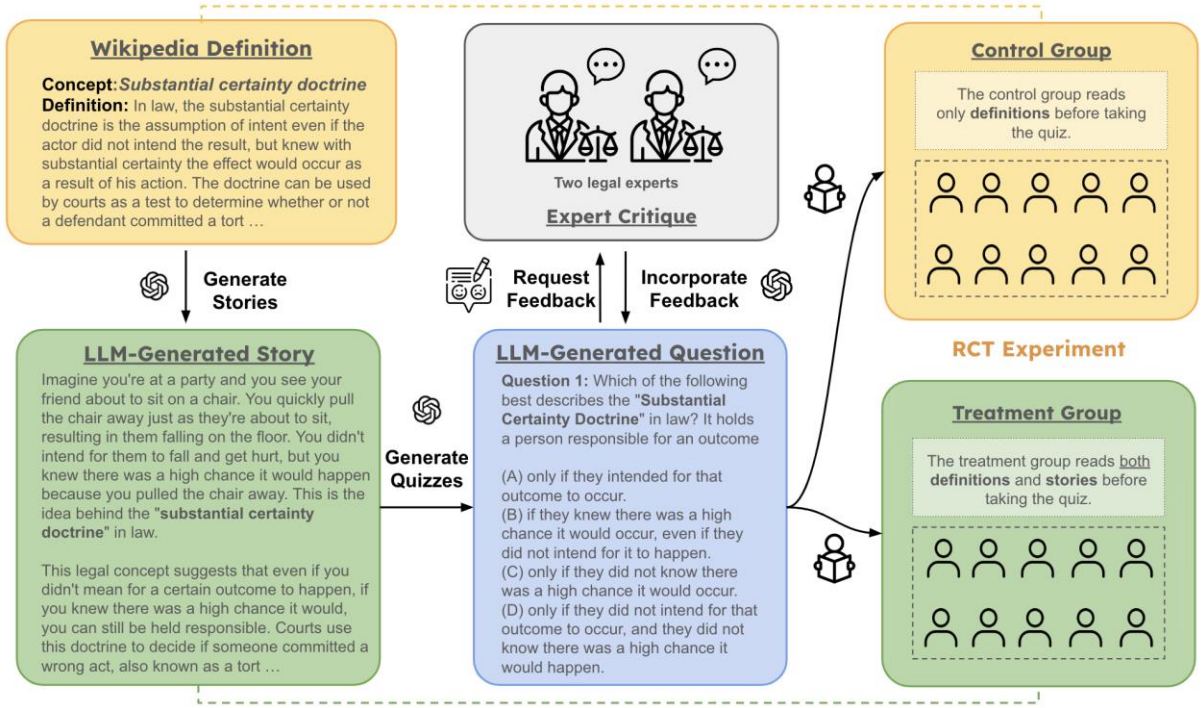

Source: Taken from: Jiang et al. (2024), "Leveraging Large Language Models for Learning Complex Legal Concepts through Storytelling".

[26] Bellemare-Pepin et al. (2024), Divergent Creativity in Humans and Large Language Models (arxiv.org), Arxiv.
[27] Rick et al. (2023), Supermind Ideator: Exploring generative AI to support creative problem-solving, Arxiv,
[28] Meincke et al. (2024), Prompting Diverse Ideas: Increasing AI Idea Variance, SSRN.
[29] Rospigliosi (2023), Artificial intelligence in teaching and learning: what questions should we ask of chat GPT? Interact. Learn. Environ. 31, 1–3 (tandfonline.com).
[30] Jiang et al. (2024), Leveraging Large Language Models for Learning Complex Legal Concepts through Storytelling, Arxiv.

However, it is not enough to focus on this increased *supply* of AI models for data analysis, writing assistance, and creative stimulation – effective *demand* is also needed to realise these opportunities. Only if there is sufficient "digital literacy" or "AI literacy" across society can the technology be used in a safe and practical way.[31] It is therefore **highly problematic that "only 55.6% of EU's population has at least basic digital skills and, at the current pace, the number of ICT specialists will reach just 12 million by 2030 – well below the EUR 20 million target and amid growing competition for digitally skilled talent"**, according to the latest EU progress report.[32] This problem is also pervasive in the sub-domain of academia: **AI literacy is increasingly becoming an issue that divides communication within scientific disciplines themselves, potentially hampering progress**. Using a natural language processing technique called document embeddings, which allow the similarities and differences of texts to be traced, researchers found a complex gap between AI-engaged and non-AI-engaged research within and across disciplines, suggesting that AI-engaged work is spreading across disciplines, but not mixing well with non-AI-engaged work.[33] This digital literacy gap, for example in knowledge of sound statistical inference, may also be an important factor in the current reproducibility crisis in machine learning-based science that has accompanied the rise of these methods in academia.[34] Preparing researchers and students for the increased integration of AI in their work therefore requires more efforts to spread knowledge about how AI works throughout society, as is already being done, for example, by the popular "Elements of AI" course in Finland.

In addition to basic AI literacy, it is crucial to develop intuitive understanding and practical application skills in students and researchers. **Rather than introducing new subjects in response to technological advances, it is more effective to integrate AI education into existing curricula**, as rightly suggested by Professor Jörn Loviscach.[35] Technical intricacies such as prompt engineering and neural network theory can be overwhelming and may not benefit those outside specialist fields, as these techniques are rapidly becoming obsolete. For example, learning how to boost prompts in a certain way may be incorporated in the next iteration of the LLM and therefore not relevant for every single student to learn. Instead, **it is more valuable to focus on equipping students to critically evaluate AI output and understand the role of AI in different contexts**. This includes fostering the problem-solving, critical thinking, and analytical reasoning skills that are essential for interpreting the outputs generated by AI systems. Students should learn to ask questions about AI-generated content, such as what data it was trained on and the robustness of its results.

