


Capacity-Building for Innovative UK AI & Education Policy

Megan Ennion, Professor Ann Kristin Glenster,
Inbar Bobrovsky, and Gareth Brinkworth



This report originates from the *Capacity-Building for Innovative UK AI & Education Policy* roundtable held at King's College, University of Cambridge on 18 November 2025, and convened by the AI and Education Community, University of Cambridge, and the Glenlead Centre in November 2025.

With gratitude for the support of the Accelerate Programme for Scientific Discovery, University of Cambridge.

This edition is published by LOGOS, with the substantive findings and recommendations of the original report preserved.

DOI: <https://doi.org/10.17863/CAM.128443>

Contents

1	Executive Summary	1
2	Introduction	3
3	Policymaking	5
4	AI Education Policies for Human Flourishing	8
5	Policies for The Purposes of Education	10
6	Competencies and Skills for The Future	12
7	Closing Remarks	13

1 Executive Summary

In November 2025 the AI and Education Community, University of Cambridge, and the Glenlead Centre convened a policy roundtable with thirty key stakeholders to discuss the need for policy to respond to, safeguard against, and benefit from the transformative effect of AI on education.

From the roundtable, the following recommendations for national AI policies emerged:

1. **Reassess assessment paradigms and support schools in navigating privacy, data protection, digital inclusion, safety, and data-access challenges.**
2. **Consider introducing national teacher training standards and ringfence time for professional AI development for educators.**
3. **Confirm a commitment to human-centric AI education policies, including contestability rights for AI-driven assessments.**
4. **Promote and enable responsible procurement, including mandatory evaluation of the quality and pedagogical effectiveness of edtech tools.**
5. **Identify and enable independent research programmes to evidence and evaluate AI's impact on wellbeing, cognition, and workload.**
6. **Strengthen government-backed, independent research infrastructure capable of supporting large-scale, longitudinal studies, alongside evidence-based knowledge-brokering, enabling schools and policymakers to make informed decisions about AI.**
7. **Develop and adopt AI definitions to be used across all AI education policies.**
8. **Embed information and AI literacy as part of the national curriculum.**
9. **Explicitly commit to digital inclusion in all AI education policies.**

The discussion centred both on the substance of AI education policies and how policies are made. First, to the substantive point, **human flourishing, digital, inclusion and human-centric AI** emerged as foundational values that should underpin the use of AI in education and all AI education policies. This would include **protecting human agency** and **empowering human educators**. It is critical to identify **the purpose(s) of education** in the age of AI. This necessitates examining questions regarding the **very nature of knowledge**. Participants emphasise that knowledge is contextual, cultural, and historically grounded, and expressed **concern that technology companies act as de facto gatekeepers** and informal arbiters of information and episteme. There is therefore a need to ensure that technological developments align with educational purpose.

The roundtable further discussed what **competencies and skills** will be needed in the future.

Participants identified the capability of **future-thinking** and the strengthening of **information and AI literacy** as key components of future education. Participants noted that AI literacy is not just about how AI tools work, but also about **understanding their impact in society**.

Secondly, the discussion circled in on how AI and education policies are made.

Participants argued that **AI education policies should be adopted early**, following a three-step sequence of identifying **(1) context, (2) purpose, and (3) need**. There is also a need for a robust lexicon of **AI definitions** to be used in policy.

Repeatedly, participants identified a pressing need for independent research and forward-leaning evidence on the role of AI in education, **noting that the current evidence base is dominated by small, isolated studies**. Participants emphasised the need for **government-backed, independent research infrastructure capable of examining AI's impacts at scale**, including strengthening funding for academic research and critical systems for knowledge-brokering and knowledge-transfers between stakeholders.

2 Introduction

In November 2025, the AI and Education Community at the University of Cambridge and the Glenlead Centre convened a roundtable discussion on **Capacity-Building for Innovative UK AI Education Policy**.

AI is reshaping educational practice, influencing integrity, assessment, wellbeing, equity, and the wider knowledge ecosystem. Department for Education (DfE) consultations and forthcoming updates to generative-AI safety guidance underscore the urgency of developing policy that is robust, informed, and oriented toward the public good.

The purpose of the roundtable was to convene and facilitate rigorous discussions between researchers, policy communities, and stakeholders in the field, and to explore how the UK might move from reactive policy cycles to proactive, adaptable, human-centred, futures-oriented frameworks.

