


Evaluating AI in Education for Core Human Capabilities

Policy Roundtable Report

Professor Ann Kristin Glenster and Megan Ennion



The Oxford-Cambridge Forum for Flourishing, Learning, AI and Research (OxBridge FFLARE) convenes research, policy and practice communities to examine the future of education in an AI-mediated world. It is a collaborative forum supported by LOGOS, University of Cambridge, and the Generation AI initiative at Reuben College, University of Oxford.

This report presents findings from the roundtable held at King's College, University of Cambridge on 7 May 2026.

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

The authors wish to record their thanks to Peter Bannister, LOGOS, University of Cambridge, and to Dr Andrew Serazin and Dr Anthony Bridgen, Reuben College, University of Oxford.

<https://doi.org/10.17863/CAM.130858>

© LOGOS 2026



Contents

1 Executive Summary	1
2 Introduction	2
3 Purpose of Education in a Time of Uncertainty	4
4 Core Human Capabilities	5
5 Need for Evidence	8
6 Adoption of AI in Education	10
7 Conclusion	12

1 Executive Summary

In May 2026, the inaugural OxBridge FFLARE policy roundtable convened leaders from academia, policy and practice to examine how artificial intelligence (AI) is reshaping education. Participants identified a central challenge: *despite rapid adoption of AI, there is no unified vision of how it should shape the purpose of education nor the development of core human capabilities.*

Examining questions around the purpose of education in fostering core human capabilities, the need for evaluation and evidence, and issues of AI implementation, the roundtable shifted the conversation from instrumental or mechanical concerns to a high-level perspective concerning the future and purpose of education itself.

Participants recognised that education is about more than simply preparing learners for future employment. It should also build core human capabilities, including agency, wellbeing, flourishing, social relationships, and trust. More than providing a mere conduit for skills, education is vital for the formation of habits and identity; indeed, for the capacity not merely for attainment, but for becoming.

Recognising a broader and deeper purpose of education necessitates a shift in assessment paradigms, empowerment of teachers, and new conceptual frameworks for pedagogy. Assessment is also about human validation and inclusion, and as such should not be modelled only on transactional and quantifiable measurements but also include room for relational deep learning.

Achieving the adoption of AI in education that promotes human learning and core capabilities requires independent, well-funded research to generate high-quality evidence at scale. Deciding how AI in education is assessed cannot be left to private technology companies. Evaluation and implementation of AI must also centre the views of those most affected by it, particularly young people. More guidance is required for teachers and parents, and schools alone should not be saddled with procuring new AI technologies responsibly. Centralised government policies and regulatory coherence are needed to ensure that AI is adopted in education in ways that promote human learning, inclusion, and equity.

The roundtable highlighted that many of the most consequential questions raised by AI in education are not purely technical, but normative and relational: what we value, what we seek to cultivate, and what kinds of learning environments we want to protect.

2 Introduction

In May 2026, the inaugural policy roundtable of the Oxford-Cambridge Forum for Flourishing, Learning, AI & Research (OxBridge FFLARE) took place at King's College, University of Cambridge. The event brought together stakeholders from academia, government, industry, teaching, and policy. The goal was to identify what core human capabilities education should cultivate in an AI-influenced world, and how AI's impact on them should be assessed.

AI is having a **significant impact** on education. Digital and AI technologies are rapidly advancing and spending in the EdTech sector is predicted to reach \$620 billion by 2030¹. The UK has announced a major expansion of the EdTech Testbed programme (involving more than 1,000 schools and colleges),² alongside plans to develop and trial tutoring tools to support 470,000 disadvantaged students by 2027.³ There is also considerable interest in technology-related policy for education, across legislation, regulation, investment, market incentives, and practice. For example, the UK recently passed the Children's Wellbeing and Schools Act 2026⁴ and has undertaken a national curriculum review.⁵ Thus, AI is **already changing education in the UK and around the world**.

The extent of this activity highlights the significant interest stakeholders have in AI. What is missing, however, is an overall vision for how these technologies should shape the future. More than individual initiatives and policies, there is an emerging need for a **deeper and more strategic conversation** about AI's implications for pedagogy, learning, and the purpose of education.

