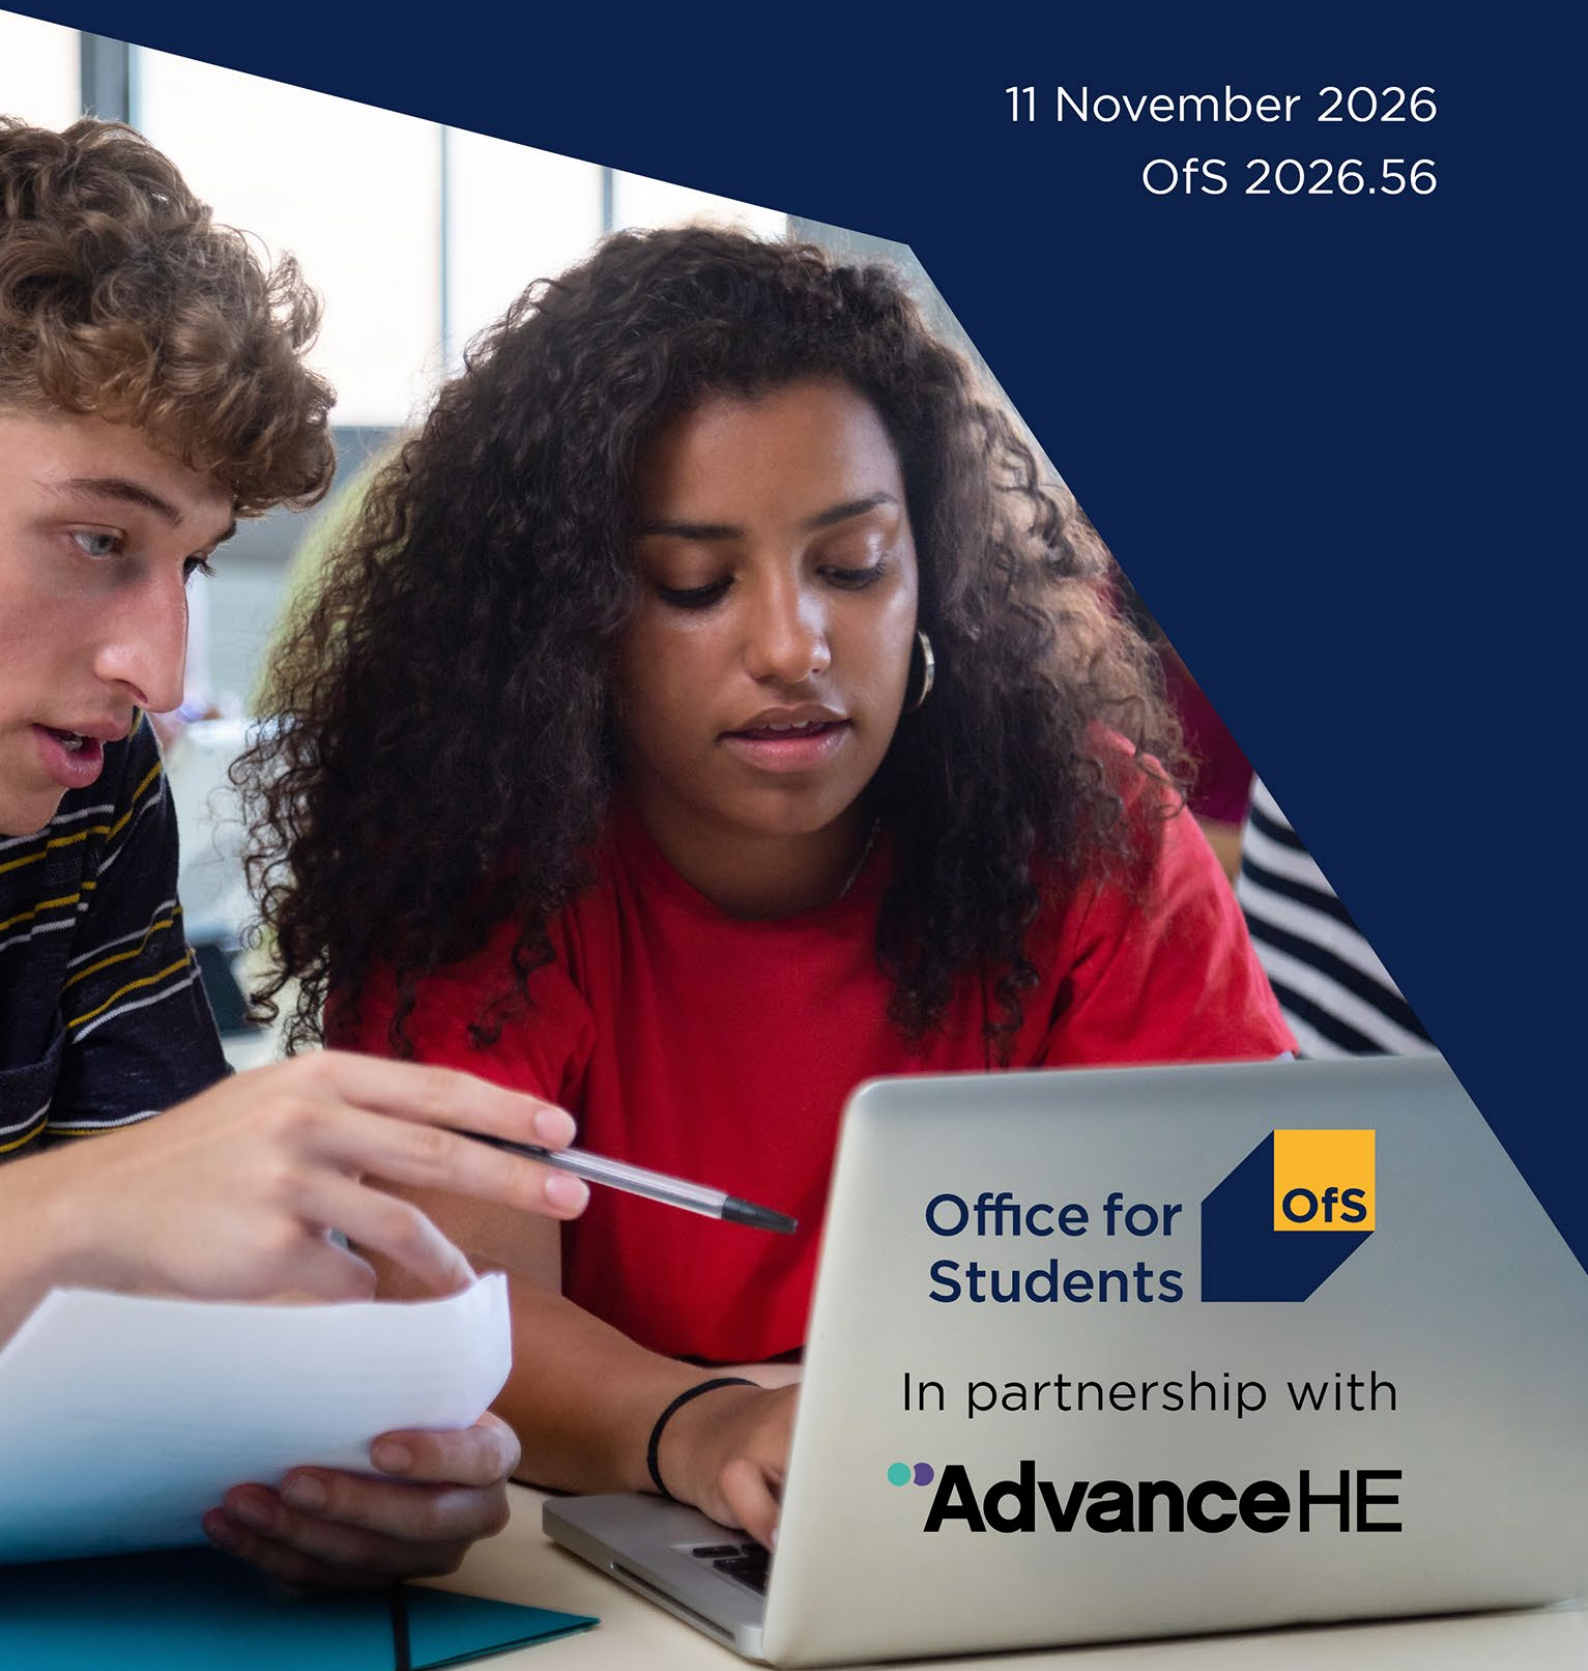
Making AI work for students

Reflections on institutional practice

11 November 2026
OfS 2026.56



Office for Students

In partnership with

AdvanceHE

Contents

Executive summary **3**

Study and teaching practice 3
Assessment 4
AI literacy and skills 4
Strategy and operations 5
Regulatory environment 5

Introduction **6**

Methodology 6
How to read this report 7

Chapter 1: Teaching and study practice **9**

Context 10
Findings 11
- Student use of AI 11
- Institutional uses of AI in teaching 12
- Provision of AI tools 13

Opportunities and risks 14
- Opportunities 14
- Risks 14

Case study: Imperial's Vision for Education in the Age of AI 15

Chapter 2: Assessment **16**

Context 17
Findings 18
- Robustness of assessments 18
- Assessment policies 19
- Assessment redesign 20
- Academic misconduct processes 21

Opportunities and risks 21
- Opportunities 21
- Risks 21

Case study: Two-lane assessment at the University of Bath 23

Chapter 3: AI literacy and skills – Students and staff **25**

Context 26
Findings 27
- Students' AI literacy 27
- Developing student AI literacy 28
- Staff AI literacy and skills 29

Opportunities and risks 30
- Opportunities 30
- Risks 30

Case study: AI platform procurement, University of Kent 31

Chapter 4: Strategy and operations **32**
Context 33
Findings 33
Strategy 33
Efficiency and operations 34
Opportunities and risks 35
Opportunities 35
Risks 36
Case study: AI oversight at Anglia Ruskin 37
Chapter 5: Regulatory and policy environment **38**
Context 39
The impact of the regulatory environment 39
Suggestions for regulatory and other sector bodies 40
Conclusion **42**
Effective institutional practice 43
Study and teaching 43
Assessment 43
AI literacy and skills 43
Strategy and operations 44
Acknowledgements **45**
Annex A: Participating institutions by tariff **46**

Executive summary

Between June and September 2026, the Office for Students (OfS) and Advance HE undertook a collaborative programme of research on artificial intelligence (AI). The research explored primarily English higher education institutions' responses to the opportunities and risks presented by the widespread availability of AI tools, particularly generative AI. The research focused on the implications of AI for teaching and learning, assessment and AI literacy. It also considered institution-wide approaches to AI, including some discussion of operational matters and implications for strategy. The research did not explore the implications of AI for academic research or engage with wider questions about how AI may shape the future of higher education.

This report sets out the findings of our research, based on conversations with over 300 students and staff. It draws on wider academic and policy literature where this supports a richer understanding of institution's approaches. The report considers:

- the ways in which institutions, staff and students are using AI and responding to the opportunities and challenges it presents
- the issues and developments that we think warrant particular attention from institutions.

The report also presents case studies illustrating the practical approaches institutions have taken to addressing the issues we identify, and offers a set of prompts intended to help institutions reflect on and further develop their own approaches to AI.

Study and teaching practice

AI is widely used by students, including to explain difficult concepts, summarise lengthy documents, generate practice questions and support revision. There may be particular pedagogical benefits for students who are neurodivergent or have a disability, or whose first language is not English. Staff adoption of AI into teaching varies significantly. Some academics described embedding AI in their pedagogical approaches, including in the design and delivery of teaching, learning activities and materials, while others use it very little or not at all.

Whether AI improves or weakens learning depends on how it is used. Used critically and carefully, it can help students identify gaps in their understanding and engage with learning in ways that cater to their needs and circumstances. There is headroom for institutions to further exploit these benefits. However, using AI to replace reading, research, problem-solving or other challenging tasks – often called 'cognitive offloading' – can circumvent the productive struggle needed to develop knowledge and skills.

Effective institutional practice in this area is likely to include considering whether and how the integration of AI tools across different courses can support teaching and learning; assuring the quality of AI-supported teaching materials and feedback; and communicating clearly to students where and how AI has been used in the process of developing them.

Assessment

AI poses significant risks to the validity and effectiveness of assessments. Where assessments are unsupervised, it is frequently possible to complete them to a high standard using AI tools without demonstrating the knowledge and skills they are intended to test. Evidence also suggests that AI detection tools do not consistently and reliably identify AI use. This creates risks to the credibility of awards and the reputation of the higher education sector. It also reduces incentives for students to engage in the process of learning.

Approaches to assessments vary widely. While many institutions have introduced AI policies, students are frequently unclear about whether and how they are permitted to use AI in assessed work, and it is usually not possible to conclusively establish whether a student has adhered to their institution's AI policy. Some institutions have redesigned assessments to reduce their susceptibility to AI use, including through increased use of closed-book exams, presentations and other assessments that take place in controlled environments.