A recent study of the impact of generative AI on learning illustrates how careful educational providers must be when integrating practical AI tools into education.[36] In a field experiment with nearly a thousand students, two GPT-based tutors were evaluated: one that mimicked a standard ChatGPT interface (GPT Base) and another with prompts designed to ensure learning (GPT Tutor). In line with expectations, the results showed significant performance improvements (48% for GPT Base and 127% for GPT Tutor). However, when access to GPT-4 was subsequently removed, students performed worse than those who never had access, suggesting that AI can be detrimental to educational outcomes if not used properly (17% reduction for GPT Base). Still, this negative effect was mitigated by some technical

---

31 Küsters (2023), ChatGPT erfordert mehr digitale Mündigkeit, cepAdhoc Nr. 1/2023; Schnöller and Privitera (2023), Politik und Verwaltung brauchen KI-Nachhilfe, Tagesspiegel Background.

32 European Commission (2024), Communication, State of the Digital Decade 2024, COM/2024/260 final, p. 13.

33 Duede et al. (2024), Oil & Water? Diffusion of AI Within and Across Scientific Fields, Arxiv.

34 Sayash and Narayanan (2023), Leakage and the reproducibility crisis in machine-learning-based science, Patterns, 4/9.

35 See his reflections: Loviscach (2024), Missing Link: Die GPT-fizierung des Studiums | heise online.

36 See: Bastani et al. (2024), Generative AI Can Harm Learning, SSRN.

safeguards in GPT Tutor, which provided hints without giving direct answers. This finding suggests that while generative AI has the potential to improve productivity and learning, it must be used with care to ensure that students continue to develop critical skills**. Simply giving every student a ChatGPT account will not automatically improve learning, and careful design, such as that used in GPT Tutor, is essential to ensure effective integration of AI in education**.

Based on these considerations, **"digital literacy" initiatives for the broader public or AI curricula for schools and universities should cover four core areas** (Table 1).[37] First, the course should explain the core technical aspects of AI, including the principles of machine learning, classifiers, transformers, and other foundational technologies. While it would be too much to ask to equip all students with the knowledge needed to potentially become developers of AI systems, the course should convey the basic concepts and foundational knowledge in an intuitive and exploratory fashion. Secondly, the course should also cover ethics, risks, and responsibilities to nudge a critical understanding of AI, including considerations such as involuntary plagiarism and algorithmic bias. This part of the course would encourage students to think critically about the implications of AI technologies, who benefits from them, and their unintended consequences. Thirdly, a more practical part of the curriculum should be geared towards understanding how AI interfaces with various types of work and everyday applications. Here, researchers and students should reflect on how individuals are likely to use AI technologies, which is crucial for becoming informed users of AI in both professional and personal lives. In addition, as AI increasingly takes over routine tasks, basic programming skills, such as understanding Python, may become more relevant.[38] This knowledge will enable users to customise AI tools to their needs and understand the underlying processes. Finally, it is crucial to adapt school curricula to emphasise essential skills, or what Azeem Azhar calls "ur-skills", that are essential for the effective use of AI tools.[39] These core skills include problem-solving skills such as analytical thinking, critical and textual analysis, quantitative and statistical literacy, and creative thinking. In addition, collaboration and judgement skills, including task delegation, are essential, too. As Azhar notes, students may currently acquire these competencies incidentally, but they should be integrated as fundamental, rigorously evaluated components of education in a timely manner to better prepare students for a future shaped by AI.[40]

Tab. 1: Proposal for a four-pillar AI literacy curriculum

Pillar	Questions	Suggestions for content	Learning goals
Core technical aspects	- What are the core technical elements behind "AI"?	- Introduction to Machine Learning literature on intuitive basis	- Understand basic components of AI and machine learning
	- How do classifiers and transformers work?	- Overview of classifiers and transformers and how they work	- Gain foundational knowledge of AI technologies
	- What are the principles guiding AI development?	- Exploration of AI foundational knowledge	- Explain AI concepts in an intuitive and exploratory manner
Ethics & risks	- What ethical considerations are associated with AI?	- Discussion on AI ethics and responsibilities	- Reflect on who benefits from AI and its broader impact
	- What are the risks of AI?	- Examination of algorithmic bias and plagiarism issues; case studies on ethical dilemmas in AI	- Identify unintended consequences of AI

[37] For the three first pillars, see also the interview with Victor Lee, a professor at Stanford's Graduate School of Education, in: Mui (2024), How to teach in a world run by AI, Politico. For the last pillar, see: Azhar (2024), Seven exponential policies for a new government (exponentialview.co).

[38] Loviscach (2024), Missing Link: Die GPT-fizierung des Studiums, Heise.

[39] Azhar (2024), Seven exponential policies for a new government (exponentialview.co).

[40] Azhar (2024), Seven exponential policies for a new government (exponentialview.co).