The roundtable was convened by Megan Ennion and Professor Ann Kristin Glenster, who chaired the discussion. This report presents the key points that emerged.

The roundtable explored three questions:

- Do we need AI specific education policies and if so, what should they cover?
- Where should policy draw the line between regulation and self-regulation?
- Are there AI specific issues that require different types of evidence and information for policymakers?

Background

As with all sectors, AI is radically transforming education. Recent international headlines have underscored the urgency: a government-sponsored ethics report abroad was found to contain [fabricated references](#); pupils in the UK have voiced fears that AI is [eroding their ability](#) to think and study independently; and, in the United States, [Big Tech companies are funding large-scale teacher training](#) in AI — prompting questions about **corporate influence and educational sovereignty**.

This conversation comes at a critical juncture for the UK. The Department for Education (DfE) has conducted several consultations on AI and education. DfE is also commencing work to update its **Safety Expectations** for generative AI in education, reflecting growing recognition of new and emerging risks. DfE has also signalled a forthcoming [White Paper to support the use of AI tutors](#), signalling a major shift in how learning may be delivered in England.

Across the education sector, AI is already **transforming** teaching practice and reshaping assessment, particularly in relation to **academic integrity**. Broader concerns centre on **digital equity and accessibility**, as well as the **unregulated** ways students are using AI for **tutorial** or **mental-health support**. Questions of bias, hallucination, and discrimination persist, alongside mounting unease that **overreliance** on AI may weaken **critical thinking** and deepen patterns of **technological dependency**. The field is grappling with how to balance **personalisation and pedagogy**. While adaptive AI systems

promise tailored learning experiences, they risk fragmenting shared knowledge, narrowing curiosity, and undermining **collective educational aims**.

At the same time, AI offers unprecedented opportunities for **fairer, more inclusive, and cost-effective education**. Its thoughtful adoption could help prepare the UK for the **challenges of tomorrow**, equipping students with the skills needed for **innovation and growth**. Yet for policymakers, the pace of technological change raises pressing questions: how to ensure policy is **robust, informed, and inclusive**; how to mitigate harms while realising **genuine educational value**. As the UK seeks to harness AI for social and economic good, this is a pivotal moment to decide what kind of education system — and what kind of **knowledge ecosystem** — we want to build.

Behind these questions are broader themes of the direction of travel whereby education policy ultimately will determine the distribution of the **power of knowledge for future generations**, and what we consider to be education for the **public good and citizenship**, which are pillars of a **democratic society**. Thus, education policy cannot just be a guessing game for the skills and jobs of the future but must be informed to lay the very foundation of our **social contract**. Striking a balance between **innovation and tradition** is vital — preserving the intellectual rigour and human values that institutions like Cambridge have long championed, while engaging critically with the possibilities of technological transformation.

Recommendations

In summary, participants identified a need for national policies to:

1. **Reassess assessment paradigms and support schools in navigating privacy, data protection, and safety, and data-access challenges.**
2. **Consider introducing national teacher training standards and ringfence time for professional AI development for educators.**
3. **Confirm a commitment to human-centric AI education policies, including contestability rights for AI-driven assessments.**
4. **Promote and enable responsible procurement, including mandatory evaluation of the quality and pedagogical effectiveness of edtech tools.**
5. **Identify and enable independent research programmes to evidence and evaluate AI's impact on wellbeing, cognition, and workload.**
6. **Strengthen government-backed, independent research infrastructure capable of supporting large-scale, longitudinal studies, alongside evidence-based knowledge-brokering, enabling schools and policymakers to make informed decisions about AI.**
7. **Develop and adopt AI definitions to be used across all AI education policies.**
8. **Embed information and AI literacy as part of the national curriculum.**
9. **Explicitly commit to digital inclusion in all AI education policies.**

3 Policymaking

A Need for National AI and Education Policies

The roundtable reinforced the view that there is an urgent and **critical need for the UK to adopt robust AI education policies at the national level**. Participants felt that in the vacuum of national policies, institutions, educators, and **students were left to fend for themselves**.

Participants thought that government has a critical role in ensuring **the structural conditions necessary for safe and digitally inclusive use of AI for learning and knowledge creation and dissemination**. To that end, participants believed that government has a responsibility to develop and adopt robust AI education policies.

Participants were particularly concerned regarding **a need for system-level policies** which could include robust procurement frameworks, to prevent dependence on proprietary vendors and multiyear lock-in of schools. Participants also identified a role for national policies in **addressing digital divides**, including promoting digital inclusion and enabling investment in equitable access. Participants also felt that national policies should also address **end-user data security**, which some saw as the most urgent risk requiring policy action.

