EUROPEAN UNIVERSITY ASSOCIATION

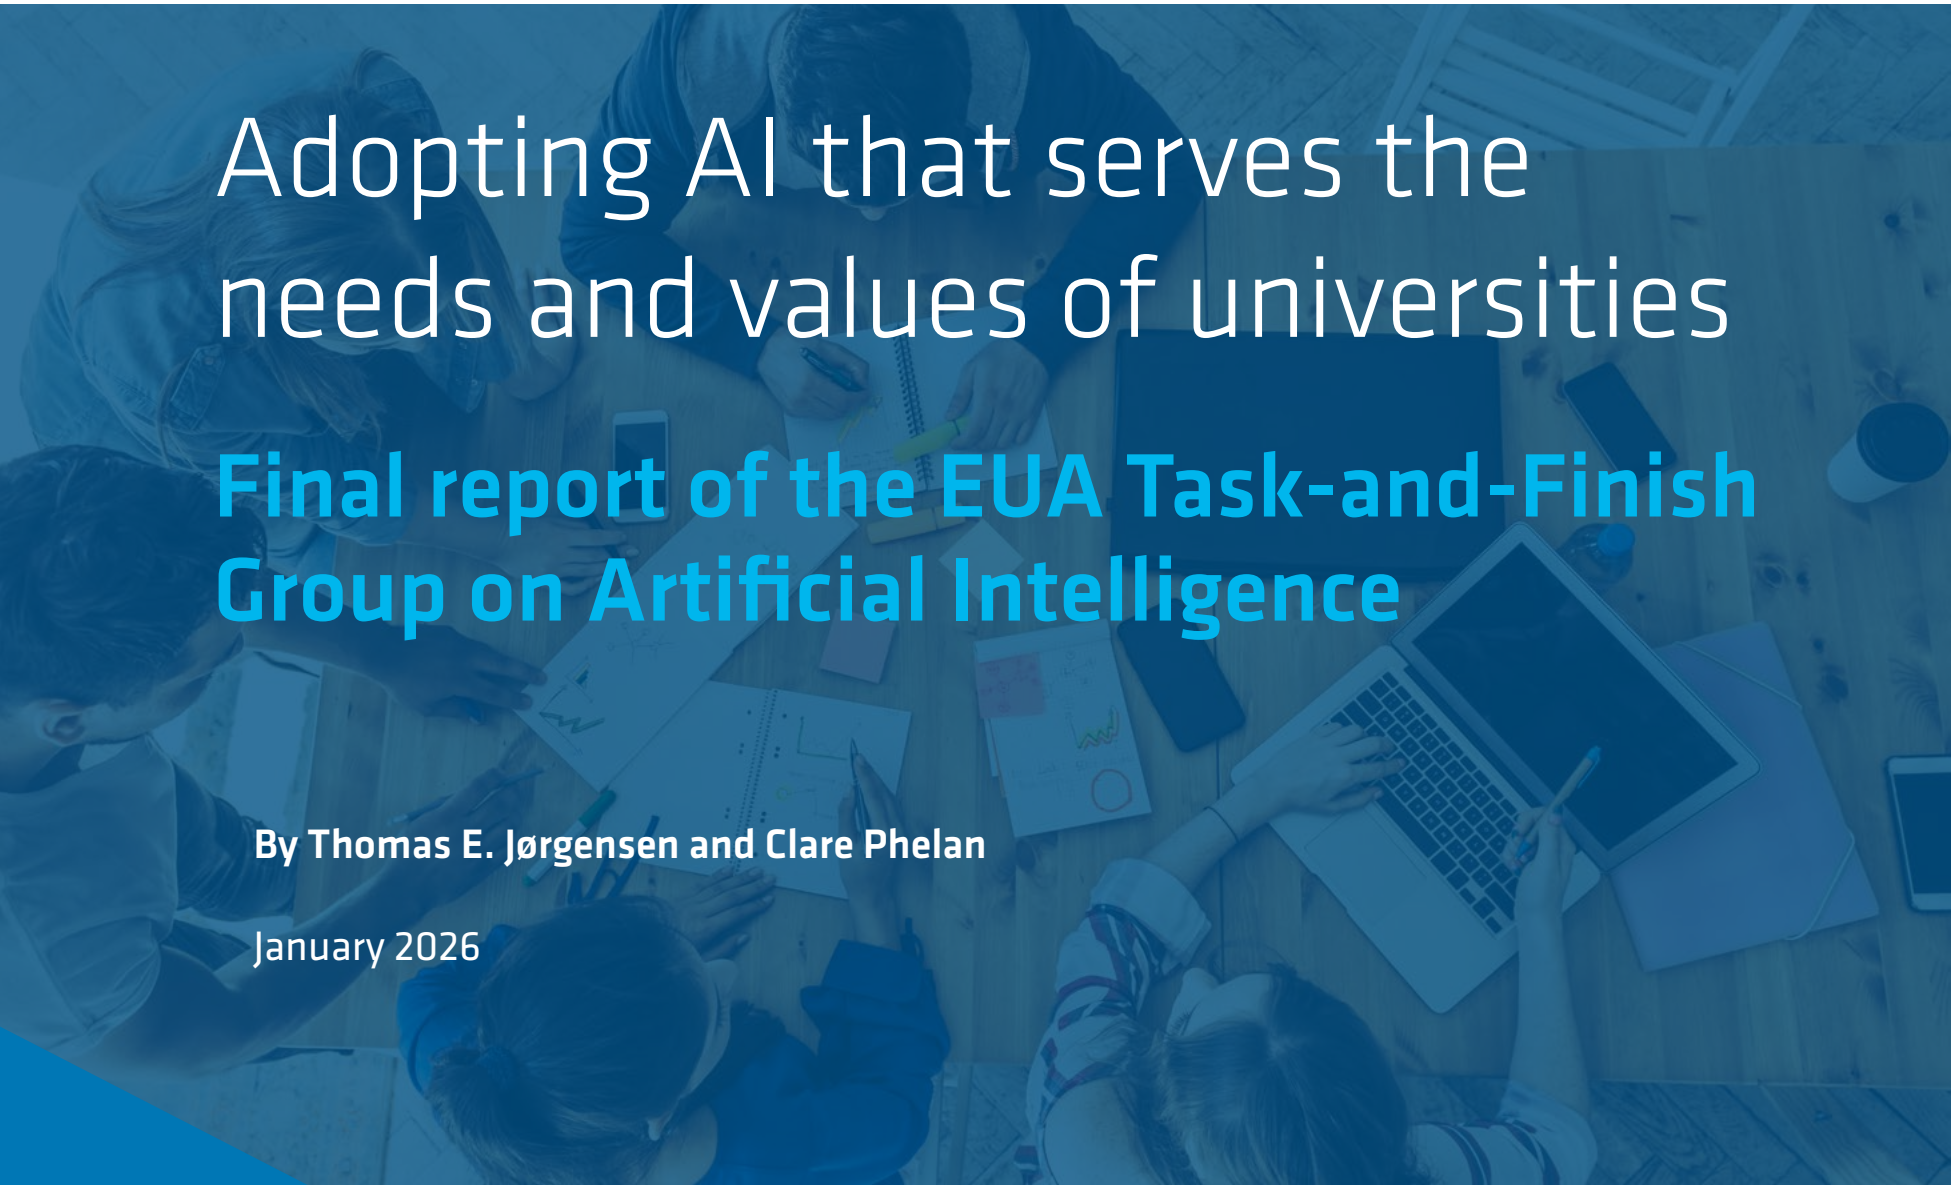

Adopting AI that serves the needs and values of universities

Final report of the EUA Task-and-Finish Group on Artificial Intelligence

By Thomas E. Jørgensen and Clare Phelan

January 2026

Contents

Introduction

Universities have come to see digital transformation as a vehicle for building institutional capacity, and therefore, worthy of sustained investment. The integration of AI, however, adds a new dimension to the transition. The rate at which technologies are advancing is disrupting this steady continuum of (digital) change, making it difficult for universities to anticipate future needs, set priorities, and commit resources accordingly.

This report focuses on the here and now. It looks at current and persisting questions for universities in five key areas – Ethics, Strategies, Training, Regulation and Sustainability – considering if and how AI is changing perceptions and practices. There are no easy answers to these questions and none are attempted in these five chapters. The report is intended as a resource for university leaders and representatives overseeing institutional change. It serves as a starting point for reflection on small adjustments that can equip university communities to approach AI mindfully.

The publication arrives at a critical time. AI systems that generate text, images, and other materials are now widely available within academic and student communities. At the same time, agentic AI – a category of AI systems designed to make decisions and carry out tasks with a degree of independence – is receiving growing attention. Alongside this, speculation continues to build about general AI, the prospect that machines could one day match human abilities across a wide range of activities. Against this backdrop, the European Union is introducing regulation that aims to anticipate and mitigate risks while enabling Europe to reap the full benefits of emerging technologies.