The inaugural OxBridge FFLARE roundtable marked the start of an initiative to explore these themes. The discussion took place under Chatham House Rule, focussing on three fundamental questions of capability, evaluation and implementation:

¹ Morgan Stanley Research, 'Global Education's \$8 Trillion Reboot', Morgan Stanley, 7 June 2023, available at: <https://www.morganstanley.com/ideas/education-system-technology-reboot>

² Department for Education and The Rt Hon Bridget Phillipson MP, 'Education Secretary speech at Bett UK Conference', GOV.UK, 21 January 2026, available at: <https://www.gov.uk/government/speeches/education-secretary-speech-at-bett-uk-conference>.

³ Department for Education and Department for Science, Innovation and Technology, '450,000 disadvantaged pupils could benefit from AI tutoring tools', GOV.UK, 26 January 2026, available at: <https://www.gov.uk/government/news/450000-disadvantaged-pupils-could-benefit-from-ai-tutoring-tools>

⁴ Children's Wellbeing and Schools Act 2026, c. 21, available at: <https://www.legislation.gov.uk/ukpga/2026/21/contents/enacted>

⁵ Curriculum and Assessment Review Final Report, Department for Education, published 5 November 2025, available at: <https://www.gov.uk/government/publications/curriculum-and-assessment-review-final-report>

- Beyond attainment, what core human capabilities should education develop and protect in a society shaped by AI, and how can we identify gaps and emerging priorities so that these are supported equitably and at scale?
- How should the impact of AI in education be assessed against these capabilities fairly and safely across contexts – including disadvantage, SEND, and digital inclusion – and what would count as “good enough” evidence to guide policy, regulation, procurement and scale-up?
- How can teacher capacity, professional judgement and research expertise be built into the adoption and ongoing use of AI in education, such that AI strengthens education while protecting workload, empathy and professional autonomy, whilst avoiding dependency?

3 Purpose of Education in a Time of Uncertainty

The roundtable discussion highlighted that beliefs about **the purpose of education** are shaped by our **cultural contexts**, geographies, demographics, and socio-economic circumstances. It also highlighted the importance of recognising that, when considering AI in education, different perspectives and cultures are not represented in dominant Western AI models. For example, the New Brunswick Model was referenced as a new pedagogical framework to promote human flourishing and well-being.⁶

Participants noted a persistent tension in the purposes of education: schooling is expected to cultivate **shared norms** that support social cohesion and continuity in a rapidly changing world, yet these aims can sit uneasily alongside the development of learner agency, independent judgement, and critical thought.

Participants recognised that, increasingly, formal education is unable to prepare students for the workplace of tomorrow. The gap between the skills needed for the jobs of the future and the skills developed through formal education appears to be widening. As a result, without a clear promise of employment, **young people have reduced motivation to pursue education**.

Furthermore, the dominant narrative of education as a pathway to employment may have come at the expense of other narratives. Educators struggle to deliver other objectives for education, such as building the capacity for **democratic resilience**. There is also an underserved opportunity for education to help foster capacities for **psychological wellbeing** and **inclusion**.

Participants also raised broader questions regarding what kind of relationship humans would like to have with AI. As some participants asserted, the starting point for using AI in education should be to ask **what problem AI is intended to solve?**

It was also highlighted that many young people do not want to have to use AI, and so the assumption that AI should be adopted and integrated into everything should not be automatically endorsed.

Some participants questioned the assumption that society will inevitably be “AI-shaped”. Rather than treating this trajectory as fixed, they argued that the role AI plays in education and society remains a matter of human choice and governance—though making deliberate choices, rather than drifting into default adoption, will require **cultural change**.

Across the discussion, participants repeatedly returned to the question of what values underpin the education system, since these values ultimately shape what should be protected, measured, and prioritised as AI scales.

