AI, Learning and Judgement: Preserving Cognitive Development in an AI-Mediated Education System

Written Evidence Submission to the House of Commons Education Committee Inquiry: The use of Artificial Intelligence and EdTech in Education

Date: 31 March 2026

Submitted by: Dr David Atkinson, BEng (Hons), CDipAF, PhD, FHEA

Position: Lecturer, York St John University

Contact: d.atkinson@yorksj.ac.uk

Table of Contents

Executive Summary

This submission examines how artificial intelligence (AI), particularly generative AI, is reshaping learning processes, teaching practice, and educational responsibility. Drawing on a structured classroom intervention in higher education, in which AI use was mandated within an assessment, it argues that the primary educational risk is not misuse of AI, but premature cognitive closure—learners arriving at answers without engaging in the reasoning required to understand them.

AI enables rapid generation of outputs, but this can compress the process through which learners explore alternatives, evaluate evidence, and exercise judgement. As a result, students may produce high-quality work with limited underlying understanding. Observed student behaviour shows significant variation: some use AI as an exploratory partner, while others rely on it in ways that reduce engagement with uncertainty and decision-making.

The central challenge for education policy is therefore to ensure that AI supports cognitive formation, not just cognitive offloading. This requires a shift from evaluating outputs to evaluating reasoning, from access to technology to pedagogical access, and from tool use to responsible judgement.

The submission recommends: redesigning assessments to prioritise explanation and ownership; supporting teachers in AI-informed pedagogy; strengthening digital literacy as critical engagement; and aligning policy and funding with the preservation of decision-making under uncertainty.

The effectiveness of AI in education will depend on whether it enhances, rather than displaces, the development of independent judgement.

1. Scope of this Submission

This submission addresses the following areas within the Terms of Reference:

- Challenges and Opportunities
- The Impact on Teaching
- The Impact on Learning

It focuses on how AI (particularly generative AI and specifically Large Language Models) is reshaping learning processes, judgement formation, and educational responsibility across higher education in a business school context, with implications applicable across other stages.

2. Basis of Evidence

This submission draws on:

- Professional experience in higher education teaching, including the design and delivery of AI-integrated assessment in undergraduate business programmes, including industrial economic strategy and entrepreneurship and enterprise development;
- A structured classroom intervention in which generative AI use was mandated within an assessed task, enabling direct observation of student interaction with AI systems;
- Reflective student submissions (n=14) documenting their (pilot) use of AI in reasoning, analysis, and decision-making contexts;
- Ongoing research into decision-making under constraint, including how individuals exercise judgement in AI-mediated environments.

In the intervention, AI use was not optional but required, allowing the role of AI in student reasoning to be made visible and assessable. The learning focus was not technical proficiency, but reasoning, evidence selection, and ownership of conclusions under uncertainty.

The intervention design was informed by ongoing theoretical research into decision-making under constraint, which examines how individuals arrive at conclusions when supported by AI systems. This work anticipates that learning outcomes depend not only on access to tools, but on how learners encounter and resolve uncertainty before committing to an answer.

3. Core Argument

Artificial Intelligence in education is often framed in terms of efficiency, access to information, and personalisation. While these are important, they do not fully capture the transformation underway.

AI is increasingly shaping not only what learners produce, but how they arrive at knowledge. It filters information, proposes interpretations, and suggests answers before learners have fully explored the space of possible understanding.

The central educational risk is therefore not simply misuse of AI, but premature cognitive closure: learners arriving at answers without traversing the reasoning process required to understand them.

This has three implications:

- Learning may shift from exploration to selection;
- Students may produce outputs without fully developing underlying understanding;
- Responsibility for judgement may weaken if decisions are effectively pre-structured by AI systems.

The key policy challenge is therefore not whether AI is used in education, but how to ensure that its use supports, rather than substitutes for, the development of individual judgement.

This perspective is informed by a research framework examining how decisions become fixed under conditions of constraint. From this standpoint, learning requires not only arriving at answers, but engaging with the space of possible alternatives prior to commitment. The classroom intervention described in this submission was therefore designed to make that process visible, rather than allowing it to remain implicit or bypassed through AI-generated outputs.

4. Challenges and Opportunities

4.1 Opportunities: Augmented Exploration

AI can:

- expand access to information,
- support idea generation,
- enable rapid iteration of approaches,
- assist learners in structuring complex tasks.

In educational settings, this can:

- reduce barriers to engagement,
- support independent learning,
- enable students to explore multiple perspectives quickly.

Where used well, AI can expand the range of possibilities that learners encounter.

4.2 The Risk: Premature Cognitive Closure

The capabilities of AI outlined so far can:

- narrow the perceived range of viable approaches,
- present answers before learners have constructed understanding,
- reduce (or even remove) the need to engage with uncertainty or ambiguity.