In this context, universities are carving out pathways towards AI adoption that best serve their needs. This publication does not provide a roadmap or recommendations; it should serve as inspiration for universities to further explore where AI adds value and where it does not. The chapters that follow touch on some of the challenges higher education institutions are contending with, as they are tasked with:

- building digital literacy while discouraging dependence;
- reimagining the classroom experience while (re)focusing attention on core academic skills;
- supporting boundary-pushing experimentation while upholding scientific integrity; and
- identifying scalable use cases while anticipating their environmental impact.

In parallel, universities are coming together to maximise their influence with policy makers and technology vendors in order to negotiate favourable outcomes for their communities.

Amid such complexity, university communities are readying to embrace change with a curious and critical mindset. This report finds that, in laying the foundations for AI integration, institutions are best guided by their core missions. By respecting the values upheld within their missions, universities can better determine where AI can improve people's everyday experiences of work and learning. Discovering the real value of technologies should be a participative endeavour, undertaken with a spirit of care and enquiry, and enabling the whole university community – leaders, learners, professional and academic staff – to play an active, conscious role in adopting AI responsibly.

What led to this report?

Following the conclusion of a two-year work programme on digital transformation, EUA turned its attention to artificial intelligence. To enhance the discourse around the integration of AI in university missions, EUA established a task-and-finish group comprising representatives of its member universities, from diverse academic disciplines, including statistics, ethics, medicine, digital humanities and computer science.

The work programme was guided by the task-and-finish group, coordinated by the EUA Secretariat and made possible through the valuable contributions of the Association's members. The Secretariat organised a series of activities with two specific aims: to create a space for members to share their principal concerns and obstacles encountered with new technologies and to showcase AI-related initiatives that they have piloted. The key work programme events were:

Spring 2024 webinar series

– *Artificial intelligence: from experimentation to institutional strategies*

– *Artificial intelligence and ethics: the place of new technologies in university missions*

Autumn 2024 workshop series

– *Artificial intelligence and universities: what futures?*

2025 EUA AI Conference

– *How universities are shaping the era of artificial intelligence*

The 2025 EUA AI Conference attracted over 200 registrations from 43 countries. The autumn workshop series was delivered with the active participation of over 45 EUA member representatives spanning 34 countries and together, the spring webinars recorded over 1500 registrations. To complement these work programme milestones, the individual perspectives of task-and-finish group members were highlighted in articles on EUA's Expert Voices blog over the course of the group's mandate.

This report distils the outcomes of the main activities and is enriched by the thoughtful contributions of event participants, moderators and speakers. The positive response among EUA members to the AI work programme is indicative of the importance of the topic across the sector, and the need for a holistic, nuanced debate about AI's adoption in universities. We hope that, by touching on some of the more salient issues, this report can usefully contribute to the conversation.

Acknowledgements

On behalf of EUA's wider community of members, the Secretariat extends thanks to members of the task-and-finish group on artificial intelligence for their commitment and expertise. This report draws heavily on the insights they generously shared over the course of the group's mandate.

The members of the group are:
- **Stefan Lang**, University of Salzburg, Austria
- **James Mackay**, European University Cyprus, Cyprus
- **Khaled Ismail**, Charles University, Czechia
- **Ekaterina Kldiashvili**, Petre Shotadze Tbilisi Medical Academy, Georgia
- **Manfred Krafczyk**, Technische Universität Braunschweig, Germany
- **Hallvard Fossheim**, University of Bergen, Norway
- **Paulo Jorge Ferreira**, University of Aveiro, Portugal
- **Štefan Stanko**, Slovak University of Technology in Bratislava, Slovakia
- **Ernest Ženko**, University of Primorska, Slovenia
- **Pedro M. Ruiz Martínez**, University of Murcia, Spain
- **Sarah Dégallier Rochat**, Bern University of Applied Sciences, Switzerland
- **Monique Mackenzie**, University of St Andrews, United Kingdom

EUA is especially grateful to Pedro M. Ruiz Martínez for chairing the group.

1. Ethics, principles and values

Risks, opportunities and a human-centric approach

The responsible use of AI in research and education has been under constant debate since universities started to deepen experimentation and test the capacity of these new technologies. The promise that AI will upend established systems and practices adds a new dimension to the conversation. However, as some have argued, AI, per se, does not radically change long-standing ethical questions within academia.[1] The real question is how to successfully adopt AI within established ethical frameworks.

Such powerful technologies can generate curiosity among users, eager to test their boundaries. However, responsible use with appropriate checks and balances helps ensure that AI adds real value for human users. Meaningful change and innovation are most likely when actions are guided by thoughtful intention rather than mere possibility, productivity, savings or convenience. Universities can foster progress by encouraging responsible experimentation and ensuring that AI is integrated with care and transparency.

In this spirit, institutions need to cultivate a sense of awareness and responsibility among users. This, in turn, can encourage reflection on how the technologies are working for the individual, the institution, the environment and for humanity in general.

> By keeping the user not just in the loop but at the centre of technological innovation and integration, universities support a wider movement to retain human agency in developing and using AI.

> Meaningful change and innovation are most likely when actions are guided by thoughtful intention rather than mere possibility, productivity, savings or convenience.

A core principle is that AI should augment and elevate human endeavours, but not control or even guide them. By keeping the user not just in the loop but at the centre of technological innovation and integration, universities support a wider movement to retain human agency in developing and using AI. This human-centric approach reinforces the need for oversight - not only to regulate generative AI systems, which have an inherent tendency for inaccuracy and hallucination, but also to uphold human values and protect citizens' rights. Embedding these values will ensure that the rights of individuals to transparency and informed consent regarding the use of their personal data by AI are fully respected.

1 Hallvard Fossheim, "Artificial intelligence: Ethical challenges for research", University World News, 30 May 2025

Ensuring justice and representativeness

As it stands, AI is typically trained on data that is widely available in society, often excluding perspectives from underrepresented cultures, disciplines, and regions. Therefore, the outputs it produces tend to replicate and reinforce stereotypes and biases that are present all around us. These could be gender and racial stereotypes where, for instance, image generations might typically depict authoritative figures as a white male, or where Western or Eurocentric perspectives are prioritised in textual responses. Analytical AI carries a similar risk of algorithmic racism and partiality. The use of AI in diagnostics, for example, can perpetuate or even amplify the bias already present in data underpinning medical studies and among members of the medical community.

Moreover, the main source of data for training many models comes from the English-speaking world, and overwhelmingly from the United States. This influence is then intensified by reinforcement learning from human feedback – even when the data workforce is based in the Global South, it is Western corporations that set the general guidelines that structure their work. It stands to reason then that output from these models will not only tend to reproduce the values inherent in these data and processes, but US-based models will also have built-in guardrails that reflect these values. For instance, US-based generative AI models will usually default to US settler colonial perspectives in accounts of US history (just as the Chinese DeepSeek is cautious with topics that are deemed problematic to the Chinese Communist Party).

Such contrasts highlight the need for European datasets, AI models and infrastructures that reflect pluralism, as promoted under the European Language Data Space and EU AI Act.

Users also need to acknowledge language disparities: generative AI is trained on widely spoken languages, and is much more accessible, therefore, for English or French speakers than for speakers of small minority languages like Sami or Frisian.

To be able to mitigate these and other biases, users first need to be able to recognise them. It will become incumbent on universities, as user communities, to ensure that staff and learners are aware of AI's inherent tendency to bias. Embedding awareness of these biases within AI-literacy training and ethics curricula can help staff and students critically assess, rather than passively accept, AI-generated outputs.

> It will become incumbent on universities, as user communities, to ensure that staff and learners are aware of AI's inherent tendency to bias. Embedding awareness of these biases within AI-literacy training and ethics curricula can help staff and students critically assess, rather than passively accept, AI-generated outputs.

This represents a first step towards countering these tendencies and working collectively to avoid the crude reproduction of stereotypes in AI-generated texts and images. Arguably, the use of AI shines a light on these biases and provides a mechanism to broaden the training corpus to include global voices and perspectives.

Protecting privacy

In the digital era, the need for privacy is strongly tied up with people's sense of justice. As individuals leave digital fingerprints through their activities, this data can generate commercial value. Additionally, exploitation of data and violation of privacy can have political, social, and even epistemic consequences – influencing how information is controlled, how power is distributed, and how individuals are profiled or governed. Increasingly mindful of this, people are seeking to keep control of their personal data and protect it from external interests. The democratisation of generative AI has, to some extent, heightened an awareness of these dangers. But within the university context, inadequate digital literacy means that the risks of feeding personal or institutional data to AI providers are not always apparent. Moreover, most of the large language models (e.g. ChatGPT, Gemini) include all user inputs in model training unless the user explicitly opts out. This might, for instance, include student papers, uploads of copyrighted academic articles, or sensitive personal data gathered via surveys. The dangers of such inputs being then replicated in future generative AI outputs are clear.[2]

To build an awareness of the risks among staff and student communities, AI training should, arguably, start with an understanding of the value of data, and the importance of keeping it in trusted hands.