⁶ Government of New Brunswick, Department of Education and Early Childhood Development, *Well-Being: A Curriculum Support Document for the New Brunswick Curriculum Framework for Early Learning and Child Care – English*, 20 February 2026, available at: <https://www2.gnb.ca/content/dam/gnb/Departments/ed/pdf/ELCC/ECHDPE/well-being.pdf>

4 Core Human Capabilities

Core Cognitive Capabilities

The discussion about core human capabilities explored the need for future cognitive attributes. These included strengthening individuals' capabilities for **metacognition**, **meta-emotion**, and **moral reasoning**.

Agency was seen as a particularly important core human capability. It was considered critical to give children agency in encountering automated and AI systems. The need for agency was also linked to the capability to exercise **stoicism** and **form social relationships**.

Notably, some discussion concerned the use of “digital twins” or avatars in learning environments. The use of digital twins was seen as important in the context of the **formation of identity** and a **sense of self**. It was thought that a digital twin could provide a learner with agency but could also supplant the “real person”, thereby diminishing agency.

Participants emphasised that deciding which capabilities education should cultivate is, first and foremost, a **value decision**: the question is not only what young people should be able to do, but what we want education to be for, and what kind of persons and communities it should help to form. From this perspective, some suggested that education should foster habits as a form of virtue formation, and framed schooling as a process of **becoming**, not merely attaining.

Within this framing, being able to **interrogate AI**—to question, test, and critically evaluate its outputs—was seen by some as a key capability that education should actively promote. Participants also questioned whether, and in what ways, AI in education can illuminate or support distinctly human capabilities such as **social understanding** and **compassion**, rather than displacing them. A number of participants argued that education should attend to **eudaimonic wellbeing** and the conditions for a good life, including **deeper meaning**, **identity**, **belonging**, **confidence**, and **quality of life**.

Changing the Assessment Model

Participants were acutely aware that the current model of individual assessments and performance means that students will use AI tools – either supervised by their educational institutions or unsupervised – to improve their results. There was some discussion regarding the difference between students who use AI to take shortcuts to obtain “the right answer” and those who use AI to help them learn.

Participants noted that some of the outcomes education may most value—such as a sense of belonging—are unlikely to be well captured by PISA scores. Thus, while participants identified a **need to radically examine the purpose of education**, this process cannot take place while the educational system is focused on individual test results. Participants expressed a desire to move beyond testing individual pupils towards assessing schools as collective learning environments or communities, and to empower teachers with a

broader range of ways to evidence and communicate their school's contribution to learning.

Participants discussed core cognitive attributes that should be promoted as the education system and its assessment structures are reimaged and transformed. It was considered that this requires clarity about which attributes education should develop and how these can be measured. However, there was also concern that one flawed assessment regime would simply replace another unless this **new framework for pedagogy was rooted in a new conceptual understanding of the purpose of education.**

Social Relationships and Trust

Participants repeatedly returned to the idea that learning and assessment are fundamentally social and relational, whereas an increasingly modular, technologically driven, and hyper-personalised model risks treating education as a transactional exchange. While micro-assessments and personalisation may help individual learners to acquire measurable knowledge, participants argued that they do not necessarily cultivate the capabilities and trust on which learning—and wider flourishing—depends. Several participants expressed concern that AI and digital technologies can interfere with human relationship dynamics in education, potentially weakening the relational fabric through which confidence, belonging, and collaborative problem-solving are built. A central claim in this discussion was that **assessment is not primarily about the data:** it is about the **relationship between teacher and student.** In this view, assessment functions **as a form of recognition**—an educator communicating, implicitly or explicitly, “I see and value who you are becoming as a person.” Participants worried that if assessment becomes increasingly automated, the validating and meaning-making function of feedback may be diminished, even where AI feedback is comprehensive, consistent, and readily available.

This tension surfaced in discussion of AI in formative assessment. While some noted potential benefits—speed, coverage, consistency—others questioned what may be lost when feedback is detached from a human relationship. Participants suggested that students often want their effort and growth to be recognised by a human teacher, not solely by an algorithmic system.