Specifically, participants found that these **responsibilities cannot be left to individual schools**, whose time, expertise, and resources are already strained. Instead, participants argued for a **high-level national framework** to support busy schools, provide coherent guidance, reduce institutional and learner risk, and ensure that shared educational aims are protected across the system. In short, effective and ethical AI adoption requires **coordinated national action**, not isolated local responses.

A Need for Early Policy

Participants expressed concern regarding the pattern of introducing technology in education because the technology exists, and not because of need.

Some participants also believed that futures-thinking should guide policymaking. According to this view, national policies should first articulate desired futures, then map backwards to determine what educational ecosystems would be needed.

They strongly believed that **technology should only be introduced, having first identified, in order: context, purpose, and need**. Several participants highlighted that the adoption of technology in education frequently fails because this sequence is reversed: technology is introduced first, and only afterwards is it given a purpose. As one participant observed, the result is that technology becomes “a tool in search of a purpose.”

By contrast, technologies in crisis settings and SEND schools often succeed precisely because context, purpose, and need are clearly defined in advance; they are need-led rather than tech-driven.

A Need for AI Definitions

Participants recognised **a critical need for national AI education policies but** had divergent views on what these policies should address. The lack of precision of the term AI was repeatedly brought up as education policy often treats AI systems as homogeneous when, in fact, their effects, risks, and governance vary significantly. Participants expressed concern that without clearer definitions, policy would be both too broad and too shallow. Some participants also thought that AI is sufficiently distinct to warrant dedicated regulation and policies.

Participants called for **more exact and robust definitions reflecting different types of AI systems, uses, and capabilities, and contexts.**

Addressing Power Asymmetries

Participants were concerned that an absence of national policies enabled further entrenchment of **power asymmetries**, particularly seen through the prism of commercial edtech companies versus public institutions. Participants felt that there was a danger that private market actors were shaping decisions about what counts as knowledge, which tools are adopted, and how systems are designed.

In the words of some participants, policymaking should not be steered by big tech interests, and the UK should avoid becoming “a cog in someone else’s machine.” As such, some participants felt that national policies were needed to **safeguard educational sovereignty** and **prioritise the rights and wellbeing of young people**. It was important to participants that policies should **ensure that technological developments align with educational purpose** rather than commercial imperatives.

The Critical Role of Independent Research and Evidence

Repeatedly, participants noted the **paucity of evidence** to inform and support AI education policymaking. Especially, they noted that:

- Many existing studies are small and isolated.
- Much US research is funded by technology companies.
- There is no large-scale UK research infrastructure for AI in education.

Participants illustrated the **need for more robust and forward-leaning evidence** by pointing to social media as a cautionary tale of how harms were identified too late. They also emphasised that technological disruption is not new. History offers lessons on how societies have adapted to prior technological shifts and how unintended consequences can emerge over time. Drawing on historical analysis was seen as vital for avoiding repeated mistakes.

Furthermore, while collaboration with edtech companies was deemed necessary, participants nevertheless cautioned against allowing commercial actors to define what constitutes

legitimate research. Participants therefore argued for **substantial government funding for rigorous independent research and independent evidence-gathering at scale** to fill emerging and widening critical evidence gaps.

There is also a critical need for systemic knowledge-brokering whereby research can inform policymaking. More research and experimentation for how this can be done is urgently needed.

4 AI Education Policies for Human Flourishing

Human Flourishing as the Normative Anchor and Measurable Benchmark

Human flourishing emerged in the discussion as an important value in education¹. Participants believed that flourishing should not be treated as an abstract ideal; rather, it was noted that global human-flourishing frameworks offer cross-cultural, empirically grounded, and actionable guidance on what education should support.

Human flourishing was presented as a practical foundation for defining the aims of AI-enabled education. Participants suggested that flourishing indicators could inform benchmarks for AI–human collaboration, helping policy remain anchored in enduring human values rather than short-term technological imperatives. Such frameworks were seen as a way to ensure that technological development aligns with human development, rather than allowing it to displace or narrow it.

Several participants highlighted **the importance of digital inclusion** and ensuring that human flourishing was for everyone. It was seen as important that AI education policies included provisions to ensure accessibility for all users across a range of needs.