Notwithstanding the ethical imperative, there is a more immediate motivation for protecting privacy: compliance with data regulation. However, navigating the often-complex regulatory landscape can present obstacles for universities. In response, some have come up with novel solutions to tackle the compliance issue. The University of Oslo (UiO) is one such example, where developers have created a walled-off version of ChatGPT to be used within the university. UiO GPT was developed via access to OpenAI's ChatGPT model, but in an environment where all data remains on servers in countries that are General Data Protection Regulation (GDPR)-compliant and will not be used for training of AI models. Data privacy and GDPR conformity were also key requirements and guiding principles when designing an inter-university solution in Austria by ACOmarket on the request of the national university association, Universities Austria.

Data generated and collected is not used to train OpenAI LLMs and remain in a closed circle, while optimised Azure AI models are deployed for specific academic purposes.

UiO GPT or ACOmarket's Academic AI are just two of many creative, homegrown solutions to protect privacy. However, not all institutions have the resources to develop similar systems, and not all learners have the possibility to subscribe to the high-end versions of the various models. Unequal access to technologies adds a social justice dimension to the AI-ethics debate. While not all AI models require advanced technical skills, many fear that, without proper checks in place, AI will continue to separate the haves from the have-nots as part of digital transformation, with the latter group possibly required to compromise on privacy and control of their data to access quality AI models. To counter this risk, universities can explore shared ethical infrastructures or public AI resources which democratise access while safeguarding privacy.

2 See, e.g. Patricia Waldron, "ChatGPT 'memorizes' and spits out entire poems", Cornell Chronicle 9 January 2024

Preserving integrity and values in academia

Another principle which is central to universities is academic integrity. Here, as in other intellectual endeavours, there is a sense that AI must be used for augmentation rather than replacement of human agency. Apart from the obvious questions surrounding learners' use of generative AI to produce essays and educators' use for assessment, there are deeper concerns related to transparency and reproducibility in the learning and experimentation process. For the most part, AI works on a non-transparent system where it is difficult, if not impossible, to retrace the steps or reasoning that led to a particular result. Such opacity complicates replication, peer review, and validation of findings. This poses serious challenges for the reproduction of scientific experiments conducted using AI technologies.

Despite these very legitimate concerns, it is generally acknowledged that AI has enormous potential to accelerate science and open up new forms of investigation.[3] But, to maximise this potential, there is need for high quality data and human oversight throughout the process. The imperative for quality data and human agency becomes even more apparent when we consider the very real possibility that generative AI could be used to produce fast but unreliable research results – a risk that is compounded by the publish-or-perish culture created by governments, research funders, rankings agencies and academic institutions.[4] This is a development that would compromise the fundamental principles of research, potentially eroding trust in science and posing serious risks to citizens and society. These risks reinforce the need to align AI-enabled research with open science standards, findable, accessible, interoperable and reusable (FAIR) data practices, and responsible-metrics initiatives such as the Coalition for Advancing Research Assessment (CoARA) and the European Code of Conduct for Research Integrity. It also reinforces the need for institution-wide training in the limitations of AI and requirements for careful human oversight. (cf. Chapter 3: Training)

Another point of contention in the debate around ethical AI is the divergence in values between those within the academic community and those developing AI systems. While it is not universally true, it is often the case that providers of AI solutions use efficiency as a main argument or selling point for embedding technology in processes. However, within an academic context, efficiency in itself is not a core value. (cf. Chapter 5: Sustainability and societal impact – Outsourcing of competence – what constitutes reasonable and responsible use?) In fact, one could argue that serendipity and the possibility of taking time to solve problems is the core of academic pursuits and that sacrificing quality for the sake of efficiency could jeopardise these much more fundamental academic values. Curiosity, reflection, and the freedom to explore are central to discovery. An overemphasis on automation can endanger these academic virtues, especially if AI is perceived as an easy answer to the constant pressure on science to be more productive.

“The imperative for quality data and human agency becomes even more apparent when we consider the very real possibility that generative AI could be used to produce fast but unreliable research results.

3 OECD (2023), *Artificial Intelligence in Science: Challenges, Opportunities and the Future of Research*, OECD Publishing
4 UK Research Integrity Office (2025) *Embracing AI with integrity. A practical guide for researchers*, UKRIO, p. 14-15

It will be critical that decisions are underpinned by a common understanding of ethics, not as fixed rules but as guiding principles.

AI adoption as a participatory exercise and whole-institutional goal

Universities managing the practical implications of AI need to consider these risks and opportunities carefully. In weighing them against each other, keeping the human at the centre should be the guiding principle. This means looking, very concretely, at where AI can augment and enhance human agency, i.e. where it benefits the individual's work or the institution's mission - and where it might mean surrendering control.

At a moment when both learners and academic staff are deepening experimentation, the real value and the viable use cases are, for many, still to be determined. In an effort to move from exploration to systematic integration, many institutions have been mapping the actual use of AI through surveys. Gaining an understanding of how AI is concretely benefiting (or disadvantaging) users is integral to the creation of institutional guidelines or policies that govern its use.

However, mapping AI use and developing guidelines to support its integration cannot be a one-off exercise, especially given the fast-developing nature of the technology and its increased embedding in the world outside the university. Rather, it should be a continuous, participatory, and proactive process. Only through open dialogue with the community (and wider society) will universities come to a conclusion on what is reasonable and productive use of AI. Moreover, it will be critical that decisions are underpinned by a common understanding of ethics, not as fixed rules but as guiding principles that encourage safe and responsible conduct.

2. Institutional strategies

Following the advent of widespread generative AI in 2022, many universities took a learn-by-doing approach. The logic for most institutions was that a pilot phase for implementation would provide time to establish how to best use the opportunities and minimise the risks. The experimentation gave way to institutional learning, but now higher education institutions face the challenge of converting experience into long-term strategy and policy.

As the sector begins to emerge from this exploratory phase and consider its longer-term approach to AI, sharing and consolidating experiences will be key.

Universities should seize the opportunity to learn from developments in other sectors, while setting out clear, sustainable, and ethical ambitions for the evolution of AI in higher education and research.

AI strategies that respect institutional missions

Universities have long seen the value in developing strategies which serve a dual purpose: to guide institutional priorities and activities and to communicate outwardly the institution's value and goals. The recent hype around AI, mostly brought about by the launch of ChatGPT in 2022, pressured universities to react immediately but likewise think strategically. Some considered fundamentally revising their strategic goals in response, but for the most part, universities have come to understand the importance of aligning AI adoption with institutions' core missions and their communities' interests.

In fact, in the few short years since generative AI became widely accessible, universities have become increasingly aware of need to keep their missions intact. Rather than uncritically following AI and all that it promises, universities should, instead, ensure the use of new technologies is aligned with their broader strategy and pedagogical goals. This alignment allows AI adoption to reinforce, not redefine, the institution's educational and societal role. As with all digital tools, the practical application of AI should, first and foremost, uphold university values and serve institutional missions.

Strategies that allow for experimentation and prioritisation

While universities increasingly acknowledge the importance of preserving agreed strategic goals, this does not – and should not – prevent them from considering their approach to AI. Many institutions have, as an initial response, articulated guiding principles on the use of AI through stand-alone strategy documents.

These broad principles continue to hold value, but the focus has shifted towards embedding AI in institutions' wider strategies. In this transition, it has become clear that, even as universities begin to consolidate their experiences, institutional strategies should continue to enable experimentation.

However, it has become equally apparent that, despite deepening appetite for exploration, it is not possible to pursue everything. To avoid hurried, short-sighted investment, universities will need to prioritise. This means identifying 'intelligent' pilots and backing initiatives that can be scaled up to provide useful services that align with each institution's strategy and add value to its mission.

For institutions with bigger capacity, this see-what-works approach will be easier. Indeed, access to funding could well be the differentiator in how institutions evolve and adapt to AI technologies. The reality is that, regardless of appetite, universities with smaller budgets will not have the same resources to develop, access or experiment with tools as their better-financed counterparts.

The same applies to institutions with restricted strategic capacity, whether this is the consequence of limited autonomy within their system or a decentralised governance model within the institution. It is no surprise, then, that examples of the strategic use of AI are more prevalent among institutions that have the capabilities and latitude to pursue AI initiatives. Similarly, the strategic use of AI is more likely seen among institutions that employ the technologies for a specific purpose, whether this be for the upskilling and reskilling of professionals or the delivery of distance learning.

This is especially evident in areas where long-term experience has been gathered in digital environments; lessons learnt from this experience can be usefully translated to comparable contexts, supporting institutions with more limited capacity to overcome their immediate obstacles or concerns in AI adoption.

One way of measuring progress and moving beyond the pilot phase is via Maturity models. A good example of such a model is the toolkit from Jisc (the UK IT provider for higher education), which aims at using AI for optimisation and automation that creates opportunities for learners and gives time for human interaction.

