

BUILDING THE FUTURE OF LEARNING IN THE AGE OF AI. THREE QUESTIONS TO HEATHER HASELEY

Ester COIS

University of Cagliari, IT

Heather HASELEY

Imperial College London, UK

Email: ester.cois@unica.it

Interview

Artificial intelligence is rapidly transforming higher education and lifelong learning. Over the last few years, AI has moved from being an emerging technological innovation to becoming an increasingly pervasive presence in educational environments, influencing how learning experiences are designed, delivered, assessed, and supported. Universities across the world are exploring its potential to personalise learning, expand access, improve efficiency, and respond more effectively to rapidly changing societal and labour-market needs. At the same time, the growing integration of AI into educational systems raises important questions.

Heather Haseley is the inaugural Chief Lifelong Learning Officer at Imperial College London, where she leads Imperial Lifelong Learning as part of the institution's Science for Humanity strategy. Throughout her career, she has held leadership positions at several internationally recognised institutions, including the University of Exeter, Arizona State University, the University of Pennsylvania, Northwestern University, and Columbia University. Her work focuses on learning innovation, educational transformation, digital credentials, skills development, immersive technologies, and the future of learning, with particular attention to how emerging technologies can support more equitable, human-centred, and future-oriented educational systems.

While technological innovation offers unprecedented opportunities, it also challenges long-established assumptions about teaching, learning, expertise, and human development. Concepts such as personalisation, scale, efficiency, and accessibility are frequently invoked in discussions about AI-enhanced education, yet they are often interpreted in simplistic or purely technological terms. Personalisation, for example, can easily be reduced to automated adaptation; scale can become synonymous with mass content delivery; and efficiency can risk overshadowing deeper educational purposes. These tensions are particularly relevant within the field of university lifelong learning. Lifelong learners represent highly diverse populations, characterised by different educational backgrounds, professional trajectories, aspirations, and learning needs. In this context, the challenge is not simply how to use AI, but how to use it in ways that strengthen rather than diminish the human dimensions of learning. Questions of equity, critical thinking, learner agency, social connection, and pedagogical quality become increasingly important as AI systems become more powerful and more widely available. These questions have become central to Heather Haseley's work on the future of learning and the evolving relationship between artificial intelligence, higher education, and lifelong learning.

Rather than focusing exclusively on technological capabilities, her research and leadership have consistently explored the broader educational and societal implications of AI, encouraging institutions to reflect on the kind of learning ecosystems they seek to create. A recurring theme throughout her work is the idea that AI should help us become more human, not less, by strengthening rather than diminishing the relational, reflective, and developmental dimensions of learning. In this interview, she reflects on three themes that have become central to contemporary debates on AI and lifelong learning: personalisation beyond automation, the responsible use of AI in educational settings, and the challenge of scaling learning opportunities without losing meaning, quality, and human connection.

Ester Cois: Personalisation is often presented as one of the main promises of AI in education. Different people need different things at different times and in different ways and therefore require highly relevant and responsive learning solutions. At the same time, we should be careful not to equate personalisation with automated adaptation or content customisation alone, recognising instead the importance of learners' contexts, motivations, and learning trajectories. How can AI support forms of personalisation that go beyond automation and truly respond to learners' diverse needs without oversimplifying learning or replacing human pedagogical judgement?

Heather Haseley: One of the first things we need to recognise is the pace at which this field is evolving. The answers I would have given six months ago would already be different today, and by the time this interview is published, what is possible will have changed again. Universities therefore need to think carefully about how they develop learning experiences and maintain their relevance in an environment characterised by rapid technological change.

A useful distinction is between personalisation and hyper-personalisation. Hyper-personalisation resembles the logic of social media feeds, where individuals are exposed only to information that aligns with their interests and preferences. While this may appear beneficial, it can create informational bubbles, encourage atomisation, and ultimately become dehumanising. Educational experiences should not become so personalised that they are detached from context, community, and meaningful interaction with others.

Instead, AI should be used to help us become more human, not less. It should create space for connection, conversation, collaboration, and practical application. This begins with good pedagogy. The emergence of AI does not invalidate decades of educational research. Scaffolding remains important. Universal Design for Learning remains important. Learners still need opportunities to encounter concepts in meaningful sequences, to be stretched, and to struggle productively, because learning often occurs precisely in those moments of challenge and growth.