Institutions should develop accessible and robust approaches to assessment, which effectively assure student achievement. Secure approaches designed to prevent students from using AI will often be the most effective as they provide a high degree of confidence that students are demonstrating their own knowledge and skills. The use of innovative assessment approaches that reflect the 'real world' application of students' knowledge and responsibly integrate AI will often be positive, but these require thoughtful design to ensure they do not compromise the validity and reliability of assessments as a measure of students' knowledge and skills.

AI literacy and skills

AI literacy refers to the knowledge, skills and capabilities students need to be effective, responsible users of AI while they study and in their future professional and personal lives. Many students see the ability to use AI well, both in general and in the context of their discipline, as important to their future employability.

The provision of AI literacy training for students is uneven. While some institutions have developed ambitious programmes intended to build students' AI capabilities, many students reported receiving limited training or none at all. As the ability to use AI effectively becomes increasingly important to students' success during and after higher education, institutional consideration of the required general and subject-specific capabilities of students also becomes increasingly important Students need opportunities to develop these capabilities, taking account of differential starting points and unequal access to AI tools.

Staff AI literacy is an essential enabler of effective institutional responses to AI, and the differing implications and applications of AI across disciplines means the most valuable AI literacy training will often be subject-specific. At present, AI literacy training for staff is frequently optional and often has low take-up.

Strategy and operations

At many institutions, approaches to AI have been driven by the experimentation of individual staff rather than strategically by senior leaders. Some institutions have established governance groups or transformation programmes, while others have appointed AI leads or ‘champions’ to help support the development of staff AI literacy and steer departmental approaches to AI. However, relatively few institutions that contributed to this research have adopted strategic approaches to AI and embedded these throughout their institutions.

Various examples of how individual staff are using AI to improve efficiency were shared. These include creating content, inputting data, automating routine tasks and deploying student-facing chatbots. However, most staff participants felt AI had not yet delivered significant efficiency gains. Barriers to the systematic integration of AI include lack of confidence in AI outputs, meaning extensive human oversight is required; limited staff capability; immature data infrastructures; and slow institutional approval processes.

Institutions have an opportunity to use AI strategically to improve services and free up staff to spend more time on high-value activities. Realising benefits at scale is likely to require a coordinated approach driven by senior leaders while ensuring individual staff and teams can continue to innovate in their own practice.

Regulatory environment

Participants called for sector, regulatory and government bodies to adopt a collaborative approach to AI, working with institutions to support improvements in their approaches. Many highlighted the value of sharing effective practice in addressing common challenges. Participants generally suggested that new conditions of registration from the OfS on AI could be counterproductive and would risk constraining innovation. Instead, some participants called for greater regulatory clarity on complying with existing conditions of registration in the context of AI.

Introduction

1. This report sets out the findings of a collaborative research project undertaken by the Office for Students and Advance HE. The organisations have distinct roles: the OfS regulates higher education in England, while Advance HE is a member-led organisation supporting enhancement in over 30 countries across the globe, including helping institutions respond positively to regulation.

2. Since the launch of ChatGPT in late 2022, use of generative AI (GenAI) chatbots and other AI tools has become widespread in higher education. The proportion of students who report having used generative AI tools to help them with assessments has risen from 53 per cent in late 2023 to 94 per cent in late 2025.[1] Many specialist tools and apps across fields such as medicine, law and the creative industries now build in AI automatically, as do many online search engines and generic tools students use every day. Staff across higher education increasingly use AI tools to support teaching, prepare learning materials, provide feedback and complete administrative tasks.[2]

3. These developments have prompted extensive commentary and research, considered throughout this report. We hope to add to existing evidence by reflecting the experiences of over 300 students and staff, from over 120 institutions, to provide an integrated picture of how the sector is navigating the changes brought by AI. We identify where institutional approaches are developing effectively and where significant challenges remain, and highlight the practical steps institutions are taking in response. We hope this will help institutions reflect on their approaches, adopt effective practice and avoid common pitfalls. The research findings will inform the OfS's regulatory approach and shape Advance HE's development and enhancement services to support institutions' ongoing responses.

4. This report considers sector approaches to AI in relation to teaching and study practice, assessment, AI literacy, and strategy and operations, with a particular focus on generative AI tools. It does not examine the impact of AI in relation to research or address broader questions about the longer-term implications of AI for the future of higher education.

5. The impact of AI also varies significantly by course. In some fields, such as medicine, computing, or mathematics, AI is transforming curricula and professional practice. In other subjects, the impact of AI may be smaller or only beginning to be understood, while essay-based subjects may face the greatest challenges to assessment effectiveness. Where appropriate, we have highlighted differences by subject indicated in the research.

Methodology

6. This report presents the findings of our engagement with students and staff from universities and colleges between June and September 2026. We held 16 live events, including seven in-person roundtables in London, Manchester and Bristol. We held four online roundtables, two in-person workshops at the 2026 Advance HE Teaching and Learning Conference, and three

[1] See Higher Education Policy Institute (HEPI), 'Student generative artificial intelligence survey 2026' and 'Provide or punish? Students' views on generative AI in higher education'.

[2] See Jisc, 'Staff perceptions of AI 2025'.

listening events as part of the OfS's annual programme of non-regulatory visits. As participants were self-selecting, they may not be fully representative of wider student and staff populations.

7. We also invited responses to a series of open-ended, online prompts, generating 64 responses, and conducted 22 interviews with representatives of higher education institutions, sector groups, technology companies and graduate employers. In total, the views of 244 higher education staff and 91 students and student representatives from at least 127 institutions have informed this work. We are grateful to everyone who participated in this research for their constructive and open engagement with us.

8. The research included participants and organisations operating across the UK, recognising that while OfS regulation applies only to English higher education, Advance HE has a broader geographical remit. Our in-person roundtables and student listening events were open to participants from English institutions, while the Teaching and Learning Conference and online roundtable events were attended by participants from across the UK nations.

9. Alongside the primary research, we draw on an extensive wider policy and academic literature relating to AI in higher education. Given the rapid pace of AI development and possible variation between England and other higher education systems, we sought to use recent literature covering England or the UK where possible.

10. Research findings were analysed by a team of researchers from the OfS and Advance HE. We also shared the draft report with a small group of experts on AI in higher education for comment.

11. The full list of organisations interviewed is in Annex A. The views set out in this report are those of the OfS and Advance HE and do not necessarily reflect the views of research participants. The views set out are conclusions based on the research and not regulatory requirements set by the OfS.