Regardless of a university's plans or ambitions for AI, implementing the technologies successfully will depend on (digital) capacity. Fears that AI will widen the gap between institutions' digital maturity may be legitimate. However, this could be mitigated by concerted action and clever coordination of strategies.

AI maturity toolkit for tertiary education

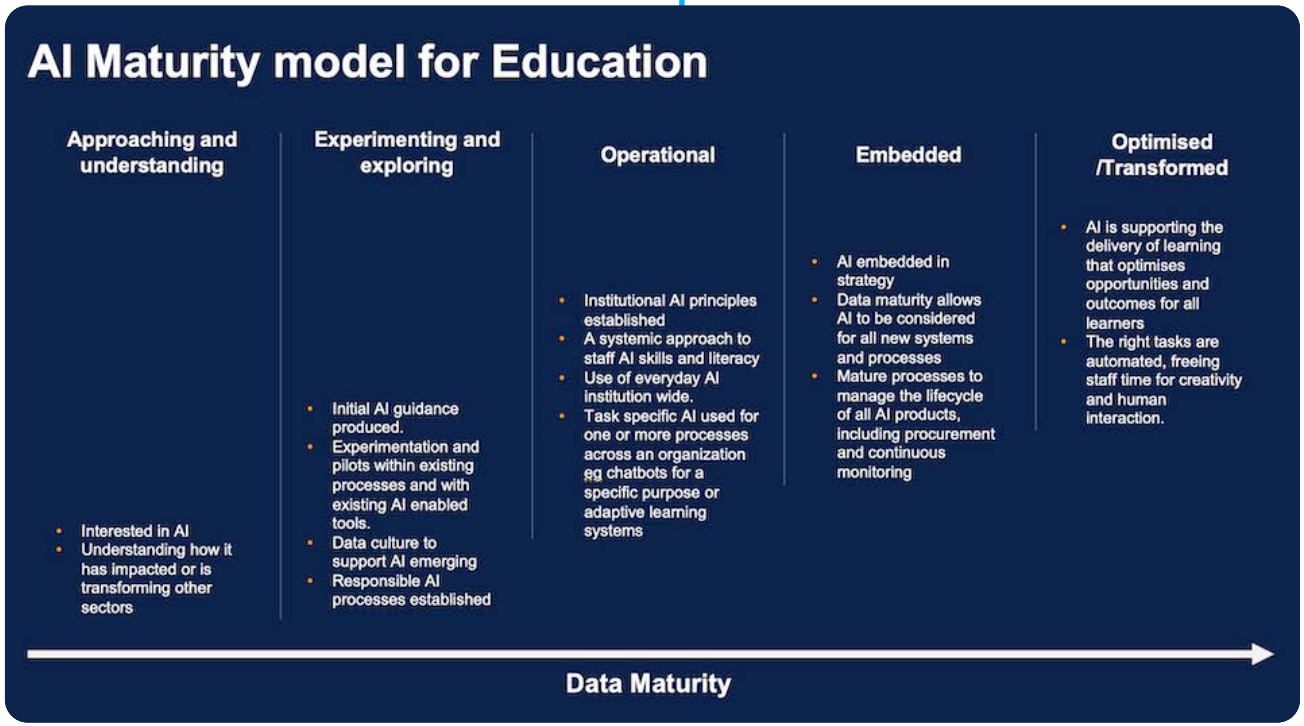

European coordination frameworks, such as the European Digital Education Hub and AI4Education initiatives, can provide valuable reference points for this alignment. Similarly, collaborative consortia or shared-infrastructure models may help smaller institutions overcome resource disparities in cases where large investments are required. In addition to the cost-saving benefits, such cooperation agreements can be useful as a vehicle for sharing experiences, negotiating with vendors and generally, strengthening institutions' visibility.

Finding complementarity to the point where each institution works to its individual capacity could well bring benefit. But this will require clarity of vision, financial imperative and sustained coordination among sectoral, or even societal, actors.

Linking strategies with national priorities and European initiatives

Widespread access to generative AI tools has centred discussion on its implications – both positive and negative – for learning and teaching. Some practitioners see AI's potential for better personalisation to fit individual learning styles, for example, transforming the classroom experience for neurodiverse learners;[5] others fear it will compromise academic integrity and undermine assessment. Generally, the immediate challenges and opportunities AI presents for learners and educators have created lively debate among learning and teaching communities. By contrast, the importance of developing strategies to support AI adoption in research and operations has, up until now, been relatively overshadowed by the immediate need to meet the challenges in learning and teaching.

As initiatives gain traction, this will likely be rectified. It is reasonable to assume that strategies will mature to encompass AI adoption in all performance dimensions of the institution, taking account of the different missions of education, research, knowledge transfer and institutional management. A more concerning trend is, perhaps, the absence of any real connection between institutional responses to AI and national-level initiatives. The integration of AI in public services is a big question for society as a whole, and the approaches to tackle it are many and varied: Scotland, for example, has established a national advisory group on AI. As another example, Spain's Ministry for Digital Transformation and Public Service approved its national AI strategy in May 2024. Similarly, Finland's AuroraAI programme and France's AI for Humanity strategy offer instructive models for higher education engagement in national AI ecosystems.

While government responses have shown broad acknowledgement of universities' role in cultivating responsible attitudes to AI, the sector has been largely absent in conversations about its adoption in wider society and compliance with European-level regulation. As these conversations advance, universities will have to look beyond their most immediate concerns and develop more outward-looking strategies and communication plans.

5 MacKenzie, M. & Elliott Spaeth, "Using AI to serve inclusive education", EUA Expert Voices 9 July 2025

In terms of pursuing and maximising technological advancement, it will be critical for universities to highlight what the sector can contribute – both on the purely scientific dimension as well as the human-level. And it will be even more critical that their institutional strategies reflect a willingness to connect with national and European-level initiatives in the digital space.

In terms of pursuing and maximising technological advancement, it will be critical for universities to highlight what the sector can contribute – both on the purely scientific dimension as well as the human-level.

Starting small – strategies that enable incremental adoption

Robust strategies are essential for setting and advancing institutional priorities. But strategising for the implementation of AI across university functions presents a new challenge. AI tools hold the promise of automating workflows, thereby freeing up capacity for more creative work among academic and professional staff. While few would disagree that the potential to accelerate laborious tasks using AI is huge, it would be premature to develop strategies that assume an AI-driven increase in productivity or try to quantify the additional workforce capacity AI will (potentially) generate.

A strategy on its own does not guarantee the successful implementation of new technologies. As with any change process, laying the groundwork is critical. And in the case of AI, the field is so dynamic that a phased approach to implementation is appropriate. While fundamental transformation may be possible (and in some cases, a legitimate goal), a more achievable 'micro-ambitious' approach with clear steps will de-risk embedding AI into an already complex ecosystem of institutional activities.

For instance, piloting initiatives without over-committing resources can yield multiple benefits: it allows communities to identify use cases that provide (whole) institutional value; it can help identify the 'right' tasks to automate and importantly, it can (quickly) reveal the hidden costs of AI.

The inclination to start small was evident among university representatives who participated in EUA-organised workshops. There was a clear appetite for investment in AI literacy and technical upskilling at the individual level as a means of equipping institutions to pursue larger-scale projects. Further, participants underlined that the introduction of AI provided opportunity to (re)focus attention on the importance of developing core academic skills, such as critical thinking (cf. Chapter 3: Training). This rediscovery of foundational competences fostering 'human intelligence' mirrors wider calls for curricula that foster critical thinking in the digital era.

A slower, more measured and mindful approach has another important advantage. It can give institutions space and time to understand the data they have and identify the data they need in order to mainstream AI. Broadly speaking, universities have been slow to mature their data management processes. But the proliferation of new technologies, and the realisation that their successful adoption depends on robust data governance, has propelled the conversation forward.

Increasingly, university leaders are turning their attention to the fuel that feeds AI; data has come to occupy a new place in strategic decision making. And with this comes an understanding that institutional strategies need to make provisions for building a healthy data culture. Formulating clear policies for assessing vendor-provided tools and educating communities on the regulatory dimensions can help prevent a scattering of non-compliant initiatives with little potential to be scaled up. Additionally, investing in data governance and management processes in parallel with the rollout of initiatives can build the foundation for a steadier, wholesale adoption of AI across the institution.

> While fundamental transformation may be possible (and in some cases, a legitimate goal), a more achievable 'micro-ambitious' approach with clear steps will de-risk embedding AI into an already complex ecosystem of institutional activities.

Balancing strategies as momentum quickens

Without the culture and capacity to support implementation, institutional strategies can fall short of expectations. This can cause disillusionment and in the case of AI, deter exploration among institutional communities. Sustaining confidence therefore depends on balancing ambition with capacity and ethical assurance as well as highlighting the benefits of AI use for individuals and our institutions.