Human-Centric

Participants repeatedly returned to the question of **what it means to be human in an AI-mediated world and educational system**. They strongly advocated for a human-centric approach to AI in education but cautioned that without a robust definition, human-centricity may become an empty placeholder term.

There were concerns regarding new ethical challenges as AI gains more advanced capabilities to detect patterns, infer mental states, or analyse behaviour. For example, participants were worried that AI systems would be used to categorise or constrain learners, or the effects of a lack of robust mechanisms for contesting AI-generated evaluations.

Participants felt that genuine human-centricity AI education policy required explicit commitments to:

- Human autonomy and independent thought
- Mental health and expressive freedom
- Human capabilities education should cultivate
- Values that ought to guide system design

¹ For the purposes of this report, human flourishing is understood in line with the Global Flourishing Study's definition: "the relative attainment of a state in which all aspects of a person's life are good, including the contexts in which that person lives." This definition foregrounds the multidimensional nature of flourishing and recognises that wellbeing cannot be separated from the social, cultural, and institutional environments that shape human lives. VanderWeele, T.J., Johnson, B.R., Bialowolski, P.T. et al. The Global Flourishing Study: Study Profile and Initial Results on Flourishing. *Nat. Mental Health* 3, 636–653 (2025). <https://doi.org/10.1038/s44220-025-00423-5>

Protecting Human Agency

Participants were concerned about diminishing student agency, particularly in younger learners. This was seen as a pressing issue as many AI systems are optimised for productivity rather than education, which risk creating environments in which students offload independent thinking to automated suggestions. It was therefore considered **important that AI education policy protect individual human agency**.

Participants felt that AI education policy should be drafted to help teachers and students to understand the power dynamics and worldviews embedded in the AI systems they use. It was thought **important that students and teachers recognise that these tools are not neutral but drive particular narratives**. Policy should ensure that students and teachers retain the **capacity to think and act independently** when interacting with increasingly persuasive and adaptive systems.

Participants recognised that the adoption and integration of AI in education also exposes students and educators to the risk of harm, especially from **data and security breaches**. They thus highlighted the need for policy guidance on data protection at the individual educator level, not just at the level of major developers. These safeguards were seen as crucial to protect students and educators' agency.

Empowering Human Educators

Participants saw **a core purpose of AI education policy was to protect and empower human educators**.

It was widely acknowledged that teachers are on the frontline of AI adoption, often without necessary time, expertise, or guidance. Currently, a disproportionate burden is placed on educators to manage students' use of AI, despite many teachers lacking foundational knowledge required to navigate AI confidently. These issues are further exacerbated by significant workload pressure.

Participants strongly **challenged the prevailing narrative that AI “saves teachers time.”** They noted that while AI may reduce workload in some areas, teachers are losing time in others, particularly from increased communication demands, the need to verify AI-generated work, and the growing responsibility to assess and oversee student use of AI. AI affects teacher's workload indirectly (for example, through AI-generated parental communications). Thus, participant found the efficiency narrative to be misleading and potentially harmful, obscuring the escalating pressures educators face.

Participants argued that **it is for policy, not individual practitioners, to provide clear guidance, safeguards, and support**.

5 Policies for The Purposes of Education

Interrogating The Purpose of Education

Participants noted that AI compels an opportunity for a fundamental re-examination of the purpose of education. As AI reshapes how knowledge is accessed, generated, and verified – and as society confronts broader transformations including democratic decline and post-truth discourse – existing educational models cannot simply be adjusted at the margins.

Participants emphasised that meaningful and responsible AI education policy cannot proceed without clarity about education's core purposes. This led participants to identify key questions:

- What should education to be in an age of AI?
- How do we know if it is being achieved?
- What should be measured, and why?
- What capacities should learners develop?
- What role should teachers play?

Participants noted that today's education systems were designed for different purposes and past conditions. The concern was not that the UK must build a system that remains unchanged for the next 50 years, but that it must build one that is resilient and agile, capable of evolving responsibly as societal needs and technologies continue to develop. Participants cautioned that redesigning education today without embedding such adaptability risks locking the system into another cycle of hurried reinvention.

Reactive policy cycles were considered insufficient in the context of accelerating technological change. Participants instead argued that **national AI education policy should be agile, forward-looking, and capable of responding to technologies, and societal conditions which do not yet exist.**

Several participants noted that although educational institutions often resist structural change, societal and technological conditions evolve regardless. As a result, participants felt that AI risks being inserted into existing systems in ways that reinforce outdated assumptions rather than advancing a coherent educational vision.