	- How can AI be used responsibly?	- Ethical frameworks and critical thinking exercises	- Learn principles of "trustworthy AI" and FAIR datasets
Practical applications	- How is AI used in work?	- Survey of real-world applications of AI	- Understand the practical uses of AI in different sectors
	- How can one use AI in the daily life?	- Showcasing, simulating, and training how to effectively integrate AI into daily tasks	- Become informed users of AI in professional/personal contexts
	- What skills are needed to interact with AI?	- Basic Python competencies	- Develop computational skills for effective use and fine-tuning of AI tools
Core problem-solving abilities	- What are general problem-solving abilities for using AI?	- Skills such as analytical thinking, critical and textual analysis	- Practice understanding of AI systems in different contexts
	- What are general problem-solving abilities for interacting with AI?	- Skills such as quantitative and statistical proficiency and creative thinking	- Practice task delegation and evaluative judgment in different contexts

Source: own table, based on cep research.

Finally, generative AI has the potential to **overcome language barriers that often impede the global dissemination and collaboration of research**. Using natural language processing, new AI writing assistants are rapidly breaking down language barriers by providing accurate translations, thus facilitating cross-cultural understanding by helping people from different cultural backgrounds understand each other's customs, traditions, and points of view.[41] Technologies like DeepL and ChatGPT play a pivotal role in this context. However, nuanced expressions such as humour and political subtleties still pose challenges for automated translation. At the moment, innovations in video AI applications are pushing the boundaries further by enabling lip-synced speech in multiple languages and generating videos in which the speaker's lip movements are synchronised with the spoken sounds, enhancing the effectiveness of cross-lingual communication. This burgeoning field of increasingly AI-driven voice services is projected to reach a market value of $72.2 billion by 2027.[42] Still, many studies exploring the use of digital in language learning have shown mixed results. For instance, while chatbots can play different roles in language learning, such as interlocutors and simulators, their unnatural responses, diminishing motivational impact over time, and potential to increase cognitive load still pose significant challenges.[43] Other research finds that AI tools can improve the language learning experience in comparison to "human peers".[44] Despite some present technological and psychological limitations, AI-powered translation and summarisation tools can thus make research accessible to a wider audience, facilitating greater participation in scientific discourse and collaboration across linguistic boundaries.

3 Con: Generative AI could exacerbate socio-economic inequalities

The integration of generative AI into education, teaching, and research raises a number of practical[45] and ethical considerations that need to be carefully addressed to ensure academic integrity and the responsible use of AI technologies. One main ethical concern is related to the behavioural effects on researchers and students, namely the potential for AI to inadvertently promote academic dishonesty.

41 For this argument, see: Küsters and Stockebrandt (2023), Digitale Brücken: Wie KI die europäische Integration vorantreibt (commongroundeurope.eu).

42 GVR (2024), Translation Services Sourcing & Cost Intelligence Report, 2030 (grandviewresearch.com).

43 Huang et al. (2022), Chatbots for language learning—are they really useful? A systematic review of chatbot-supported language learning, J. Comput. Assist. Learn. 38, 237–257.

44 Fryer et al. (2017). Stimulating and sustaining interest in a language course: an experimental comparison of Chatbot and human task partners. Comput. Hum. Behav. 75, 461–468, ScienceDirect.

45 E.g., the immense volume of scientific data, such as from telescopes and satellites, can reach petabytes, posing challenges for data sharing and interoperability of AI tools. See: The Royal Society (2024), Science in the Age of AI, Report.

For instance, some researchers may over-rely on AI-generated content without proper attribution or critical engagement with the material.[46]

Fig. 6: Estimated increase in AI-written papers

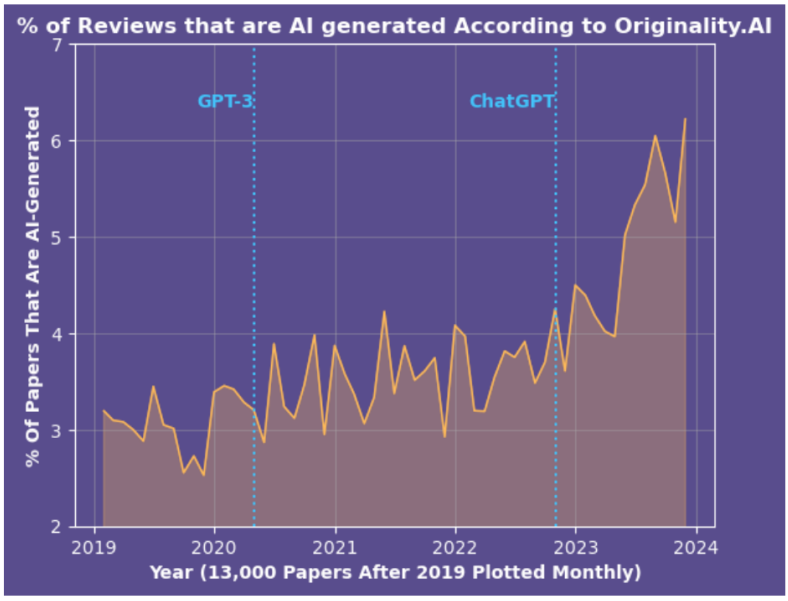

Source: Taken from: Akram (2024), "Quantitative Analysis of AI-Generated Texts in Academic Research". Note: Percentage of paper that are likely generated by AI after year 2019 scored by originality's detection tool.