As experimentation deepens, institutions – and society – have a vested interest in keeping this momentum going. Through a trial-and-error process, we can expect to better understand where AI adds value and which initiatives are worthy of long-term investment. Without doubt, universities should seek to accommodate this discovery process. By developing strategies that encourage and reward (rather than limit) possibilities for exploration, universities can build digital literacy, drive technological advancement and crucially, shed light on the value of interdisciplinarity in optimising AI.

> By developing strategies that encourage and reward (rather than limit) possibilities for exploration, universities can build digital literacy, drive technological advancement and crucially, shed light on the value of interdisciplinarity in optimising AI.

To do all this, universities will have to get the balance right. Strategies will have to serve the institutional community in the first instance but also align with Europe's broader goals for open data and digital sovereignty. Similarly, institutions will have to weigh up the likely benefits versus the cost of investing in pilot initiatives. But perhaps most importantly, they will have to generate a risk-appetite without compromising ethical standards or the community's wellbeing.

3. Training

While university missions might remain largely unchanged in the era of AI, many institutions acknowledge that AI will become integral to all scientific, economic and social endeavours. Therefore, AI literacy is a key competence that can enrich learners' academic and future professional lives.

With the right training and resources, AI presents clear opportunities for growth – both for universities and for individuals that make up their communities. But for institutions to successfully adopt or mainstream AI technologies, they should consider three key questions:

- What benefits are anticipated from the tool?
- How might it impact its users?
- Who requires training, and in which areas?

Ultimately, AI literacy is about more than technical readiness and developing practical skills; training needs to guide user individuals and communities towards moderate and mindful use, based on the idea that not everything that is possible makes sense or is sensible.

AI in the curriculum – a common base and discipline-specific knowledge

In certain disciplines – astrophysics or computational chemistry, for example – practices in using AI are established in the academic community. Similar trends are now emerging in the humanities and social sciences, where AI is transforming areas such as linguistics, history, and philosophy through new modes of analysis and interpretation.

In general terms, the uptake of AI in research is accelerating and with this, knowledge of relevant technologies becomes an integral part of research-based education. As is the case for all digital tools, AI technologies should serve – rather than steer – the specific goals of the relevant university department, research team or pedagogical initiative. This discipline-centred approach has the benefit of maximising existing competence and encouraging boundary-pushing exploration into how AI can support research practices.

> At the institutional level, working towards universal literacy should, in the first instance, cultivate an awareness of AI's capacity and limitations and, where appropriate, encourage deeper engagement among users.

However, recent developments have made clear the need to build baseline AI literacy across the university community in parallel. By this logic, students, administrators, leadership and academics would acquire sufficient knowledge to work beyond passive use and meaningfully contribute to the debate around the integration of new technologies in their domains. This idea aligns closely with European policy frameworks such as the Digital Competence Framework for Citizens (DigiComp 2.2), which promotes basic digital AI awareness for all.

At a societal level, the demand for AI skills across many sectors is putting pressure on universities to ensure employees enter the labour market with at least a basic understanding of the technologies. At the institutional level, working towards universal literacy should, in the first instance, cultivate an awareness of AI's capacity and limitations and, where appropriate, encourage deeper engagement among users. By tailoring training programmes to serve the diverse needs of their communities, universities would, in theory, build their own digital capacity while equipping learners with sound judgement and a critical attitude towards AI.

The challenge for universities – and for society more generally – is to ensure that training on the technical side – regardless of the level – consistently challenges the user to consider the ethical implications of using AI.

Educating beyond the technical skills

Aside from ensuring basic competence and cultivating judgement among users, universities should be granted a large degree of autonomy in adopting AI technologies. Institutional – or even departmental – priorities will influence attitudes and necessarily determine how new technologies are approached.

Moreover, the adoption of AI should be rooted in a digital culture that goes beyond the development of technical capabilities. As powerful as technologies are, they do not exist in isolation. Digital, legal and AI literacy are all connected. It follows, therefore, that frameworks and mechanisms designed to encourage engagement should address the legal, strategic, and regulatory aspects of AI. For instance, training modules can combine General Data Protection Regulation (GDPR) awareness with practical data-governance exercises, ensuring that users understand compliance as part of everyday practice. At the same time, the general principles of Open Science (including data accessibility and sharing) should not be jeopardised by exercising a non-reflective and pseudo-accurate mode of data protection.

It is increasingly understood that building a digital culture supportive of responsible AI integration starts with educating users on the ethical dimension. Proficiency in using AI technologies should be underpinned by an understanding of their limitations and the risks they carry, whether these relate to inadequacy of data, inherent bias or the environmental impact of AI.

Technologies will advance and the regulatory environment will evolve, requiring a reasonable flexibility in the design of training programmes. It is the ethical dimension that is the one constant which can ensure longevity and adaptability of training programmes across disciplines and contexts. (cf. Chapter 1: Ethics, principles and values)

Finding and scaling use cases – universities as safe spaces for discovery

AI technologies are evolving quickly, but current models have clear limitations, and their full potential—particularly in real-world use cases—remains uncertain. For this reason, there is a strong argument that concrete use cases, especially those that carry an ethical imperative, and are technically and economically feasible, should be driving the conversation. There is much speculation about the possible uses of AI with much larger capabilities than what we see today, including projections about reaching a general artificial intelligence. However, it would be risky to let speculation guide implementation in the immediate term. Instead, once a solid and convincing use case is established, relevant training should follow to support the initiative's scaling and mainstreaming.

However, the challenge for universities is determining which specific uses add value to their missions. Of the various, mostly small-scale pilots – which include, for example, the creation of AI avatars of teachers – few have taken off to the point of demanding universal attention or upskilling. Therefore, universities have not been investing in extensive training for fear this will return little reward for the institutional community.

Instead, the focus has been on formulating guidelines that encourage controlled and responsible experimentation with AI in the university setting. These guidelines have largely been welcomed among academic and student communities. They signal the importance attributed to the responsible use of AI throughout the institution and have, for many, eased uncertainty around what constitutes (in)appropriate use. However, in many cases guidelines are not aligned with the academic or broader institutional strategy, nor do they shed light on where experimentation could or should be deepened.

Moreover, broad-brush principles are often criticised for providing too little instruction on how to usefully and safely integrate AI into operations with 'transactional' tasks and repetitive workflows. Generally, universities recognise this and understand that, as experimentation deepens, guidelines will have to be updated to reflect the proliferation of technologies and probable uses within the institutional context.

As these guidelines evolve to promote and support use-case experiments, it will be critical that outcomes are properly evaluated based on impact (e.g. are students learning more or better?) rather than metrics (e.g. how many students use the learning assistant?).

Demystifying AI – whose job is this?

Early responses to the rollout of generative AI in recent years have served as an opportunity for universities to reaffirm their values. The sector was quick to acknowledge the impact but emphasised that universities' approach to AI should be underpinned by respect for fundamental values.

As institutions that tend towards values-based change, universities' stance on what constitutes responsible development and use of technologies tends to be respected. This is especially true among those concerned about the social and ethical implications. A values-based approach is an opportunity to embed principles and encourage reasonable use among all citizens. So how should universities approach training to set the stage for AI's ethical integration in wider society?

Universities are often seen as key players in demystifying AI. Demystification is a crucial first step toward building understanding and digital literacy, giving users more than just familiarity with interfaces, but insight into how these technologies work, what they can realistically achieve, and where their limitations lie. This understanding allows users to make informed decisions about when and how AI adds value, fostering discerning and responsible use.

Whether demystifying AI should start with training in how the technologies are built, what they can do or where the risks lie is up for debate. What is certain is that universities need to sit at the table, alongside government and industry representatives, to establish where the gaps are and design training that equips citizens to engage with AI safely and effectively.

A concerted effort to upskill – who should design and deliver training?

Many applications of AI, including within the university context, are at the limit of what might be considered responsible or compliant. ChatGPT, as an example, can be useful for enhancing assessment questions. But for evaluating students, greater trust and transparency is needed than is the case with off-the-shelf AI. Given this current uncertainty around what constitutes high-risk use (cf. Chapter 4: Regulation and compliance), universities tend to advocate for a multi-dimensional approach to training and upskilling that deals explicitly with the risk factors.

Technical, legal, organisational and societal factors can influence how successfully digital tools are taken up, both within the university sector and across society more broadly. As such, these various elements should be addressed in AI-related training programmes. But with the agendas of various stakeholders – government, industry and civil society – at play, the question of how training programmes are designed and delivered is an important one.

One approach would be to design training programmes which draw on the combined expertise of these various stakeholders. Universities, for example, might be better able to educate users on the environmental or regulatory factors in partnership with civic organisations and public authorities. Similarly, training in specific skills might be more usefully delivered in cooperation with industry partners.

However training is approached, there is broad consensus on two points. The first is that users (as opposed to developers) should be prioritised, i.e. training should focus on where AI can be useful and discourage use where there is no clear benefit. The second is that training should promote critical engagement with AI, enabling learners to use the technology in ways that enhance their skills and agency, rather than reducing their capacity to think independently.