Within this pedagogical framework, AI can support learners in several valuable ways. One important application is helping individuals analyse their existing skills and capabilities and identify personalised pathways towards specific goals. Rather than decontextualising learning, AI can help learners understand which courses, experiences, or learning opportunities are most relevant for achieving their personal and professional aspirations.

AI can also facilitate just-in-time learning. Learners often need support at the exact moment a problem emerges, whether that is during a class, at work, or in the middle of the night. AI can provide immediate guidance and access to information when it is needed most. At the same time, learners must develop the capacity to critically evaluate the information they receive rather than accepting it uncritically.

Another important area concerns personalised support and intervention. Educational institutions already collect large amounts of information about learner engagement and behaviour. AI can help identify patterns that indicate when learners are struggling and can do so earlier than traditional approaches. This allows educators to provide timely and targeted human support, making interventions more effective and helping learners achieve success more quickly.

AI can also contribute significantly to accessibility and inclusion. Whether through translation, adaptation of materials, or support for learners with disabilities and diverse learning needs, these technologies can help make educational opportunities more equitable and accessible.

At the same time, institutions must intentionally design opportunities for learners to engage with one another. Personalisation does not eliminate the need for community. Learners need spaces where they can discuss ideas, challenge assumptions, collaborate in teams, and benefit from perspectives that differ from their own. Human peers often provide forms of critical feedback that AI systems cannot currently replicate. AI tends to be highly self-affirming; it rarely tells us that an idea may be flawed or that we may have overlooked an important perspective. Other people do. They bring alternative experiences, different viewpoints, and constructive disagreement, all of which remain essential for meaningful learning.

For the same reason, learners need opportunities to engage with authentic problems, real-world debates, and complex challenges that require dialogue, judgement, and collective reflection. These are fundamentally human dimensions of education that cannot simply be delegated to AI.

Finally, AI can serve as a connector between learners. It can help match individuals with peers who share similar goals, with coaches and mentors who can provide targeted support, and with people who have already succeeded in the areas to which they aspire. In this way, AI can help build relationships, strengthen learning communities, and connect learners with sources of inspiration and guidance.

Ultimately, the most meaningful forms of personalisation are not those that isolate learners, but those that help them navigate pathways towards their goals while strengthening the human relationships that support learning and development.

Ester Cois: How can we make responsible use of AI in educational settings and support learners in engaging with these tools in a balanced way—avoiding both fear and over-reliance—while fostering awareness of the importance of maintaining such balance?

Heather Haseley: I think this is a particularly important question because discussions about AI often become polarised. On one side, there are those who believe AI will solve nearly every problem and fundamentally improve society. On the other, there are those who view it as a threat that will undermine education, work, and human relationships. The reality is much more complex. AI offers significant opportunities, but it also raises important challenges that require careful governance, critical reflection, and deliberate choices about how these technologies are used.

The research landscape around AI and learning is evolving rapidly, and recent studies highlight a critical tension at the heart of AI-assisted learning. Shaw and Nave (2026) introduce the concept of “cognitive surrender,” referring to the tendency for individuals, particularly under time pressure, to defer decision-making entirely to AI rather than engaging their own reasoning. Shen and Tamkin (2026) demonstrate that AI use affects skill formation

differently depending on the learner's stage of development: experienced users may risk deskilling through cognitive offloading, while novices who rely excessively on AI may never develop foundational competencies in the first place. Gerlich (2025) similarly warns that habitual cognitive offloading can negatively affect critical-thinking capacity over time.

At the same time, Wang and Zhang (2026) provide an important nuance, showing that learners who engage AI as a genuine pedagogical partner rather than as a shortcut can simultaneously exercise critical evaluation and strategic offloading, thereby freeing cognitive resources for deeper and more transformative learning. The crucial distinction is therefore not whether AI is used, but how and why it is used.

Together, these studies reinforce a foundational principle of learning science: growth occurs within what Vygotsky (1978) termed the Zone of Proximal Development: the space between what learners can accomplish independently and what they can achieve with appropriate guidance and support. As educators, we have an ethical responsibility to remain informed about emerging evidence and to design learning experiences that preserve this space for development. AI should function as a scaffold rather than a substitute. We must intentionally create opportunities for reflection, critical thinking, and productive struggle because the stakes involve nothing less than the long-term cognitive development of learners.