This creates a form of premature closure, where learners accept or refine outputs without fully interrogating them.

In practice, this may be recognised where learners adopt or rely on AI-generated structures or conclusions with minimal adaptation, show limited evidence of alternative approaches being considered, or are unable to explain how a conclusion was reached beyond reproducing the output.

4.3 Uneven Effects Across Educational Stages

These dynamics are likely to vary, as follows:

- In *early years and schools* there is a risk of over-reliance before foundational skills are established;
- In *further and higher education* there is risk of superficial sophistication—high-quality outputs with limited underlying understanding; and
- In adult learners, there is the potential for strong augmentation where prior judgement frameworks exist.

The key variable is not access to AI, but the maturity of the learner's judgement capacity. This variation across educational stages reflects the Committee's concern that the benefits and risks of AI adoption are not uniform, but depend on learner maturity, context, and the presence of supporting pedagogical structures.

Together, these observations highlight both the key opportunity—AI as a tool for expanding exploratory learning—and the central risk of premature cognitive closure, directly addressing the Committee's focus on how AI can improve educational outcomes while introducing new forms of risk.

5. The Impact on Teaching

5.1 From Knowledge Delivery to Judgement Conditioning

AI reduces the need for teachers to act as primary sources of information. Instead, their role increasingly becomes:

- structuring learning environments,
- guiding how students engage with uncertainty,

- supporting the development of judgement.

Teaching shifts from delivering knowledge to conditioning how learners navigate and evaluate knowledge.

5.2 Evidence from Classroom Practice

In a controlled higher education assessment setting where AI use was mandated:

- All students (2nd year undergraduates) engaged with AI critically rather than passively;
- Many described negotiating with AI, refining prompts and rejecting outputs;
- Students frequently emphasised ownership of their final conclusions; and
- A subset demonstrated clear recognition that AI can support analysis but should not make decisions.

Importantly, this boundary recognition was not explicitly taught, and it did not emerge uniformly. Rather, it appeared through structured exposure to *AI-supported decision-making*.

5.3 Workload and Practice

AI can reduce administrative workload and support content generation. However, the more significant shift is pedagogical:

- assessment design becomes more complex,
- marking must focus on reasoning and process rather than output, and
- teachers must interpret signs of judgement, uncertainty, and ownership.

This may increase cognitive and professional demands on educators, even where routine tasks are reduced. However, arguably, this should play to the strengths of a professional educator.

5.4 Teacher Confidence and Training

Effective use of AI in education requires teachers to:

- understand how AI shapes reasoning processes,
- design tasks that reveal rather than obscure student thinking,
- support students in challenging AI outputs.

Current training provision could best be described as uneven; it often focuses on tools rather than pedagogical implications.

6. The Impact on Learning

6.1 Cognitive Offloading vs Cognitive Formation

AI enables cognitive offloading:

- rapid generation of text,
- structured arguments,
- suggested approaches.

However, education requires cognitive formation:

- developing reasoning,
- engaging with uncertainty,
- constructing understanding.

The risk is that offloading displaces formation.

6.2 Effects on Critical Thinking and Creativity

Where AI is used uncritically:

- critical evaluation may decline,
- originality may be reduced,
- learners may default to plausible but unexamined outputs.

Where AI is used reflectively:

- it can support exploration,
- expose alternative perspectives,
- enhance creative iteration.

The difference lies not in access to AI, but in how learners are required to engage with it.

6.3 Evidence of Variation in Student Engagement

In the controlled higher education assessment setting using mandated AI, observed student behaviour varied significantly:

- Some students used AI as an exploratory partner, challenging and refining outputs;
- Others relied more heavily on AI-generated structures;
- A minority produced highly polished but less reflective work.

This variation suggests both that access to AI does not guarantee learning, and that judgement development is uneven and experience-dependent. It can be illustrated through two common patterns.

In one case, a student used AI as an exploratory partner—iteratively refining prompts, rejecting outputs, and restructuring arguments to reflect their own reasoning. The final submission showed clear ownership and adaptation.

In another case, a student produced a highly polished output with minimal evidence of interrogation or revision; the work was coherent but lacked visible engagement with uncertainty or alternative approaches.

This contrast highlights that output quality alone is no longer a reliable indicator of learning.

6.4 Assessment Integrity and Validity

Traditional assessment methods are increasingly challenged:

- outputs can no longer be treated as evidence of understanding,
- coursework completed outside controlled environments is difficult to interpret,
- detection-based approaches are unreliable and misaligned with educational goals.

Assessment must therefore shift towards:

- requiring explanation and justification,
- evidencing reasoning processes,
- demonstrating ownership of conclusions.