Rather than asking how AI can be integrated into current structures, the more **urgent question is what education itself must become in an AI-mediated world.** Effective AI education policy must therefore begin not with inherited and legacy systems, but with **the futures we expect and want to create.**

Defining the purpose of education was seen as a prerequisite for building human-centred, future-resilient AI education policy and for ensuring that technological adoption aligns with broader societal and developmental aims.

What is Knowledge?

Related to the purpose of education, participants brought up **questions on the very nature of knowledge.** Participants asked what knowledge educators is intended to instil

in students, and what human knowledge means in the age of AI.

Participants emphasised that knowledge must be understood as contextual, cultural, and historically grounded, not simply modelled as information.

The hierarchy of knowledge was described as increasingly compromised, with **technology companies acting as *de facto* “keepers of knowledge”** by determining what enters models and what is excluded. Tech companies as informal arbiters of what “counts” as knowledge raises profound questions about epistemic authority, cultural positionality, and the depth of meaning that distinguishes human understanding from LLM-generated text.

Beyond this, participants noted that **technology companies increasingly influence** which pedagogies are embedded in tools, who gains access to specific forms of knowledge, what knowledge is taught or prioritised, and how knowledge is evaluated or measured. The consequence may be that **AI systems risk reshaping educational epistemologies** in ways that may not align with societal or pedagogical values.

This led the discussion onto the **indispensable role of philosophy** in shaping AI education policy. Participants stressed the need to interrogate foundational questions: **What counts as knowledge? Who decides?**

Participants believed that philosophical expertise was crucial to interrogate foundational questions about AI, education, society, and the very nature of knowledge. These questions were seen as foundational to educational policy, not peripheral.

Participants stressed the need to engage with diverse philosophical traditions, including eastern perspectives, to better understand concepts such as agency, knowledge, relationality, and the human–technology interface. Philosophy was described as **essential not only for interpreting technological capabilities but also for expanding our imaginative and conceptual possibilities**. This was seen as critical to developing a culturally sensitive, conceptually robust AI education policy.

6 Competencies and Skills for The Future

The discussion not only questioned the purpose of education and knowledge, but also addressed what capacities or skills education should teach.

Futures-Thinking

Participants argued that **teaching students to think about the future** explicitly should be a central goal of education as a core competency. This would include cultivating:

- Foresight and anticipation
- The ability to analyse multiple possible futures
- Critical thinking and sense-making in complex environments
- Skills for imagining, designing, and building future societies

Participants considered these core competencies for democratic citizenship and long-term societal resilience. Futures-thinking was described as **needed to navigate uncertainty, shape society, and respond to emerging global challenges**. As one participant put it, “We need to anticipate and build our future, not just react to it.”

Strengthening Information and AI Literacy

Participants reported growing concern about LLMs “poisoning” information environments and the inability of many students to recognise when information is distorted. It was recognised that these are not novel challenges. Participants called for **reinvigorating information literacy** – which has existed since the 1970s – as a core part of the curriculum, extending to AI.

The focus should continue to be on teaching individuals:

- Critical engagement with information
- Checking sources
- Understanding how information is produced
- Evaluating accuracy and reliability

In addition to information literacy, a basic, shared understanding of what AI is and how it works was viewed as essential to conducting mature policy conversations.

Participants also emphasised that **AI literacy must include understanding the impact of integrating AI into society, not just how AI tools operate**. They noted that many existing AI literacy frameworks focus on the mechanics of AI, while neglecting the broader social, ethical, epistemic, and behavioural implications of widespread AI adoption, an omission that must be addressed in any meaningful educational response.

7 Closing Remarks

This roundtable identified a pressing need for national AI education policies to ensure that the adoption of these new technologies is **human-centric**, promoting **human agency**, and empowering **human educators**. AI is a transformative technology that may lead us to question the **very nature of knowledge** and the **purpose of education**.

Participants urged policymakers to lean into this debate by adopting **clear AI definitions, early policies, identifying context, purposes, and needs, and** ensuring that the policies promoted **future-thinking** and **digital inclusion**. Participants also acknowledge the challenges with robust policymaking in a rapidly changing environment and called for a strengthening of **independent research** and **evidence** to be part of the **knowledge-brokering** between stakeholders which is essential for impactful, forward-leaning, and responsible AI education policies for the UK in the future.



www.logoscambridge.org