This means that learners should not simply be given unrestricted access to AI tools without guidance. Just as learning itself requires scaffolding, AI use requires scaffolding. Students need opportunities to develop foundational knowledge and competencies before relying extensively on

AI support. If AI consistently bridges the gap between what learners can currently do and what they are expected to achieve, they may never fully develop the underlying skills required to evaluate, challenge, or improve the outputs they receive. Critical thinking cannot be reduced to telling students to "check the answer;" learners must possess sufficient knowledge and understanding to judge whether an answer is actually good or bad.

For this reason, educational design should continue to create spaces where learners grapple with complexity, encounter uncertainty, and engage directly with challenging tasks before turning to AI for support. We should use these technologies to extend learning, not to bypass it. Equally important is helping learners understand how learning itself happens, fostering reflective practice and metacognitive awareness alongside technological fluency.

More broadly, responsible use requires open and informed conversations about AI. Universities, educators, learners, and policymakers need sufficient levels of AI literacy to engage in meaningful debate about the values that should guide these technologies and the roles we want them to play in society. If we do not actively discuss these issues, we risk allowing technological development to proceed without sufficient consideration of its educational and societal implications.

Another important issue concerns ownership and the development of AI models. At present, much innovation is driven primarily by commercial interests and profit incentives. History reminds us that technological progress does not automatically translate into social benefit. Universities therefore have an important role to play in shaping alternative futures. Rather than simply adapting to technologies created elsewhere, higher education institutions should collaborate to help create and influence the kinds of tools they want to see in the world.

This includes supporting the development of open models and smaller, more sustainable systems that require fewer resources and can be used in a wider range of contexts. Such tools could be particularly valuable in regions with limited connectivity or infrastructure,

enabling local deployment, offline use, and broader accessibility. Equally important is ensuring transparency about how these systems operate and what their costs are.

Learners and educators should understand not only the benefits of AI but also its human, environmental, economic, and societal implications. Many AI systems are becoming embedded in everyday technologies so seamlessly that their presence often goes unnoticed. Yet every interaction carries costs in terms of energy consumption, infrastructure, labour, and social impact. Making these implicit dimensions visible is essential if we want learners to become informed and responsible users.

Finally, we are beginning to observe signs of societal concern regarding the economic consequences of AI, particularly in relation to employment. Public attitudes towards AI appear to be shifting as people begin to experience some of its wider social and economic effects. Whether these concerns represent a temporary adjustment or a longer-term challenge remains to be seen. What seems clear, however, is that new forms of work and new skill requirements will continue to emerge. Universities therefore have an ethical responsibility not only to educate learners about AI but also to support large-scale reskilling and upskilling efforts that enable individuals and communities to participate successfully in these evolving futures.

Ultimately, responsible use is not simply a matter of regulating technology. It is about making conscious choices regarding how AI is integrated into learning, how learners develop alongside it, and how educational institutions contribute to shaping technological futures that enable both human flourishing and social good.

Ester Cois: Scale is another recurring theme in current debates on education and lifelong learning. AI is frequently described as a way to expand access and provision at scale, yet scaling learning is not the same as merely scaling content delivery. In your view, how can AI help scale learning opportunities while preserving pedagogical quality, relational dimensions, and meaningful engagement with learners?

Heather Haseley: The first question we should ask is why we want to scale in the first place. Educational research consistently shows that smaller learning environments, personalised support, and even one-to-one tutoring can often be more effective. Therefore, scale should never be treated as an objective in itself. We need to be clear about why we are scaling any learning activity before deciding whether scale is the right approach.

If the goal is to increase access to learning opportunities that would otherwise be unavailable, then scaling can be highly valuable. Access does not automatically equate to learning, but making knowledge and learning opportunities available in contexts where they would otherwise be inaccessible is a worthwhile objective. At the same time, many institutions pursue scale because it reduces the cost of delivery. If more learners can be reached with fewer resources, provision becomes more efficient. The important question, however, is at what cost. If efficiency comes at the expense of deep learning, meaningful engagement, or educational quality, then we need to reconsider what we are doing.