To illustrate, experiments have shown that generative AI can simulate student essay writing in ways that are undetectable to teachers, while the latter are overconfident in identifying sources.[47] Further work has shown that AI-generated exam submissions go largely undetected and receive higher marks than real student work, with only 6% of ChatGPT answers submitted by 33 fake psychology students flagged by examiners.[48] A recent study developed a method for detecting intentionally fabricated content that academic organisations use to publish online, with a 98% success rate, and found a clear increase in papers written by AI, particularly after the release of ChatGPT in November 2022, rising from 3.61% to 6.22% in one year (Figure 6).[49] In May 2024, Wiley publishing house discontinued 19 scientific journals to mitigate against **systemic manipulation of the publishing process by AI**.[50] Overall, this evidence highlights the urgent need for clear guidelines on the ethical use of AI in academic work, delineating the boundaries between legitimate assistance and misconduct.

The content generated with the help of generative AI may also impact the education system on an institutional level, as quality-issues can already be seen in **peer review processes that were corrupted by generative AI tools**. A study from Stanford University analysing over 146,000 peer reviews for major computer science conferences suggests that up to 17% of reviews may have been significantly altered

[46] See: Barnett (2023), ChatGPT Is Making Universities Rethink Plagiarism, Wired.
[47] Fleckenstein et al. (2024), Do teachers spot AI? Evaluating the detectability of AI-generated texts among student essays, Computers and Education: Artificial Intelligence 6.
[48] Scarfe et al. (2024), A real-world test of artificial intelligence infiltration of a university examinations system: A "Turing Test" case study, PLOS ONE 19(6): e0305354.
[49] Akram (2024), Quantitative Analysis of AI-Generated Texts in Academic Research: A Study of AI Presence in Arxiv Submissions using AI Detection Tool, Arxiv.
[50] Claburn (2024), Wiley shuts 19 scholarly journals amid AI paper mill problem, The Register.

by LLM chatbots.[51] This is highly problematic, as sound peer review is a cornerstone of modern academia. The AI-influenced texts were identified using a novel method that detects an increased frequency of certain adjectives, such as "commendable" and "innovative", which are used more frequently by AI than by humans. Using similar estimation methods, in a systematic, large-scale analysis of 950,965 papers published between January 2020 and February 2024, the same team found that higher levels of LLM modification are associated with papers whose first authors post preprints more frequently and papers of shorter length,[52] suggesting an over-emphasis on quantity than quality. Overall, the findings raise concerns about the transparency and accountability of the use of AI in scientific publishing, particularly when texts are produced under time constraints.[53]

What does this imply for the academic writing process of the future? **A complete automation is not to be expected**, despite all dystopic visions, because good writing and effective communication require something meaningful and original to be expressed – a quality that AI, despite its capabilities, cannot (yet) replicate.[54] For instance, professional editors and experts from across disciplines recently raised concerns that generative AI could promote a homogenised, predominantly English-language discourse that overshadows regional and cultural specificities in science communication and might reduce the diversity of science narratives.[55] In addition, the model's lack of critical thinking and contextual understanding limits its role to assistive functions in writing, such as linguistic support or brainstorming, rather than substantive content creation. In light of these systems' potential to create so-called hallucinations (a euphemism for errors), Amanda Alvarez from the Finnish Center for Artificial Intelligence wonders "whether generative AI is in fact completely antithetical to the very purpose of their work".[56] Based on this, it seems likely that human insight and creativity will remain indispensable in producing effective and reflective research, despite the appeal of AI's technical support.

Next, AI models may inadvertently perpetuate biases present in their training data, leading to distorted or discriminatory results in their outputs. For example, a recent study highlights the subtle but pervasive issue of covert racism, manifested as dialect bias, in these models.[57] It empirically demonstrated that LLMs tend to assign negative stereotypes to speakers of African American English to an extent that exceeds historical human biases. The research shows that language models can inadvertently propagate these biases by making prejudicial assumptions about individuals based on their speech patterns, leading to discriminatory outcomes such as suggesting less prestigious jobs or harsher legal penalties for speakers of African American English. Moreover, attempts to correct these biases, such as human feedback training, were ineffective and sometimes even exacerbated the disparity between overt and covert biases. These findings underscore the imperative to promote AI literacy among users, enabling them to critically evaluate and question the outputs of AI systems, and to advocate for the development of AI models that are interpretable and fair.

More generally, non-English LLMs face significant challenges that affect their development and utility. This is evident, for instance, in the nascent state of Vietnamese LLMs, which are predominantly commercial and lack rigorous evaluation. Researchers have highlighted that current assessment methods,

---

51 Liang et al. (2024), Monitoring AI-Modified Content at Scale: A Case Study on the Impact of ChatGPT on AI Conference Peer Reviews, Arxiv.
52 Liang et al. (2024), Mapping the Increasing Use of LLMs in Scientific Papers, Arxiv.
53 See the interviews in: Chawla (2024), Is ChatGPT corrupting peer review? Telltale words hint at AI use, Nature.
54 See also: Citi GPS (2024), Human skills to thrive in the age of AI.
55 Nature Editorial (2024), AI is no substitute for having something to say, Nat Rev Phys 6, 151.
56 Alvarez et al. (2024), Science communication with generative AI, Nat Hum Behav 8, 625-627.
57 Hofman et al. (2024), Dialect prejudice predicts AI decisions about people's character, employability, and criminality, Arxiv.