Managing the gap between theory and practice

Successful change often depends on the readiness of communities to adapt their ways of working. Communities are more likely to embrace new technologies when users perceive them as beneficial and manageable, feel confident in their ability to use them, and are ready to adjust their workflows to take advantage of new possibilities. AI undoubtedly offers lots of new possibilities, but real, positive and sustainable change will depend on expectation management, effective communication, and adequate support for user communities.

“Communities are more likely to embrace new technologies when users perceive them as beneficial and manageable, feel confident in their ability to use them, and are ready to adjust their workflows to take advantage of new possibilities.

While some foresee AI having a strong influence on university practices, the real impact on teaching, learning, and research is still unclear. Early experiences suggest that perceptions vary: some worry AI may reduce problem-solving skills, while others see potential to enhance critical thinking when used thoughtfully. Thoughtful use, however, depends on competencies such as metacognition and reflective judgment – skills that often need to be developed through training and guided practice.

Training will be central to reconciling different experiences and understanding the gap between theory and practice. If universities can sustain consistent, open dialogue on AI's possibilities and limitations, their communities will be better able to discern where technologies add value – both to users' lives and institutional missions.

“If universities can sustain consistent, open dialogue on AI's possibilities and limitations, their communities will be better able to discern where technologies add value – both to users' lives and institutional missions.

4. Regulation and compliance

The most significant upcoming challenge to universities in terms of AI regulation is the implementation of the European Union's AI Act (2024). The EU's approach to AI has been orientated towards trust. The premise is that only when users trust AI will they begin to engage with it and reap the benefits. This approach forms the basis of the AI Act, which regulates the technology according to the level of risk that particular uses carry. In practical terms, this means that if the use of AI has considerable influence on human lives, there must be guarantees that the systems used are trustworthy. As education influences lives at many levels, using AI here is to be more tightly regulated than, for instance, using AI to optimise cooling systems in data centres.

Risk levels, controls and compliance

In the context of education and risk categories, the use of AI for access, evaluation and monitoring of learners is considered high risk in the AI Act. Therefore, related practices must adhere to high standards in terms of training and human oversight. Where the risk is considered unacceptable – in emotion tracking, for example – the use of AI is forbidden.

While risk levels are clearly delineated in the AI Act, and uses categorised according to the risk they carry, there are still grey zones for universities in terms of which practices are AI Act-compliant, and which are not. For example, according to the Act "AI systems intended to be used for the purpose of assessing the appropriate level of education that an individual will receive or will be able to access"[6] carry high risk. This risk level could, in theory, apply to adaptive learning using AI, meaning that related activities would have to comply with requirements for the data used for training and include provision for human oversight.

> While risk levels are clearly delineated in the AI Act, and uses categorised according to the risk they carry, there are still grey zones for universities in terms of which practices are AI Act-compliant, and which are not.

Similarly, AI systems used for evaluating learning outcomes could be subject to requirements for high-risk use. The precise guidelines as to where the stricter rules apply are expected in 2026. In the meantime, universities are becoming accustomed to a certain level of ambiguity and working towards the cause-no-detriment principle in their experimentation with AI.

6 AI Act, p. 127

Overseeing compliance at the national level

The precise guidelines will come into force as part of an AI governance system that is, to some extent, under construction. The AI Act foresees a system wherein national bodies are mandated to implement and monitor compliance with the new regulation. As it stands, many of these national bodies are working to get a better understanding of the actual implications for various sectors. For some of these authorities, their regulatory power has traditionally been limited to product safety and market surveillance.

Overseeing compliance in higher education represents a new competence for many of these national bodies; they will have to learn fast to roll out implementation nationally and at the same time ensure that steps are in line with European standards and expectations. It is important that this be a coordinated effort to avoid repeating some of the mistakes in the adoption of General Data Protection Regulation (GDPR) where, for instance, the same proctoring tools were considered acceptable by some agencies while not by others.

An opportunity for dialogue

The newness of the approach necessitates early-stage dialogue between universities and authorities. By opening dialogue channels early, universities can put forward their case for AI systems in education that are safe and trustworthy, but not so restrictive as to block the advantages that AI can offer learners and university staff alike. Cultivating a positive relationship with national surveillance authorities would serve the higher education sector well. This is especially true given that compliance with the AI Act is just one of many regulatory challenges universities are currently facing. The recent wave of European digital regulation, for example on cybersecurity, or the need to work with national authorities in the field of research security, are examples of similar areas.

Member state governments are tasked with designating authorities to deal with the AI Act at the national level. In many cases, these will be established agencies whose scope of activities in monitoring and implementation will be expanded. These agencies will also act as a bridge to the European AI office, tasked with coordinating implementation among member states.

“By opening dialogue channels early, universities can put forward their case for AI systems in education that are safe and trustworthy, but not so restrictive as to block the advantages that AI can offer learners and university staff alike.

Pooling expertise for the benefit of all

Likely, it will be advantageous for universities to engage with these agencies via the bodies that represent them at national level (i.e. university association or rectors' conference). This would not only serve to amplify universities' common concerns or priorities, but it could also afford individual institutions the benefits of pooling their academic, legal and technical expertise.

Such pooling of resources could be useful for monitoring compliance requirements, understanding their implications, and generally, laying the groundwork for optimising and harmonising universities' use of AI within the regional or national context.

Challenges at the institutional level

Fruitful dialogue within a regional or national higher education system will, undoubtedly, ease the transition. Nevertheless, challenges will present and have to be met within the individual institution. Looking to the future, it is conceivable that education technology providers will develop AI-based systems for the purpose of assessment, for example. Here, it is important that the units within the university responsible for buying this product are well equipped to judge whether it is AI Act-compliant. What is more, the relevant unit or user community will have to be adequately resourced to meet the compliance obligations required for deployers of high-risk systems, including human oversight, keeping logs and having the overall technical capacity to use the system.

Besides managing the more formal institutional-level partnerships with technology providers, the informal use of AI as 'shadow IT' will still carry risks for universities. For example, many teachers will have already built up their personal routines for using AI to assist with examinations, using readily available models like ChatGPT or Copilot – models that by no means meet the European standards for high-risk uses. Despite the very real risk individuals' practices pose for universities, it will be difficult to convince users within the institutional community to give up their established routines and switch to a compliant product the university has available. This transition from free experimentation among students and staff to adherence to institutional rules will have to be carefully managed. It will require active participation among all parts of the university's community and adequate training of users for the best chances of success. This stresses the great importance that AI governance is going to have in ensuring that universities get the best possible benefits of AI while keeping risks under control.

5. Sustainability and societal impact

Using AI in universities has consequences outside of those directly linked to the core missions of learning, research and innovation. It impacts sustainability, in its broadest sense, making it impossible for universities to ignore the knock-on effects of integrating new technologies. This raises questions that are not easy to answer:

- How can universities reconcile AI's climate impact with their ambitions to reach carbon neutrality?
- How can AI technologies preserve and enhance diversity and wellbeing within the institutional community?
- And importantly, how can universities' actions encourage a deeper understanding of technologies and a healthy attitude towards AI in society?

Environmental impacts

Looking at the global impact of AI, there is an increasing awareness that small, individual actions have repercussions that stretch far beyond the limits of the user's personal computer. Here, evidently, the immediate concern is the cumulative impact of one-off uses on resource consumption and ultimately, climate change.

At a global level, environmentalists tend to focus heavily on emissions generated using AI. While the impact is undoubtedly real, it is hard to measure; estimates of the power used today by big data centres are imprecise, and projections for the future will depend both on how widely the technologies are taken up and how efficient they become in the years ahead.[7] Attempts to calculate the long-term impact of AI are further complicated by the fact that energy consumption varies greatly depending on the specific use: for example, the researcher training a foundation model for a discipline-related purpose will generate a large, immediate impact in terms of use of resources.

By contrast, the learner using generative AI sporadically will use fewer resources, but the net impact of all learners using generative AI will, likely, be much larger. Importantly, projections for resource consumption need to take account of the fact that dedicated systems for specific tasks require much less energy than general purpose systems.

Despite the many conveniences AI offers, users are often removed – both geographically and psychologically – from the engines and resources that power it. Those casually using generative AI to summarise articles are, to some extent, alienated from the mechanics behind it – that is, the local impacts of the construction of large scale data centres in terms of water use and environment destruction, the mining of minerals to make batteries and chips for the actual machines, and the production of e-waste in their continual operation and eventual dismantling. This not only means a linear increase in power demand but also a shift in magnitudes such that society's current ambitions for sustainable energy production are challenged.

7 Chan, Sophia (2025), How much energy will AI really consume? The good, the bad and the unknown Nature 639 5 March 2025, p. 22-24

The shift is dramatic enough that providers (and governments) are considering the commission of new nuclear power plants despite the still unsolved issue of long-term treatment of nuclear waste.

Few will think about the rare earths being mined in Inner Mongolia or Tanzania when writing a prompt to ChatGPT. Nevertheless, these rare earths are integral to the whole, complex system that drives AI.