For this reason, I believe we should always design for impact first and then determine whether scale is the most appropriate mechanism for achieving that impact. Rather than designing for scale and only later asking what impact it might produce, we should begin by identifying the outcomes we hope to achieve. Is the goal greater reach? Improved competency development? Faster creation and dissemination of knowledge? Once we are clear about the desired impact, we can decide whether and where scale is appropriate.

AI can support scale in several important ways. For educators and instructional designers, it can accelerate the creation and updating of learning materials, enabling more timely and contextually relevant content. The pace of discovery and change is significantly faster than it once was, and AI can help us shorten the development cycle so that learners receive relevant knowledge more quickly. Human expertise and pedagogical design remain essential, but AI can help increase the pace at which high-quality learning experiences are created.

Feedback represents another domain in which AI can add considerable value. AI can facilitate faster feedback, more frequent assessment opportunities, and earlier identification of learning difficulties. The feedback may not always be as deep as that provided by a human expert, but it can be delivered more quickly and more often. This creates more opportunities for iteration and improvement, while also enabling educators to identify where learners need additional support. In this sense, scale and personalisation become closely connected: AI can help us target human interventions more effectively and make those interventions more meaningful.

Beyond individual feedback, AI can help identify broader patterns of misunderstanding and recurring learner difficulties. These insights can inform redesign efforts and guide educators towards the areas where human support will have the greatest impact. In this way, technology can help ensure that human contact is used where it matters most.

Particularly interesting opportunities emerge in the design of simulations, game-based learning, and iterative practice environments. AI can support the creation of learning experiences in which learners repeatedly attempt tasks, receive feedback, revise their approach, and try again. It can help generate multiple variations of challenges within defined parameters, making it easier to create opportunities for practice, experimentation, and gradual skill development. These forms of test–feedback–revision cycles are often difficult to provide at scale through traditional approaches.

From the learner's perspective, one of the most significant benefits concerns opportunities for practice, particularly in areas such as language learning. One of the longstanding challenges in language education is that learners rarely receive enough opportunities to practise speaking. AI can help provide additional conversational practice, immediate feedback, and low-stakes opportunities to build confidence before engaging in real-world interactions. It does not replace human communication, but it can help learners arrive better prepared for it.

It can further support mastery at an individual pace, recognising that learners progress in different ways and at different speeds. Learners vary in the amount of practice they need, the speed at which they progress, and the kinds of support they require. By providing additional opportunities for practice, feedback, and iteration, AI can help learners identify areas where they are struggling and make the pathway towards their goals more visible and more achievable.

Perhaps one of the most transformative possibilities lies in contextualisation. One of the challenges of large-scale and global learning is that educational examples and case studies are often developed for a generic audience. Yet the realities faced by learners differ significantly across countries, professions, and communities. The challenges encountered by a learner in the United Kingdom may not be the same as those faced by a learner in Italy or elsewhere. AI can help generate contextually relevant scenarios, case studies, and problems that reflect learners' own environments and make learning more meaningful and immediately applicable. Ten years ago, creating this level of contextual adaptation for individual learners would have required considerable time and effort. Today, AI makes it increasingly feasible.

At the same time, we must remain vigilant. Scaling learning should not mean scaling content alone. The most important challenge is ensuring that we scale opportunities for practice, reflection, interaction, experimentation, and genuine learning rather than merely increasing the volume of information delivered. If we lose context, relationships, and human engagement, we risk undermining the very purpose of education. The goal should not be to scale content delivery, but to scale meaningful learning.

Ester Cois: Thank you very much.

REFERENCES

- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), 6. <https://doi.org/10.3390/soc15010006>
- Shaw, S. D., & Nave, G. (2026). Thinking-Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender. *The Wharton School Research Paper*. https://doi.org/10.31234/osf.io/yk25n_v1. Available at SSRN: <https://ssrn.com/abstract=6097646>
- Shen, J. H., & Tamkin, A. (2026). How AI Impacts Skill Formation. arXiv preprint arXiv:2601.20245.
- Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press.
- Wang, S., & Zhang, H. (2026). Pedagogical partnerships with generative AI in higher education: How dual cognitive pathways paradoxically enable transformative learning. *International Journal of Educational Technology in Higher Education*, 23(11). <https://doi.org/10.1186/s41239-026-00585-x>