mainly multiple-choice questions, do not reflect real-world applications, leading to public mistrust and under-utilisation in Vietnam.[58] This scepticism undermines the democratisation and accessibility of technology, as local people doubt the relevance of LLMs to their needs. In addition, barriers to access, such as the requirement for a phone number and payment fees that are prohibitively expensive relative to local living costs, limit the availability of advanced models such as GPT-4 in non-English speaking countries. These barriers contribute to wider technological and economic disparities, as non-English speaking countries face potential delays in productivity. The significant challenges faced by AI models in low-resource languages can now be studied using the IrokoBench benchmark recently developed by the Masakhane project.[59] This benchmark evaluates AI performance in 16 African languages in a structured way, covering different regions such as West, East, Southern and Central Africa, as well as English and French.[60] The results show that proprietary models (such as those underpinning ChatGPT), while superior to open models, still suffer significant performance degradation due to limited monolingual web data for African languages. However, the issues of bias and discrimination inherent in generative AI are currently difficult to address through regulatory requirements.[61] Therefore, the best way forward at this time is to **conduct more extensive testing across multiple languages to improve the capabilities of AI models in low-resource settings**.

A more structural concern about the integration of AI into education and research relates to **asymmetry and market concentration in access to and development of state-of-the-art generative AI models and services**.[62] Recent progress in AI comes mainly from the vast data and computing power concentrated in a few large technology companies.[63] According to the latest figures, industry is currently leading in AI innovation, having produced 51 notable machine-learning systems last year, significantly outpacing academia, which contributed only 15.[64] The nature of today's generative AI is thus deeply intertwined with the resources and operations of these companies, resulting in a growing dependence that gives a small group of tech giants excessive control over societal processes, potentially including research. A report by the UK's Competition and Markets Authority shows that a small number of dominant technology companies are increasingly influencing the markets for foundational AI models by controlling both the development of basic technologies, including resources such as computing power, data, and talent, and the distribution of models through major platforms and applications.[65]

This extent of centralisation and market power extends to shaping the trajectory of AI advances and influencing academic research in the field. For instance, the black-box nature and proprietary limitations of AI tools hinder the reproducibility of AI-based research. Issues such as inadequate documentation, limited access to essential infrastructure, and lack of accountability make it difficult for

---

58 See: Kannan (2024), Improving Equity and Access to Non-English Large Language Models, Stanford.

59 See: Adelani et al. (2024), IrokoBench: A New Benchmark for African Languages in the Age of Large Language Models, Arxiv.

60 In particular, IrokoBench includes three main tests: AfriMGSM for primary school mathematics, AfriMMLU for multiple choice questions in various subjects, and AfriXNLI for sentence classification.

61 For instance, legal experts argue that many cases of algorithmic bias, particularly those involving black-box models whose decision-making processes are opaque, are difficult to capture under the category of direct discrimination. This poses significant challenges for the use of generative AI in critical sectors such as education and research, where fairness and equity are paramount. The impact of protected characteristics in AI-driven decisions requires more robust legal and ethical scrutiny. See: Adams-Prassl et al. (2022), Directly Discriminatory Algorithms, Modern Law Review.

62 See: Küsters and Kullas (2024), Competition in Generative Artificial Intelligence, cepInput No 6.

63 OMI (2024), Civil Society Groups Urge UK to Investigate Microsoft's Monopolistic Partnership with OpenAI.

64 See: AI Index Steering Committee (2024), "The AI Index 2024 Annual Report," Stanford University, HAI_AI-Index-Report-2024.pdf (stanford.edu).

65 CMA (2024), CMA outlines growing concerns in markets for AI Foundation Models.

independent researchers to scrutinise, verify, and replicate experiments.[66] As a result, the tech sector's control over AI research and knowledge creation puts critical scholars and advocates in a precarious position and risks depriving them of essential insights into the impact of AI and the industry behind it.[67] Therefore, policymakers need to support open-source models and avoid restrictive licensing schemes for developing generative AI.[68] Open-source AI models allow researchers from around the world to use cutting-edge technologies without the need for extensive infrastructure or funding and reduce the risks of concentrated economic and political power in a few generative AI platforms. This democratisation can level the scientific playing field, a point taken up below.

Finally, there are also concerns regarding the data protection of researchers, as generative AI relies on large datasets that often contain sensitive information.[69] There is a risk that personal data of students or teachers used to train generative AI could be exposed, leading to data breaches and unauthorised data extraction.[70] To mitigate these risks, institutions and AI providers should implement strong data protection measures, such as data anonymisation and de-identification, to ensure that individuals' privacy and confidentiality are maintained. In addition, research organisations should limit the use of proprietary data during AI training, instead relying on synthetic datasets or pre-trained models to prevent data breaches. Regular "red teaming" exercises, which simulate potential security breaches, can help identify and address vulnerabilities in AI systems, thereby better protecting personal rights.

Still, the effectiveness of ethical guidelines and privacy measures for AI in education is fundamentally limited by the fact that **AI security is not an inherent property of the model**.[71] LLM applications can easily be repurposed for malicious purposes because an adversary or student can use a model out of context, bypassing any security measures encoded in the model itself. These models, by their very nature, lack the ability to recognise the nuances of their context of use, such as distinguishing between legitimate academic use and non-malicious misuse, such as students cheating on their homework. Consequently, it is not feasible to delegate the entire responsibility for addressing misuse to the developers. This highlights the need for a more holistic approach to AI safety in education, including more AI literacy courses and ongoing collaboration between educators, AI developers, and policymakers.

4 Policy recommendations

To realise the potential of generative AI in education and research while addressing the associated practical and ethical challenges outlined above, policymakers and actors in the provision of research and education, such as university leaders and publishing houses, should focus on the following areas.