12. While we hope this report makes a meaningful contribution to understanding AI in higher education, it is not intended to be the final word. We hope it will prompt further discussion, reflection and research as the technology continues to evolve, and institutional approaches continue to change with it.

How to read this report

13. This report has five chapters:

    a. **Teaching and study practice.** This chapter examines how students and staff use AI to support learning and create teaching materials, and how institutions are embedding AI tools in their provision.

    b. **Assessment.** This chapter examines the challenges AI presents to the validity and security of assessment, and how institutions are responding through policies, assessment redesign and academic misconduct processes.

    c. **AI literacy and skills.** This chapter examines the knowledge, judgement and practical capabilities students need to use AI effectively, critically and responsibly. It considers how

institutions are developing these capabilities in their curricula and through extra-curricular approaches. It also considers the approaches institutions are taking to develop AI literacy among staff.

d. **Strategy and operations.** This chapter examines the extent to which institutions are developing and deploying strategic approaches to AI, including its use in administration and student services.

e. **Regulatory and policy environment.** This chapter examines how regulation, professional requirements and sector guidance are shaping institutional responses to AI.

14. A summary of the findings from the research, and key issues for institutions to consider as they develop their approaches to AI, are discussed in the conclusion.

15. Reflections on the impact of AI on equality of opportunity are integrated throughout the five chapters.

16. The first four chapters are structured as follows:

    a. **Context** introduces the key issues by setting out selected, established evidence and wider developments where relevant.

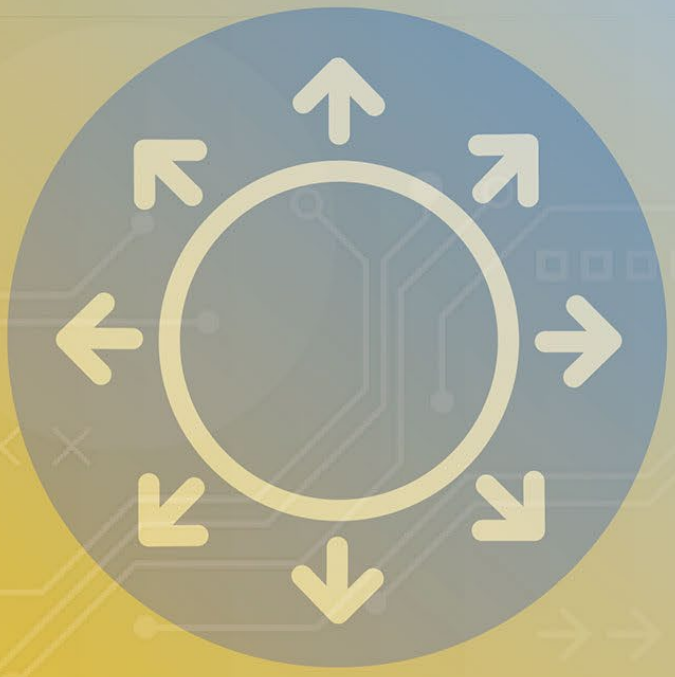
    b. **Findings** describes how institutions, staff and students are using AI and responding to the opportunities and challenges it presents, drawing on the primary research conducted for this report.

    c. **Opportunities and risks** reflects our view of issues and developments that warrant particular attention, whether because of their potential benefits or the challenges they present.

    d. **Case studies**, which are included throughout the report, illustrate practical approaches institutions have taken to address some of the issues we identify.

17. The fifth chapter summarises research participants' views on the impact of the policy and regulatory environment and explores the changes they believe would enable universities and colleges to make the most effective use of AI.

Chapter 1: Teaching and study practice



Context

18. AI is not the first technological development to disrupt existing methods of teaching and learning. However, generative AI combines the following features in ways that distinguish it from earlier technological developments:

    a. **Speed.** Substantial outputs can be produced in seconds.

    b. **Scale.** The technology is available to millions of users simultaneously, at relatively low cost, and without requiring specialist knowledge or skills to use it.

    c. **Scope.** The technology has a very wide range of uses including research, writing, coding, design, planning, feedback and facilitating reflection.

    d. **Specificity.** Outputs can be tailored to an individual task and repeatedly refined through conversation.

    e. **Separation.** The final product can become separated from the learner's own knowledge, reasoning and capability.[3]

19. Existing research suggests AI can support personalised learning and reduce barriers to learning for some student groups in particular.[4] However, an emerging body of evidence also suggests that depending on AI to complete tasks can displace the cognitive effort required for deep learning, a process sometimes called 'cognitive offloading'.[5] Studies of school pupils have found that students using AI initially outperform those who do not, but perform worse in subsequent closed-book assessments.[6] The overall impact appears to depend heavily on how AI is used by teachers and learners and the learning context in which it is deployed.

20. AI tools are widely used by students during their studies. Surveys suggest that between 74 and 94 per cent of students now use AI to help them complete academic work.[7] Students are using AI tools to explain difficult concepts, teach them how to solve problems, summarise literature, generate and edit images, and draft, edit and structure text.[8]

---

[3] See Future Proof Teaching, 'AI cheating isn't the real crisis. It shows how fragile some assessments have become'.

[4] See Vieriu, Aniella Mihaela, and Petrea, Gabriel, 'The impact of artificial intelligence (AI) on students' academic development'; Cooper, Grant et al, 'Using generative artificial intelligence to reimagine feedback in higher education: A collaborative autoethnography';, HEPI, 'How AI-generated feedback could assist student learning'; Zhao, Xin et al, 'The use of generative AI by students with disabilities in higher education'.

[5] See Kosmyna, Nataliya et al, 'Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task'; Lodge, Jason M et al, – 'Artificial intelligence, cognitive offloading and implications for education'; Gerlich, Michael, 'AI tools in society: Impacts on cognitive offloading and the future of critical thinking'.

[6] See Stromberg, David et al, 'The Generative AI Learning Penalty: Evidence from Chinese Secondary Education'; Bastani, Hamsa et al, 'Generative AI without guardrails can harm learning: Evidence from high school mathematics'.

[7] See HEPI, 'Student generative artificial intelligence survey 2026'; YouGov, 'How are UK students really using AI?'.

[8] See HEPI, 'Student generative artificial intelligence survey 2026'.