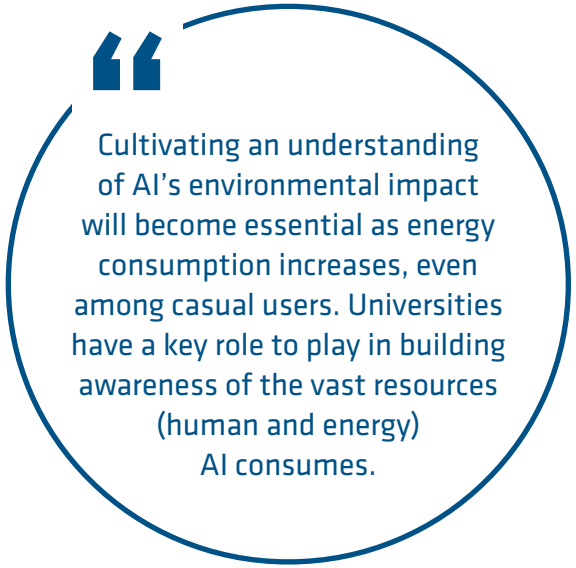

And while ignorance or indifference might be convenient, cultivating an understanding of AI's environmental impact will become essential as energy consumption increases, even among casual users. Universities have a key role to play in building awareness of the vast resources (human and energy) AI consumes.

Social impacts

Confronting users with some of the uncomfortable realities of AI use brings the wider social justice dimension into sharper focus. The use of human trainers in large generative AI models is one such example of AI's damaging social impact. These trainers tend to be workers in the Global South, who are systematically exposed to highly disturbing content when moderating the models, often to the point where it is damaging for their mental health.[8] Companies often rely on paying low wages for data-labelling and content-moderation work outsourced to countries like Kenya, the Philippines, and Venezuela. This embeds a structural inequality in which value accrues to AI firms in the Global North while the health risks and low wages are borne by labour in the Global South.

The quicker uptake of AI in emerging countries exposes another vulnerability of the Global South, where users have embraced these technologies with a strong belief that they bring immediate, personal benefits.[9] While some benefits may indeed be quick and tangible, emerging countries have tended to adopt AI technologies with comparatively little concern for protecting privacy and retaining control of their data. By contrast, countries in the Global North have, in many cases, the luxury of taking a long-term view, where a slower, steadier uptake of AI ensures privacy and data sovereignty are, at least, considered (if not guaranteed).

On the other hand, AI technologies also bring benefits to disadvantaged countries, not just in the form of skilled employment but also in increased access to free informational and educational resources. AI is already being deployed in government initiatives to improve agricultural productivity, improve healthcare outcomes, and support climate resilience among other initiatives. Universities have a duty to encourage and enable such initiatives while mitigating social harms.

8 Niamh Rowe "It's destroyed me completely': Kenyan moderators decry toll of training of AI models", The Guardian 2 August 2023
9 Google | IPSOS (2025) *Our life with AI: from innovation to application*, p.12

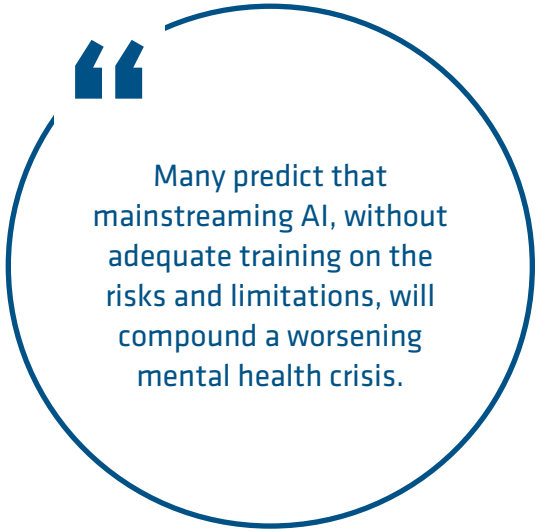

Impact on the institution – safeguarding wellbeing among users

The upsurge in new technologies, combined with the promise from private providers that these will improve efficiency, puts pressure on universities to move forward and move fast. This is often compounded by a fear of being left behind in the AI race.

This pressure and fear can prompt university leaders to take hasty decisions around implementation, with insufficient thought given to how AI actually improves people's everyday experiences of work or learning.

Identifying where AI adds real value, and where it is potentially damaging requires competencies in psychology, human-machine interaction and, potentially, pedagogy and cognition.

Many within the academic community fear that an overreliance on AI will ultimately diminish students' autonomy and sense of agency. This risk is, perhaps, particularly pronounced for learners who assume a certain rationality is embedded in AI-generated output. These are students who, for example, come to depend on AI for inspiration, tutoring or brainstorming, overlooking the importance of human-to-human interaction in their academic (and social) development.

Concerns that (over)dependence on AI will lead to apathy or social isolation feeds into the wider debate around digital culture, digital sobriety and wellbeing. Among university professionals, there is a widely held concern that learners' mental health has been in decline, particularly since the Covid-19 pandemic. Many predict that mainstreaming AI, without adequate training on the risks and limitations, will compound a worsening mental health crisis.

To mitigate this, it is critical that universities implement structures that allow for systematic dialogue and exchange among users. Establishing where there are vulnerabilities among user communities will go a long way towards safeguarding human relationships and prioritising wellbeing in the institutional setting.

Widening participation in higher education – can AI be the great 'leveller'?

Fears that a diminishing sense of agency among learners will lead to disengagement in academic pursuits may be well founded. However, there is also evidence that systems based on AI can be beneficial in terms of engaging, for example, neurodiverse learners, who can struggle in more traditional classroom environments.[10] By offering personalised pathways that can be adapted to individual styles, needs or ambitions, AI can open up possibilities for learners with diverse profiles. Technologies can, for example, be used to create online content – through avatars and chatbots – that is made available to lifelong learners, whose priority is often flexible timing when accessing material.

10 Mackenzie, M., Elliott Spaeth (2025), "Using AI to serve inclusive education", EUA Expert Voices 9 July 2025

While some argue that AI will amplify differences or inequalities in student populations, others see the potential for it to make the classroom experience more inclusive and personalised. Since there is growing evidence that AI may exacerbate the digital divide, it becomes all the more important to design inclusive AI solutions–ones that intentionally account for differences such as neurodivergent learning styles, linguistic backgrounds, and cultural contexts. When developed responsibly, such systems can support more personalised, equitable, and meaningful learning experiences for all students. However, close monitoring of these solutions is essential to detect and address any hidden biases or discriminatory effects.

Outsourcing of competence – what constitutes reasonable and responsible use?

As generative AI rapidly developed, governments and industry scrambled to be seen as a 'first mover'. Universities are trying to balance this pressure with a healthy degree of caution, often tending towards cultures which promote the reasonable and responsible use of new technologies while also recognising the need to experiment.

However, in terms of approaches to AI, the question of what constitutes reasonable or responsible use becomes more complex. It is wrapped up in the wider debate around which functions we want to preserve as essentially 'human' and which ones can be automated without risk.

While aiming towards efficiency can indeed support universities' more human-orientated missions (an efficient administration, for example, serves learners in a timely manner with high quality support), it is important to steer the conversation about AI away from a pure efficiency discourse.

At the institutional level, there might be a convenience in using AI to produce essays more quickly or as an integral part of planning and executing scientific experiments.

But if repeated use compromises fundamental academic skills or research principles, universities need to rethink their approach. Ultimately, AI should not be used only to make work easier; instead, its potential to enhance learning and deepen knowledge should be central.

At a societal level, AI raises questions about our individual and collective capacity for critical thinking, and whether this will be diminished, eventually, if we externalise what we now consider 'human' competence to machines that mimic human behaviour. So, while technological advancement is often celebrated for enhancing productivity, we need to acknowledge, as users, that not everything (curiosity, for example) can – or should – be measured or judged in terms of efficiency, and that efficiency is rarely a means for achieving a pedagogical goal.

What is more, when much of our work is automated using non-transparent AI models, questions have to be asked as to which processes require human intervention and which can reasonably be performed using AI, without compromising our reasoning, judgement or integrity.

Conclusion

> By building mechanisms for participatory dialogue and decision making into technology integration, universities equip themselves to determine where AI can support everyday practices, where it has transformative potential, and where it may carry disproportionate risk.

As an integral part of the social fabric, universities' approaches to AI will shape attitudes towards new technologies in the workplace, in civil society, and in cultural and domestic spheres.

Technologies have advanced quickly, with the potential to transform the way universities work for the better. But positive, meaningful change takes time. Rather than blindly pursuing efficiency and productivity, universities can – and should - take a more deliberate and reflective approach to AI adoption. Rooted in human values and guided by their core missions, universities are well positioned to advance and embrace new technologies responsibly.

This begins with building awareness of possibilities and risks, cultivating shared ownership of the process and creating space for reflection on the real and lasting value of digital tools. By building mechanisms for participatory dialogue and decision making into technology integration, universities equip themselves to determine where AI can support everyday practices, where it has transformative potential, and where it may carry disproportionate risk.