This raised a further concern about **equity:** a future in which less advantaged schools rely more heavily on AI-mediated feedback and support, while more advantaged schools preserve the time and relational attention of human teachers. Participants stressed the importance of keeping the right goal in view: not efficiency alone, but **the role of education in providing validation, trust, and the conditions for human flourishing.**

In discussions of adoption, participants argued that **trust** is central to education but is being **eroded.** They suggested that teacher–student trust has been “rattled” by the arrival of AI as a third actor in learning and assessment—an actor introduced without clear consent, accountability, or shared norms for its use. Participants noted that people can place near-automatic trust in AI outputs, even as responsibility for errors remains unclear. They stressed the need to protect trust and accountability between humans.

Participants also linked this to a wider concern that education is becoming increasingly **transactional** rather than **relational**, driven by efficiency, measurement, and a culture of constant documentation. They warned that when AI is used invisibly for speed—by students or teachers—effort, authorship, and recognition can be eroded, and relationships may thin. In this view, the risk is not only inaccurate outputs, but a gradual weakening of the social fabric that makes learning possible: trust, rapport, and the time required for reflective professional judgement.

5 Need for Evidence

The Difficult Question of Data

A key theme to emerge was the need for **stronger evidence** and **better data** to understand how AI is reshaping education and the development of core human capabilities. Participants questioned how cognitive and other human attributes are currently measured, noting a **dearth of high-quality, longitudinal data** on the impacts of AI in schools and on the capacities education aims to cultivate.

At the same time, calls for more granular, continuous, and “robust” data were treated as inherently risky. Participants cautioned that intensifying measurement could slide into greater **surveillance** of children and their development, with potential chilling effects—particularly on agency, experimentation, and the freedom to learn without constant monitoring. This created a **central tension**: data may be needed to support a shift away from today’s assessment-driven model towards one focused on broader human capabilities, yet the same data infrastructures could easily be used to further entrench the current logic of personalised micro-assessment and behavioural optimisation.

Concerns about trust and accountability were sharpened by discussion of the political economy of evidence. Participants questioned the extent to which **research funded by technology companies** can be treated as independent, and asked what data commercial tools are collecting, how it is being inferred, and who is ultimately accountable for the standards applied. A key provocation was: **who assesses the assessor?** Some participants suggested that private actors may already recognise that many tools are not yet “ready” for high-stakes educational contexts, making it even more important that public institutions define acceptable practice before norms are set by default through deployment.

In response, participants argued for clearer boundaries and governance: the education system needs to decide what it will and will not accept in data collection and inference before it is designed and deployed by commercial technology companies. This includes clarifying which kinds of data points enable which inferences, where the “line in the sand” should be drawn, and how those lines are communicated to educators, students, and families. There were calls for evidence gathering and evaluation to be taken seriously and adequately funded through **independent mechanisms**, with **research designs that include teachers and the voices of those most affected, including young people**. Participants also emphasised the need to define **evidentiary standards** more clearly—both for research and for policy and procurement decisions.

Assessments of AI in Education

The discussion moved to how **AI itself could be assessed**. Participants felt that the starting point should be to identify a specific problem or issue, then evaluate AI according to whether it is able to address this problem or issue. As opposed to introducing an AI tool and subsequently trying to find a use case.

Participants queried whether there should be AI benchmarks or certification schemes, and what these should entail. It was suggested that benchmarks should move away from a focus on output data to assessments that reflect underlying values. There were also discussions about the potential to move beyond static benchmarks to more fluid and clustered benchmarking.

The discussion explored the need to assess AI in relation to the purpose of education today, assessing AI as a means to deliver education for the future, and the broader need to change how education itself is assessed.

Participants also noted that human judgement often functions as the default “gold standard” benchmark. They raised a philosophical question, however, about whether we actually want AI to assess in a human-like way. Should evaluations of AI be benchmarked against human judgement, or should AI be developed and assessed against criteria that aim for something better than current practice?

Participants highlighted the need for young people to be included in the development and adoption of AI in education. Participants asked for **more research giving weight to students’ voices and experiences and understanding how students are affected by AI**. As such, the conversation recognised that questions around agency are also tied to questions of oversight. Participants argued **that those affected by AI in education should be included in decision-making and governance regarding how these technologies are developed and used**.