21. Research also suggests AI tools are now widely used by teaching staff in their work. Jisc research from the 2024-25 academic year found that more than a third of teaching staff in the UK used AI tools in their work, with more recent evidence from the USA suggesting 90 per cent do so.[9] Initial research in this area suggests AI tools can produce learning materials of comparable standard to human staff, but are prone to inaccuracies and require iteration.[10]

22. Existing research raises concerns about unequal access to AI tools. Jisc research reports concerns among students that the ability to pay for premium versions of AI tools gives some students a 'competitive edge'.[11] Free plans of generative AI models typically have strict daily caps, while paid plans grant higher usage limits and often access to more powerful or specialised models. Surveys of students also suggest significant differences in AI usage by socioeconomic background.[12]

Findings

Student use of AI

23. Most student participants reported using AI in some way to support their studies and described a high proportion of their peers doing the same. Uses of AI include explaining difficult concepts and technical terminology, helping them grasp challenging material more quickly; summarising lengthy documents, such as lecture notes and research papers, helping them to identify key information more quickly and manage large volumes of reading; and generating practice questions and quizzes, supporting revision.

24. Some student groups told us they find AI tools particularly valuable in reducing barriers to learning:

    a. Neurodivergent students and students with disabilities described using AI for transcription, spelling support, translating notes into braille, explaining marking criteria in simpler language and helping with revision planning.

    b. Students whose first language is not English described using AI to help them engage with complex academic language and understand institutional policies and processes.

25. Student use is not universal. Some students, sometimes described by staff as 'conscientious objectors', avoid using AI for principled reasons including its environmental impact, concerns over how large language models are trained and perceived threats to human creativity. Some students described avoiding AI tools because of previous experiences with academic misconduct processes, a lack of confidence, a lack of clarity around permitted use or concerns it would negatively impact their learning. Some staff reported student protests in response to AI integration and described a polarised atmosphere.

[9] See Jisc, '2024-25 UK higher education teaching staff digital experience insights survey findings'; Educause, 'The Impact of AI on Work in higher education'.

[10] See Leinonen, Juho et al, 'AI-generated slides: Are they good? Can students tell?'.

[11] See Jisc, 'Student perceptions of AI 2025'.

[12] See HEPI, 'Student generative artificial intelligence survey 2026'.

26. Some students said there is a stigma around AI use among students, with AI use taken to imply minimal effort or inadequate capability. A small number described feeling pressure from their peers not to use AI, particularly students studying creative subjects. This discourages some students from being transparent about how and where they use AI. It also discourages some from discussing AI with their peers, although others, including students studying medicine and computer science, described open discussion about AI on their courses.

27. Both staff and students raised concerns that some uses of AI may be reducing student engagement with parts of the learning process. These uses include overdependence on AI summaries at the expense of direct engagement with core academic materials, which can lead to surface-level understanding of materials and subjects; relying on AI rather than using study materials provided by teaching staff; using AI to search for key texts, rather than reading and deciding themselves whether a text is relevant; and to replace productive struggle in many other areas. As a result, many research participants suggested AI use is causing students' learning to atrophy.

28. This was considered particularly detrimental during the earlier stages of higher education, when students would normally be developing foundational knowledge. Some staff also suggested that by pushing students to focus primarily on the output, rather than the process of learning, AI was diminishing students' enjoyment of study.

29. Some staff described students whose first language is not English using AI to translate work completed in their preferred language into English. They raised concerns that this could reduce students' engagement with English and, over time, weaken the value of an English higher education qualification as a signal of English language proficiency.

> 'One big risk is inappropriate cognitive offloading and use of AI, particularly GenAI to replace critical engagement with study materials and course content. On the other hand, there are opportunities to enhance learning and teaching if AI is truly used as a companion/assistant for both staff and students.'
>
> (Staff respondent, online survey)
>
> 'What GenAI promises to students (and to society at large) is quick outcomes without effort or struggle; yet we know that struggle and failure are instrumental to growth in education.'
>
> (Staff respondent, online survey)

Institutional uses of AI in teaching

30. Teaching staff are using AI to support teaching and learning in a range of ways, including generating practice questions and quizzes, creating case studies, producing model answers for critical evaluation, generating a diverse range of learning resources, such as podcasts or flashcards, and generating feedback based on recorded lesson observations.

31. Some universities and colleges are using external, AI-enabled tools to support teaching and learning at the institution level. Examples include tools that provide students with on-demand learning support and feedback and the integration of AI tools into virtual learning environments.

32. There is significant variation in staff approaches to AI, often within the same institution or programme. Some staff allow or encourage students to use AI where it supports their learning, others discourage or prohibit students from using it and others do not discuss it at all. Combined with the absence of clear policies or guidance at some institutions, this means students can struggle to navigate the requirements and expectations of teaching staff.

33. During our research there was significant debate about whether teaching staff should use AI in the production of feedback. Students were often critical of AI being used in this way, frequently because it was poor quality or because they suspected staff had not been transparent about whether or how AI was used. Some described access to feedback from expert teaching staff as a key part of higher education's value proposition. Conversely, staff participants suggested that AI-generated feedback could be a beneficial starting point where it reduces the time spent producing feedback, allowing more time to be spent with students. One staff member had trained a pedagogically informed AI tool to mark student work and shared the source code of the tool with students.

> 'You want proper feedback from your teaching figure, because if you wanted AI feedback, you could just go and get it yourself, like just go and ask AI yourself.'
>
> (Student participant, online roundtable)

Provision of AI tools

34. Most research participants reported that their institution provides AI tools for staff and students, most commonly Microsoft Copilot Chat as this is included within existing Microsoft 365 licensing arrangements. Participants highlighted the advantages of Copilot's closed and secure environment, which enables institutional data to be used within the system while maintaining compliance with data security and information governance requirements. Some institutions have signed agreements to provide students and staff with to other AI tools, including the University of Oxford, the London School of Economics and the University of Kent, as these tools are considered more powerful or have dedicated education versions.

35. Some staff and students at institutions that had not purchased licences reported using tools they considered more powerful than those provided by their institution. However, some staff we spoke to raised concerns that colleagues are unaware of the potential risks associated with inputting sensitive data, including student data, into free online tools where it may be used to train AI models.

36. Participants raised concerns that students who can access more powerful tools would be able to benefit from additional features and functionality, widening existing inequalities as AI becomes increasingly integrated into students' study practices.

Opportunities and risks

Opportunities

37. AI presents many opportunities to enhance teaching. AI can help staff produce high quality learning materials and resources more efficiently, improve accessibility and keep content up to date. By generating instant, personalised feedback on teaching resources or recordings of live teaching, it can support staff to make improvements to their teaching practice. When trained on course content, AI tools can also help ensure students have access to accurate information relevant to their studies. This has the potential to drive pedagogical innovations in teaching with benefits to teaching quality.