Ultimately, adopting AI in universities should serve their immediate communities as well as the public interest. Proficiency in AI tools can, in many cases, support the work of individual users or whole university departments, but digital literacy goes far beyond technical capability. Training on the ethical dimension should foster a sense of personal and collective responsibility, prompting users to consider the social and environmental footprint of creating and interacting with the technologies.

While the most immediate risks and rewards of using AI are becoming apparent, the longer-term capabilities and cumulative impacts of these technologies remain unknown. Within higher education and research, the sensible use of AI can widen access to knowledge, support the academic endeavour and deepen our understanding of the world. But to fully exploit the potential of these technologies, universities – and the people who work, study and conduct research within them – must take a balanced, human-centred approach that upholds ethical principles, protects academic values and safeguards the wellbeing of academic and student communities.

EUROPEAN UNIVERSITY ASSOCIATION

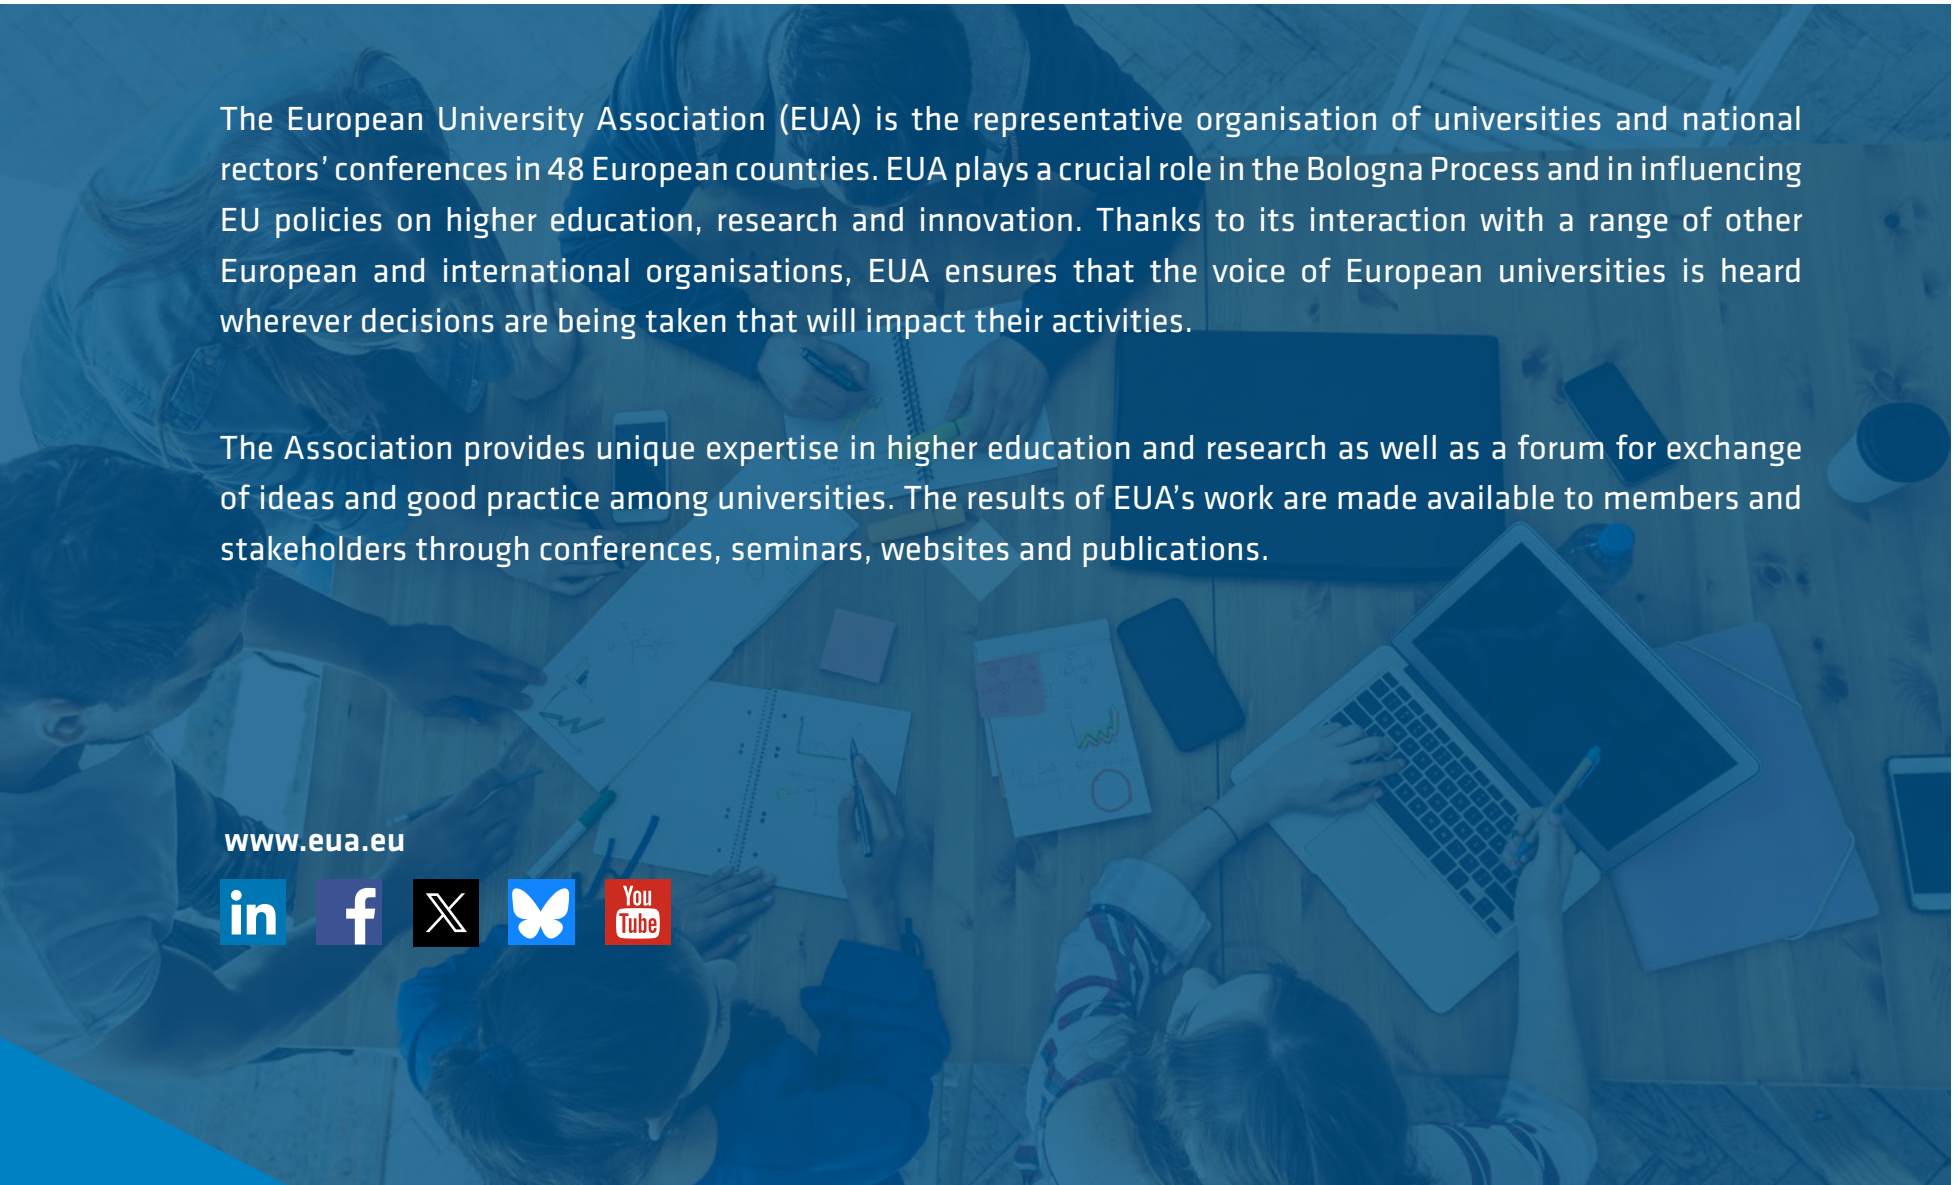

The European University Association (EUA) is the representative organisation of universities and national rectors' conferences in 48 European countries. EUA plays a crucial role in the Bologna Process and in influencing EU policies on higher education, research and innovation. Thanks to its interaction with a range of other European and international organisations, EUA ensures that the voice of European universities is heard wherever decisions are being taken that will impact their activities.

The Association provides unique expertise in higher education and research as well as a forum for exchange of ideas and good practice among universities. The results of EUA's work are made available to members and stakeholders through conferences, seminars, websites and publications.

www.eua.eu