First, tailored "AI literacy curricula" should be developed as quickly as possible that can complement existing courses. As AI scholar Gary Marcus pointed out, the rapid development of AI has outpaced education systems, leaving the public largely uninformed about the risks and benefits of AI.[72] As a result, there is an urgent need to improve AI literacy to prevent scholars and students from being exploited and unprepared. To address this, policymakers should mandate and fund AI education at all levels, including schools and adult education programs. As described in more detail above (section 2),

[66] The Royal Society (2024), Science in the Age of AI, Report.
[67] Whittaker (2021), The Steep Cost of Capture, Interactions 28, 6, 50–55
[68] Küsters (2023), The Future of AI: Open-Source Utopia or Licensed Monopoly? - (commongroundeurope.eu).
[69] Steevens (2024), ChatGPTs Memory-Funktion: Jetzt wird OpenAI richtig dreist, Heise.
[70] Borkar (2023), What can we learn from Data Leakage and Unlearning for Law?, Arxiv.
[71] For this argument, see: Narayanan and Kapoor (2024), AI safety is not a model property (aisnakeoil.com).
[72] Marcus (2024), Opinion | How to Protect Americans From the Many Growing Threats of AI - POLITICO.

a comprehensive "digital literacy" or AI curriculum should cover the technical foundations of AI, such as the principles of machine learning, while also introducing ethics, risks, and societal implications to encourage critical thinking about the impact of the technology. In addition, the curriculum should include practical applications of AI to encourage reflection on its use in different professional and personal contexts and train general problem-solving abilities. Overall, governments and educational institutions should therefore allocate resources to the integration of AI technologies, including the development of AI literacy programmes and the development of open-source AI tools that are accessible and equitable. In this context, it is crucial to establish collaborative frameworks that bring together educators, researchers, AI developers, and policymakers to share best practices, develop standardised ethical guidelines, and address the evolving challenges associated with generative AI in academia. AI developers should be required to contribute to the creation of AI-related curricula and fund public service announcements that highlight the risks associated with their products.[73]

Going beyond adding specific AI courses, **we need a larger paradigm shift in educational practices, from traditional task-oriented educational models to a competency and process-oriented approach that fosters lifelong learning, creativity, and democratic values**. What does this mean? Traditional task-based education is defined by its focus on specific, discrete tasks or assignments that students must complete. The primary aim is to measure the ability to perform certain tasks correctly and in a certain (quick) time, which leads to an emphasis on memorisation and repetition, assessment through standardised tests, a curriculum focused on specific subjects and pre-defined learning outcomes, and, by implication rather than by design, limited scope for creativity and critical thinking. This approach is particularly prevalent in China due to the legacy of Confucian thinking.[74] In contrast, process-oriented education prioritises the development of broader competencies and the processes by which students learn and actually apply knowledge. It focuses on understanding underlying concepts and principles, encourages critical thinking, problem solving and creativity, and uses a variety of methods, including project-based learning, portfolios, and reflective practice. In this regard, current research and discussion about AI in education "focuses on the digital infrastructure without considering the pedagogical strategies to facilitate the inclusion of these technologies in the educational environment".[75] A recent position paper by participants in Edunautika, a barcamp focused on modern pedagogy in the digital age, rightly criticised the current use of AI in education, arguing that it often entrenches traditional, rigid learning frameworks rather than fostering more innovative and flexible educational practices.[76] Using AI solely to perpetuate task completion and reinforce the memorisation of fixed knowledge contradicts AI's potential to enable more personalised and exploratory forms of learning, the authors of the Edunautika paper argue. They sketch an educational paradigm that uses AI to create open, engaging, and learner-driven environments, where the focus is on the process of learning itself. As shown in some of the studies quoted above (section 2), AI can indeed be used to enhance creative problem solving and facilitate the understanding of complex concepts such as legal doctrines. However, utilizing these possibilities would require a fundamental redesign of learning spaces to accommodate continuous feedback loops, reflective practices, and collaborative learning. To give an example, assessment methods could be developed to assess not only knowledge but also AI-relevant skills such as critical

[73] As proposed by: Marcus (2024), Opinion | How to Protect Americans From the Many Growing Threats of AI - POLITICO.
[74] Wang (2022), Resurgence of Confucian education in contemporary China: Parental involvement, moral anxiety, and the pedagogy of memorization, Journal of Moral Education, 52(3), 325–342.
[75] Triberti et al. (2024), "Better than my professor?" How to develop artificial intelligence tools for higher education, Front. Artif. Intell.
[76] Samel et al. (2024), KI-Impulspapier (ebildungslabor.github.io).

thinking, creativity, and collaboration. Likewise, teacher training should focus on facilitating process-oriented learning and the effective integration of technology in the classroom. In short, continuous professional development focused on innovative pedagogical practices and AI literacy is needed.