38. When used critically and carefully, AI has the potential to improve learning by providing students with immediate, personalised, tailored support and feedback. This may help students identify areas for improvement, address gaps in their understanding and engage with learning in ways that reflect their individual needs and circumstances. The ability to draw on AI instantly, at any time, has the potential to help address historical imbalances in access to tutoring and personalised support.

39. AI reduces barriers to seeking academic support and allows for the provision of support at near zero marginal cost. This is likely to be particularly beneficial for students with diverse learning needs, such as disabled students or those whose first language is not English, both in their academic studies and across their wider student experience.

Risks

40. Academic research suggests that whether AI improves or weakens learning depends on how it is used. The offloading of challenging tasks to AI tools, such as summarising reading materials, conducting research, solving problems, drafting text, or writing in English when it is not a student's first language, can weaken learning by circumventing the productive struggle needed to develop knowledge and skills.[13] Use of AI in this way presents a serious risk to student learning. Recent research suggests assessment design can create incentives for students to use AI in ways that support learning rather than replacing it.[14] The role of robust assessment practice that incentivises deep engagement by students with their course materials is discussed further in Chapter 2.

41. There is a risk that students accept AI-generated outputs uncritically, without assessing their accuracy or recognising their limitations. AI literacy and subject knowledge are complementary: students can only critically evaluate AI outputs when they have a deep understanding of their course content.

42. As the gap in performance between the most powerful paid-for tools and the free versions can be large, where AI is used to support learning, students with access to more powerful tools may derive greater educational benefits. However, where AI is used in place of learning,

[13] See Stromberg, David et al, 'The Generative AI Learning Penalty'; Kosmyna, Nataliya et al, 'Your brain on ChatGPT'.

[14] See WonkHE, 'Trained to stop learning: How students are experiencing assessment and learning in an age of AI'.

students with access to more powerful tools may be more likely to offload the cognitive effort required for learning.

43. Without adequate scrutiny, AI tools create risks to the quality of teaching materials. The speed and scale of AI adoption in higher education is likely to expose existing weaknesses in systems for assuring quality. If AI-generated content is not subject to critical review, inaccuracies, bias and weak pedagogical design may become embedded within the curriculum. This risks reducing the quality of teaching and undermining confidence in academic standards, particularly where institutions and academics are not transparent about where and how AI has been used.

Case study: Imperial's Vision for Education in the Age of AI

Imperial's response to generative AI was spearheaded by seven AI 'futurists', who together represent the institution's faculties and central education teams. The futurists led conversations with staff and students on the institution's approach to AI, which contributed to a published vision for Education in the Age of AI. This vision is now being implemented via an in-depth transformation programme lead by a new Centre for AI in Learning and Teaching, and in collaboration with a network of Departmental AI in Education Leads.

Imperial provides a GenAI workspace, dAIsy, through which students and staff have access to multiple generative AI models such as ChatGPT, Claude and Deepseek through a single, data-secure interface. It is also launching 40-hour courses for students and staff focusing on the application of AI in the workplace, culminating with a capstone project.

Finally, Imperial is developing discipline-specific pedagogical approaches that will ensure students graduate with the ability to work confidently, critically and responsibly with and without AI.

Chapter 2: Assessment



Context

44. AI tools are extremely powerful and are continuing to improve rapidly. Attempts to benchmark the performance of AI tools show the most powerful tools perform better than human experts at a range of tasks, including answering PhD-level science questions, and the recent release of AI-generated mathematical proofs demonstrates AI's potential to advance the frontiers of human knowledge.[15] In many cases, with only limited input, AI tools can be used to produce outputs that score highly in a range of higher education assessments.[16]

45. The power, speed and accessibility of AI chatbots, and their ability to generate responses tailored to specific tasks, pose serious challenges to designing effective high-stakes assessments. If it is possible to produce a passing or even high mark in an assessment with only limited student engagement, the assessment no longer effectively tests students' knowledge and skills. If widely repeated, this would raise significant concerns about the credibility of awards issued by higher education institutions that depend on these assessments.

46. Academic research suggests that humans cannot reliably identify AI-generated text.[17] Some companies have developed AI detectors, which aim to establish the proportion of text that was AI-generated, and there is evidence that some newer detectors may be able to identify AI-generated and modified text more accurately than established tools.[18] However, research suggests that improvements in detection are often temporary and typically followed by advances in AI systems that make AI-generated text harder to identify.[19] Advances in so-called 'adversarial' strategies that aim to trick AI detectors into rating output as human generated, such as the use of 'humaniser' tools, may further reduce their accuracy.[20] While some technology companies have deployed techniques described as 'watermarking', which mark text as AI-generated, the impact of these techniques is not yet clear.[21]

47. The limits of human and tool-based detection mean it is extremely difficult to determine whether and how far AI was used in producing student work. This poses a risk to established

[15] See Stanford University Human-Centred Artificial Intelligence, 'The 2026 AI index report: Technical performance'; OpenAI, 'On the Navier–Stokes Millennium Prize Problem' and 'Sharing AI progress in mathematics'.

[16] See Scarfe, Peter et al, 'A real-world test of artificial intelligence infiltration of a university examinations system: A "Turing Test" case study'; Ibrahim, Hazem et al, 'Perception, performance, and detectability of conversational artificial intelligence across 32 university courses'.

[17] See Fiedler, Alexandra, and Döpke, Jörg, Do humans identify AI-generated text better than machines? Evidence based on excerpts from German theses'; Cheng, Adam et al, 'Ability of AI detection tools and humans to accurately identify different forms of AI-generated written content'.

[18] See Van Vlasselaer, Marijke et al, 'Who wrote this? Evaluating the reliability of AI detection tools in higher education'.

[19] See Li, Qizheng et al, 'Dialectical analysis and developmental considerations of the "Cat-and-mouse game" in AI-generated text detection'; Laurier, Linda et al, 'the cat and mouse game: the ongoing arms race between diffusion models and detection methods'.

[20] See Al Roken, Noora et al, 'On adversarial attack detection in the artificial intelligence era: Fundamentals, a taxonomy, and a review'; Perkins, Mike et al, 'Simple techniques to bypass GenAI text detectors: Implications for inclusive education'; Liang, Weixin et al, 'GPT detectors are biased against non-native English writers'.