These issues intersect with current UK policy concerns regarding assessment integrity, coursework design, and the reliability of out-of-class assessment, as well as wider expectations from bodies such as Ofqual and Ofsted around critical thinking and digital literacy.

6.5 Digital Literacy and Responsible Use

Digital literacy must evolve beyond tool use to include:

- critical engagement with AI outputs,
- recognition of (knowledge) limitations and uncertainty,
- responsible use and ethical awareness.

Learners must be equipped not only to use AI, but to interrogate and resist it where appropriate.

This has implications for curriculum development, suggesting a need to place greater emphasis on critical engagement, reasoning, and decision-making skills alongside technical competence.

6.6 Inequality and Access

Inequality is not only about access to tools, but access to:

- supportive learning environments,
- effective assessment design,
- guidance on responsible use.

Students with stronger prior educational support, or better guidance, are more likely to develop effective AI-supported learning practices.

This suggests a shift from inequality of access to technology, towards inequality of "pedagogical access": the extent to which learners benefit from well-designed tasks, informed guidance, and teaching that supports effective engagement with AI.

7. Safeguarding and Ethical Considerations

AI introduces new tensions:

- efficiency vs learning depth,
- access vs over-reliance,
- automation vs responsibility.

These tensions cannot be resolved solely through restriction or access policies. Instead, education systems must:

- make the role of AI explicit,
- require accountability for its use,
- preserve spaces where learners must still make decisions under uncertainty.

A central safeguarding concern is the risk that learners become over-reliant on AI systems without developing the underlying skills and judgement required for independent decision-making. Addressing this requires not restriction of access, but careful design of learning and assessment environments that preserve responsibility and cognitive engagement.

8. Government Framework and System Readiness

Current policy discussions are seen to focus on infrastructure, access, and innovation. These are necessary, but insufficient.

There is a need for greater clarity on:

- what educational outcomes AI should support,
- how learning processes are preserved,
- how assessment systems adapt.

Without such clarity, there is a risk that efficiency gains are prioritised over learning quality, and that educational objectives become misaligned with technological capability.

This suggests that current policy frameworks lack sufficient clarity of objective in relation to learning processes, and that quality assurance mechanisms are not yet fully aligned with the realities of AI-mediated assessment and teaching.

9. Recommendations

The following recommendations are offered, based on the argument extended herein.

9.1 Assessment

- ***Redesign assessments to prioritise judgement and ownership*** Require students to demonstrate how they arrived at conclusions, not only what they produce.
- ***Ensure assessment reflects AI-mediated conditions*** Move away from output-based evaluation towards evidence of reasoning, justification, and decision-making.

Example policy lever: Incorporate requirements within assessment guidance that high-stakes tasks include evidence of reasoning processes and student authorship.

9.2 Teacher Capability

- ***Support teachers in AI-informed pedagogy*** Provide training focused on learning design, assessment, and student engagement with AI.
- ***Define responsible AI use in education*** Establish clear expectations that AI may support analysis but not replace student judgement.

Example policy lever: Develop CPD standards for "AI-informed pedagogy" focused on judgement, assessment, and critical engagement.

9.3 System-Level Aims

- ***Adopt a learning-process focus in AI policy*** Ensure policy considers how AI affects reasoning and understanding, not only outputs.
- ***Strengthen digital literacy as critical engagement*** Equip learners to question, challenge, and evaluate AI outputs.
- ***Address inequality in learning environments*** Ensure all learners have access to effective pedagogical support, not only tools.
- ***Preserve spaces for decision-making under uncertainty*** Ensure education continues to develop independent judgement.

Example policy lever: Link EdTech funding and pilot programmes to demonstrable support for reasoning, judgement development, and assessment redesign.

10. Conclusion

AI presents significant opportunities for education, but also introduces a fundamental shift in how learning occurs.

If learners increasingly arrive at answers without engaging in the process of understanding, there is a risk of diminished judgement capacity and over-reliance on AI systems.

The purpose of education is not only to produce correct answers, but to develop the ability to make decisions under uncertainty, take responsibility for those decisions, and engage critically with the world.

Ensuring that AI supports rather than undermines this purpose is central to the future of education.

Without a clear conceptual understanding of how AI reshapes learning processes, there is a risk that promising interventions are replicated at scale without preserving the conditions that make them effective.

This is not simply a question of technology adoption, but of preserving the conditions under which learning—and responsibility—remain meaningful.

A related position paper outlining the underlying research framework is available via Zenodo: Atkinson, D. (2026). Material Entropic Dynamics: Why Faster Prediction Cannot Deliver Freedom in the Age of AI. Zenodo. https://doi.org/10.5281/zenodo.18902944.