Next, improving equity and access in the development of LLMs for underrepresented languages as well as mitigating data fabrication, poisoning, or contamination. Research highlights the importance of building high-quality, diverse datasets that are tailored to the unique linguistic structures of each language, ensuring the absence of toxic or grammatically inconsistent content.[77] In particular, it demonstrate the effectiveness of using so-called paired datasets (such as Wikipedia articles that exist in both English and the target language) for speeding up the fine-tuning process for non-English LLMs. More generally, open-source LLMs, paired high quality datasets, and better evaluation frameworks are needed to improve the accessibility and performance of LLMs for all languages. Moreover, to fully meet the challenges of biases and discrimination discussed above, data curators and information managers are critical to maintaining data quality and mitigating risks such as data fabrication, poisoning, or contamination.[78]

Complementing these improvements with respect to (academic) research, policymakers could support increased equity and access to LLMs through subsidised licensing for historically disadvantaged education institutions and universities, open-source initiatives, or partnerships between technology providers and educational institutions. For instance, incentivising or even mandating open foundation models, i.e. models with widely available weights, is highly relevant because "closed model developers exert greater power in defining and restricting use cases they deem unacceptable, whereas downstream consumers of foundation models can better make these decisions for themselves with open models".[79] In the context of Europe, a substantial step would be the financing and creation of a large European language model for science, emphasising open-source accessibility for research and adaptability for the education sector, while ensuring security to mitigate the risk of misuse.[80] It should meet ethical and legal standards, including compliance with the EU's new AI Act, focusing on transparency and societal impact, and should be mainly trained on academic texts.

Finally, concrete internal guidelines are required for all education-related institutions and potentially regulatory measures that oversee the ethical use of AI in education, protect privacy, and ensure transparency and accountability in AI-assisted research and education. A good example is the recent guidelines from the BBC, the New York Times, and others for using generative AI in the newsroom, which are closely aligned with their respective ethical journalism handbooks and revolve around notions of transparency, human guidance, and complementing, not replacing, human input.[81] Similar internal guidelines for European research institutes and schools should closely follow the ethics guidelines for "trustworthy AI" developed in 2019 by the High-Level Expert Group on AI, covering, e.g. transparency and enhanced societal impact assessments.[82] In addition, these guidelines should explicitly take into account the inherent trade-offs of using or avoiding generative AI, i.e. the implicit opportunity costs. For example, as the UN advisory body on AI has noted: "Leveraging AI to improve access to

[77] Truong et al. (2024), Crossing Linguistic Horizons: Finetuning and Comprehensive Evaluation of Vietnamese Large Language Models, Arxiv.
[78] The Royal Society (2024), Science in the Age of AI, Report.
[79] Bommasani et al. (2023), Considerations for Governing Open Foundation Models, p. 5.
[80] Praas (2024), Launching an 'AI moonshot' to develop a European large language model is the game changer that Europe needs, CEPS.
[81] See: Ethical Journalism - The New York Times (nytimes.com).
[82] See: High-Level Expert Group on AI (20199; Ethics guidelines for trustworthy AI.

education might raise concerns about young people's data privacy and teacher agency. However, in a world where hundreds of millions of students do not have access to quality education resources, there may be downsides of not using technology to bridge the gap."[83] Addressing such trade-offs will benefit from governance mechanisms that enable to share lessons learnt as well as from empirical studies that examine the effects of using generative AI in educational settings.

5 EU initiatives in cross-border education reform

In the European Union (EU), education remains primarily a national competence, with each Member State responsible for the structure and content of its education systems. Formally, EU Member States are tasked with implementing policies and curricula that respond to the unique cultural, social, and economic contexts of their populations. They retain autonomy over key educational decisions, including teacher training, pupil assessment, and resource allocation. However, the EU facilitates cross-border cooperation, the exchange of best practices, and the funding of initiatives that promote innovation and inclusiveness in education. Accordingly, the EU can act as a "catalyst for national reforms and investments"[84] in supporting and coordinating efforts to improve the quality and accessibility of education across the continent, as illustrated, for instance, by its higher education package from 2022. As argued in this section, this collaborative approach is particularly important as AI reshapes research and education processes.

As shown above, the integration of AI in education, if well tailored and adapted, can enhance personalised learning for students, streamline administrative tasks for teachers and professors, and provide advanced analytical tools for researchers. Therefore, aligning the next European Commission's strategic focus on digital skills and updating education for AI and digitalisation is crucial to maintaining its competitive edge in the global market. To date, the EU has mainly sought to boost productivity through significant investment in hardware, i.e. frontier technologies such as supercomputing, semiconductors, the Internet of Things, quantum computing, and space technology. Leadership in AI innovation is a more recent component of this strategy, with initiatives such as ensuring access to tailored supercomputing capacity for AI start-ups and industry. However, a more holistic and strategic approach to soft factors such as education, skills, and digital literacy is needed, along the lines of the argument outlined above. According to the State of the Digital Decade 2024 report, only 55.6% of the EU population has basic digital skills and the current pace of growth in ICT professionals is insufficient to meet the 20 million target by 2030, highlighting the urgent need to accelerate existing efforts to develop digital skills.[85] The EU focus on digital skills is also linked to, and justified by, the broader objective of promoting inclusion and participation in all European regions, in particular to bridge the digital divide between urban and rural areas and between different socio-economic groups.[86]

In this respect, there are some promising signs in the recent Political Guidelines 2024-2029 presented by President Ursula von der Leyen, which explicitly focus on education and skills development for the coming digital transformation.[87] The establishment of a so-called "Union of Skills" will focus on investment in lifelong learning, skills retention, and the recognition of different types of training across the

[83] UN AI Advisory Body (2023), Governing AI for Humanity, interim_report.pdf (un.org), p. 12.
[84] Murphy and Claeys-Kulik (2022), The European Strategy for Universities: a catalyst for national reforms and investments? (theparliamentmagazine.eu).
[85] European Commission (2024), Communication, State of the Digital Decade 2024, COM/2024/260 final, p. 13.
[86] European Commission (2024), Communication, State of the Digital Decade 2024, COM/2024/260 final, p. 13.
[87] For this paragraph, see: Von der Leyen (2024), Political Guidelines for the Next European Commission: 2024-2029, p. 12.