[21] See Anthropic, 'How Claude's text watermarking works'.

methods of detecting plagiarism, in which software or an assessor review a completed assessment. Research suggests there was a significant increase in AI-related academic misconduct cases up to the 2024-25 academic year.[22]

48. Early research suggests there is the potential for AI tools to play a constructive supporting role in high-stakes marking, but current AI tools are not ready to replace human judgement or act as the sole mechanism for determining a student's mark.[23] An early 2026 study testing generative AI tools' marking of student essays found significant gaps between the grades awarded by the tools and by human reviewers.[24]

Findings

Robustness of assessments

49. A significant proportion of assessments used in the sector remain in unsupervised, non-controlled formats where AI can be accessed by students completing them, and many of these can be completed to a high standard using AI tools with limited engagement from students. Programmes assessed through remotely submitted essays were described as being particularly vulnerable. Student research participants frequently told us that students they knew had completed assessments using AI, sometimes in just a few minutes, and received a similar grade to those who had completed the assessment themselves. Some students described their peers being unable to recall basic details about the assessment they had just submitted.

> 'I've had people admit, peer-to-peer, "I wrote this whole essay with ChatGPT." […] Then they don't get caught, they get a good grade probably, and then because it worked the first time, they do it on their next assessment.'
>
> (Student participant, online roundtable)
>
> '… It is incredibly disconcerting when I plan, and write an essay completely by myself, and get the exact same grade as someone who used AI to write it.'
>
> (Student respondent, online survey)
>
> 'I've sat at the back of a lecture and seen people on ChatGPT writing assignments during a lecture for a different assignment. And it's like, why are you doing this? Perhaps if you'd paid attention in the lecture that you're currently using AI to help you write […] you wouldn't need [AI].'
>
> (Student participant, online roundtable)

50. Many staff participants told us they believe AI is used extensively in the process of completing assessments. Some described receiving high volumes of submissions that look the same,

[22] See The Guardian, 'Revealed: Thousands of UK university students caught cheating using AI'.

[23] See Gov.UK, 'Principles of AI use in marking'.

[24] See University of Cambridge, 'AI not yet good enough to mark university essays, rewarding "style over substance"'.

reference the same literature and are written in the same style. Some said explicitly they believe the assessments run by their institution are not fit for purpose or do not accurately measure students' knowledge and skills. Some saw the risks that AI poses to effective assessment as an opportunity to review assessment practice more broadly and ensure that it effectively measures the intended learning outcomes of each course.

> 'In our School, I think we're realising that a number of our assessments may be unfit for purpose, and that pervasive use of AI may render our assessment of students close to impossible in some areas.'
>
> (Staff respondent, online survey)
>
> 'AI has put academic integrity at serious risk. […] Coursework is now overwhelmingly written in part, or exclusively, with the help of AI. Presentation slides fare no better.'
>
> (Staff respondent, online survey)

51. Participants raised concerns that, given the large capability gaps between the most powerful, paid tools and the cheapest or free versions discussed in Chapter 1, a student paying for the most advanced model may receive a higher grade than a student with equivalent knowledge and skills using a cheaper model or not using AI. More powerful AI models may also be better able to evade detection tools and 'humaniser' tools, which rephrase text to avoid detection, also often have paid versions.

52. Some participants described how tools provided to particular students as part of reasonable adjustments, for instance to check punctuation and grammar, can contain powerful AI features not available to others. They were concerned that, rather than levelling the playing field, these features could confer a significant advantage in assessments.

Assessment policies

53. Most institutions we spoke to have assessment policies or guidelines setting out how students can use AI, which may be set at the institution, department or programme level. Many use 'traffic light' systems or assessment scales, which classify assessments by whether, and how, AI may be used in producing them: red is prohibited, green is allowed or encouraged, and amber allows some uses but not others.

54. Many student participants, however, said they were unclear about when and how they were permitted to use AI in assessed work. Some reported rapidly changing policies or policies introduced partway through the academic year; unclear, poorly communicated policies that vary across departments and staff; or no formal institutional policy.

55. Some institutions require students to declare how AI was used, for example through signing an 'AI declaration' or disclaimer at the point of submitting an assignment. Staff participants expressed concern that students do not accurately disclose how they use AI. Many students told us they would not declare they had used AI, even if they had done so in accordance with the AI policy, for fear that this would cast suspicion on their work and lead to them receiving a lower grade.

Assessment redesign

56. Some institutions have undertaken exercises to understand the vulnerability of their assessments to AI misuse. One participant, for example, told us their institution had 'stress-tested' assessments by having staff try and complete them using AI tools.[25] Another described developing a diagnostic tool to determine the vulnerability of their institution's courses based on assessment type and recommend mitigations.

57. Many institutions have redesigned their assessments to reduce their susceptibility to AI misuse. Approaches include:

    a. Using 'secure' assessments, assessments designed to prevent students from using AI while completing them. These include in-person written exams, presentations or vivas, short discussions to test students' understanding of their submitted work, and practical or performance-based assessments.

    b. Process-based assessments, which assess how students produce their work alongside assessing the final output. Some participants described approaches that assess students' skills – for example, their ability to complete a literature review or summarise a text – at multiple checkpoints or 'stage gates', to verify students have gained the intended knowledge or skills.

    c. Assessments where AI use is permitted or required, such as critiquing and interrogating AI-generated output or using AI to generate a first draft, which students edit and improve. These assessments may be used to evaluate students' AI literacy.

58. Alongside being secure, staff participants stressed that assessments need to be effective, meaning they accurately assess how far students have met the intended learning objectives. Some participants described undertaking curriculum redesign in parallel with changes to assessments to ensure alignment between the two.

59. We heard that redesigning assessments to be secure presents challenges. Barriers cited include resourcing concerns, with in-person exams and vivas being particularly resource-intensive. We also heard concerns that some secure assessments, such as closed-book exams, bear little resemblance to the activities students will complete in professional environments, and concerns that their structured format and rewarding of long-term memory may disproportionately disadvantage some student groups, such as disabled students. Professional, statutory and regulatory bodies may also prescribe particular assessment approaches, or require institutions to seek their approval before changing them. The increasing development of wearable technology may increase the feasibility of cheating in secure assessments.

60. Participants said that on many courses, assessments have not been redesigned and remain in forms susceptible to AI misuse, such as an open-book essays or problem sets.

[25] See HEPI, 'Using AI to think with students, not for them'.

Academic misconduct processes

61. Many institutions have adapted their academic misconduct processes. Some institutions conduct vivas or disciplinary hearings with students suspected of academic misconduct to verify whether the student has a deep understanding of their submission.

62. Despite research suggesting reliability issues with AI detection tools, many staff we spoke to said their institutions use them. Research participants said that some institutions include scores produced by AI detection tools as supporting evidence in academic misconduct processes, while others were aware of colleagues using them informally to identify assessments, which may then be flagged for misconduct or awarded a lower grade. Some institutions had stopped using AI detection tools because of concerns about their efficacy.