6 Adoption of AI in Education

The Market Aspects

Participants queried the necessary market conditions to introduce AI tools in education. While overall they were positive about the use of AI, participants also issued cautionary warnings against the unregulated adoption of AI in education. For example, participants warned against using human students to experiment with AI, whether for commercial or research purposes, but also recognised that some experimentation was necessary to move forward.

Participants noted that **industry was not best placed to answer the question of what the purpose of education should be**. As these tools are being developed and introduced into schools at scale, there is a risk that they set the pathway for the future of education by building education experiences and models based on today's flawed – and soon to be outdated – assessment paradigms. The uptake of these products at scale risks freezing both the use of technology and the future of education in a paradigm where the gap between what is taught and the skills and competences needed in the future continue to widen.

Need for Friction

Participants noted that market incentives do not lead to the desired level of friction in education. Participants noted that **AI may reduce or eliminate opportunities for students to make mistakes, take cognitive detours, or engage in mental struggles that are essential for learning**. Friction is key to knowledge accumulation and thus AI in education must be designed and used with friction in mind for learning to take place.

However, the attention economy incentivises the creation of products and experiences that cause frictionless use. This is a problem in relation to “weak” and “strong” students as the first group is likely to rely on (addictive, lower quality) frictionless products, while the latter is more likely to seek out friction to ensure that they learn.

The unresolved tension between agency and friction was exemplified by advances in neurotechnologies. As such, participants acknowledged the challenges for education in empowering learners with agency in a world where everyday reality increasingly involves interacting with emerging technologies, which will be increasingly cognitively invasive.

Policy Expertise

While aiming to address how these tools can be designed to be used by, work for, and help solve problems for people, participants expressed concern about the limited social science expertise within government, among policymakers, and in technology companies.

Participants emphasised that human teachers remain critical and noted that many teachers and parents lack clear guidance on how best to respond to AI. They argued for a broader ecosystem of support that recognises teachers, students, and parents as an interconnected system, and for approaches to AI in education to be co-developed from the outset with all three.

Related concerns were raised about procurement: participants argued that schools and individual educators should not be expected to procure AI tools without specialist support. It was highlighted that procurement should not be the responsibility of schools. Instead, it should be handled at a higher level with the necessary expertise and accountability.

It was suggested that AI in schools could be governed through a **test-and-learn** approach—supporting carefully scoped, real-world implementation with clear success criteria and independent evaluation, and using rapid feedback loops to adapt, pause, or stop tool deployment based on outcomes. Noting that **national policy did not sufficiently focus on teacher agency**, participants also emphasised that teachers could—and should—be supported to take part in this kind of test-and-learn research, with the time, training, and practical infrastructure needed to contribute to evidence generation in their own settings.

There were calls for a more **centralised government policy approach**. Participants expressed concern about the digital divide, noting that better-resourced schools are already at an advantage—able to offer stronger guidance, run AI-focused teacher workshops, and in some cases draw on in-house AI expertise to support tools, training, and school-wide approaches. Participants argued that teachers should have access to support and resources across the system, and that this should not be left to individual institutions. Ensuring a level playing field nationally, they suggested, will require coordinated government policy.

Participants noted that the regulatory debate to date was not primarily about pedagogy, but about how AI is used in support functions for teachers and students. There was a view that **regulatory coherence and regulation to prevent harm** are needed. In terms of risks, participants called for more regulatory clarity. For example, some participants suggested that edtech could be labelled in a similar fashion to nutrition labels on food products.

7 Conclusion

The roundtable highlighted that many of the most consequential questions raised by AI in education are not purely technical, but normative and relational: what we value, what we seek to cultivate, and what kinds of learning environments we want to protect.

Participants emphasised that if policy is not guided by a clear account of education's purposes and the core human capabilities it should develop, default pressures—measurement, efficiency, and market momentum—may end up defining the direction of travel.

This report therefore closes not with a settled conclusion, but with a clearer map of the tensions, risks and priorities that participants felt must shape the next phase of debate and decision-making as AI becomes part of everyday learning.

OXBRIDGE
FFLARE