Union. Embedding lifelong learning in education and supporting the career prospects of teachers will be central to these efforts. A strategic plan for STEM education, as proposed in the Guidelines, will address the underperformance and shortage of qualified teachers in science, technology, engineering, and mathematics, with a particular focus on increasing the participation of women in these fields. Vocational education and training (VET), which prepares individuals for work and equips them with the digital skills needed by businesses, is also listed as a priority. This will include a European strategy for vocational education and training to increase the number of graduates in secondary education, but will require more precise information on how this preparation of individuals should take place and on which competences it should focus. Skills funding must be better aligned with labour market needs, focusing on sectors critical to the green and digital "twin transition". A skills portability initiative will ensure that high-quality skills are recognised across the EU, regardless of where and how they were acquired. This could help the proliferation of AI talent.

So far, many of these initiatives, as outlined in the Political Guidelines, are simply labels and repackaging of older programmes, without substantial funding or publicity. However, bringing them together under a more comprehensive, unified umbrella, providing the initiatives with sufficient funding, and initiating a deeper dialogue with EU Member States with their education and labour market competencies could indeed help ensure that Europe's workforce is prepared for the evolving AI future, including the likely changes to the labour market.[88] Not mentioned in the Guidelines are the European Digital Innovation Hubs (EDIHs), which could likewise support the digital transformation of the European public by providing training and digital maturity assessments for SMEs. To date, this network has grown to more than 200 hubs in 90% of European regions and has made a significant contribution to improving digital skills.[89] As highlighted in the State of the Digital Decade 2024 report, cooperation and sharing of best practices between Member States will play another important role in achieving the ambitious skills targets set for 2030.[90] This is also true with respect to more general labour policies.

Going forward, a larger part of the EU's digital transformation strategy should therefore focus on addressing the digital skills gap and promoting process-oriented learning to ensure that the workforce can meet the demands of an increasingly AI-driven world. To this end, EU programmes and initiatives can build on the multi-faceted, four-pillar "AI literacy" curriculum outlined above. Moreover, such a comprehensive approach would not only prepare Europe's workforce for the evolving demands of the digital economy, but also reduce digital divides and foster a more inclusive and resilient digital society.

6 Conclusion

Based on current prototypes and emerging empirical evidence, the use of generative AI in education has significant potential to accelerate data analysis, streamline writing processes, stimulate creativity, and improve accessibility to learning. This advance, as explored in this paper, offers an opportunity to transform outdated task-oriented educational paradigms by personalising learning experiences and democratising access to high-quality educational tools. However, simply giving every student a ChatGPT account will not automatically improve learning, as experiments show that careful design is essential to ensure effective integration of AI in education. Moreover, the widespread use of large

[88] For predictions, see the literature review in: Küsters and Poli (2024), Resisting or Rebooting the Rise of the Robots?, cepStudy No. 2/2024.
[89] European Commission (2024), Communication, State of the Digital Decade 2024, COM/2024/260 final, p. 9.
[90] European Commission (2024), Communication, State of the Digital Decade 2024, COM/2024/260 final, pp. 17f.

language models (LLMs) poses challenges in terms of privacy, security, and the perpetuation of existing prejudices.

At the core, realising the opportunities offered by LLMs depends on sufficient digital and AI literacy across society. From this perspective, it is highly problematic that EU citizens still lack in digital skills and that even within academia, AI literacy is becoming a divisive issue within scientific disciplines, potentially hindering progress. To address this, educational institutions need to integrate updated, comprehensive AI curricula that cover four core areas: the technical aspects of AI, ethical considerations and risks, practical applications, and essential core skills for the effective and reflective use of AI tools. This paper provides a concrete framework for what such an AI curriculum could look like, which could also serve as an inspiration for the various EU initiatives in cross-border education reform recently announced by President Ursula von der Leyen in her Political Guidelines for 2024-2029.

Moreover, evidence of systemic manipulation of publication processes and quality issues in peer review due to new AI tools underscore the urgent need for mandatory ethical guidelines. These guidelines should clearly delineate the boundaries between legitimate assistance and academic misconduct, ideally by giving concrete examples. In addition, it must be reflected that AI models risk perpetuating biases from their training data, leading to distorted results, and that non-English LLMs face significant development challenges. Structural concerns also arise from asymmetry and market concentration in access to advanced AI models. To mitigate these issues, policymakers should improve equity and access for underrepresented languages and address data-related challenges. Supporting historically disadvantaged institutions through subsidised licensing, open-source initiatives, and industry partnerships can further promote equity. By addressing these elements, the education sector can foster a continuous learning environment in the AI age without widening existing socio-economic divides, thus ensuring the long-term resilience of science as a cornerstone of European competitiveness.

Centrum für Europäische Politik
FREIBURG | BERLIN

Author:

Dr. Anselm Küsters, LL.M., Head of Division Digitalisation and New Technologies
kuesters@cep.eu

Centrum für Europäische Politik FREIBURG | BERLIN
Kaiser-Joseph-Straße 266 | D-79098 Freiburg
Schiffbauerdamm 40 Räume 4205/06 | D-10117 Berlin
Tel. + 49 761 38693-0