63. A spokesperson for Turnitin, which operates an AI detector widely used in the UK sector, told us that its tool's output should be used as only one data point when working to understand how work was authored.

64. Student participants were strongly critical of AI detectors, arguing they are inaccurate, overstate AI use and cast suspicion on student work even when they do not provide strong evidence of AI use. Some students were frustrated that once an AI detector had given them a high score, there appeared to be no route to proving their innocence. Students frequently said it was unfair that some students were penalised for using AI while other students were not. Students described a variety of strategies to lower their AI detection score, including removing characteristics of AI-generated text such as em-dashes, using 'humaniser' tools or deliberately adding spelling and grammar errors.

Opportunities and risks

Opportunities

65. AI has created an exciting impetus to update assessment approaches and strategies so they more effectively test students' knowledge and skills. As the OfS has previously set out, there are opportunities for AI to drive pedagogical innovation in assessment design, including through the use of new technologies.[26]

66. As the technology develops, it is possible that AI will be able to support the accurate and consistent marking of student work. Innovation by institutions in developing effective, reliable approaches to AI marking is welcomed, although any use of AI would require careful testing and oversight to ensure that poor implementation does not undermine marking quality or students' access to high-quality feedback and support.

Risks

67. Ensuring high-stakes assessments are effective in the age of AI is a generational challenge. As many assessments currently used by higher education institutions are unsupervised and not held in controlled conditions, it is frequently possible for students to use AI tools in place of their own knowledge and skills in ways that cannot be reliably detected. The resulting scores are therefore likely to reflect students' access to AI tools and their proficiency in using them, rather than reliably demonstrating the knowledge and skills the assessments are

[26] See OfS, 'Embracing innovation in higher education: Our approach to artificial intelligence'.

intended to measure. This presents a significant risk to the credibility of assessments awarded by higher education institutions and the reputation of the sector.

68. To date, many changes to assessments have focused on giving students instructions rather than making structural changes to the way assessments are designed. It is important that institutions set out clear and consistent AI policies and support students to adhere to them. However, without structural changes, it is not possible to verify whether students have followed those policies or to meaningfully deter AI misuse.[27] Additionally, the boundaries between acceptable and unacceptable uses of AI are often difficult to define in practice, as AI tools can generate ideas and reasoning in ways that blur the distinction between students' own work and AI assistance.[28]

69. The effective assessment of students depends on robust assessment approaches that provide assurance students have achieved the intended learning outcomes. Secure assessment methods, such as closed-book exams, presentations and vivas, performances, and forms of practical and observed assessments, will often be the most effective forms of assessment because they provide a high degree of confidence that students are demonstrating their own knowledge and skills. While the use of AI may still be hypothetically possible, for instance through wearables, it is much more difficult than in unsupervised assessments.

70. The use of non-secure assessments, during which students can use AI, creates risks to student learning and the credibility of assessment. Our research shows that many students now see AI as the most effective route to a high grade, and students may feel they are at a disadvantage if they do not use AI. This is creating pressure to outsource more of the assessment process to AI tools even when students believe that doing so is detrimental to their learning.

71. There is an urgent need to ensure that any non-secure assessments maintain their validity despite the availability of AI tools, meaning they effectively assess the knowledge and skills they are designed to measure. If it becomes possible to complete an assessment using AI without the knowledge and skills being tested, the assessment should be replaced. The rapid development of AI and wide variation in the capabilities of different models means ensuring non-secure assessments are robust is likely to require thorough and ongoing testing.

72. Whether assessments are secure or non-secure, they should enable all students to demonstrate their knowledge and skills. This includes removing avoidable barriers through inclusive design and making reasonable adjustments as required.

73. Securing assessments against AI misuse may sometimes involve trade-offs, including the potential for reduced ‘authenticity’. Efforts to develop assessments that are authentic while ensuring assessment is accessible and robust to AI are welcome. A typology of secure assessments developed by the University of Melbourne includes a range of practical options alongside more traditional approaches such as closed-book exams, and a set of case studies

[27] See Corbin, Thomas et al, ‘Talk is cheap: Why structural assessment changes are needed for a time of GenAI’.

[28] See Slimi, ’Frontiers | A systematic critical review of generative AI's impact on authorship, pedagogy, and integrity (2023–2025)’.

from Curtin University includes a mix of secure and non-secure assessments that may be effective in the context of AI.[29] The feasibility and efficacy of these approaches will vary by subject and course. Assessing students' knowledge, rather than skills, will usually require assessments to be secure.

74. There is a risk that some institutions and courses attempt to use misconduct processes to retrospectively fix the problems caused by assessments that are not resilient to AI. These processes are time- and resource-intensive, stressful for students, and unlikely to be effective in securing the validity of an assessment process. As it is difficult to conclusively establish whether AI was used, there is a further risk that students subject to academic misconduct processes feel they have not been treated fairly. Assessments should be designed to minimise the opportunities for academic misconduct and facilitate the detection of misconduct where it does occur.

75. There is a risk that some institutions put disproportionate weight on the outputs of AI detection tools to determine whether AI was used in producing an assessment, despite evidence they are often ineffective. The output of detection tools should be, at most, one piece of evidence to support a thorough assessment of whether a student has committed academic misconduct. Relatedly, suspicion of AI use without strong evidence should not be used as a basis to award students a lower mark. While there may be exceptions, effective assessment design will frequently mitigate the risks presented by AI better than retrospectively attempting to verify AI input.

Case study: Two-lane assessment at the University of Bath

The University of Bath had previously adopted a traffic light approach to assessment, with three categories:

a. AI entirely prohibited.

b. Some specified uses of AI allowed.

c. AI has an integral role in the assessment.

As some staff sought to prohibit or overly regulate AI use in coursework assessments, it became clear that it was not possible to police specified uses of AI in Category B assessments in a way that was reliable and equitable.

In response, in September 2026, Bath adopted a 'two-lane' approach to assessment first developed at the University of Sydney. Assessments must fit either into a 'closed' lane, where assessment conditions ensure that students cannot access AI tools, or an 'open' lane, where the use of AI tools is permitted. Open-lane assessments must therefore be designed assuming students may use AI tools.

[29] See University of Melbourne, 'Secure assessment'; Curtin University, 'Educator case studies: Explore innovative assessments'.

Recognising the challenges posed by increasingly capable AI tools, the university has focused on helping staff design robust open-lane assessments that promote meaningful learning. This has included a retreat for academics to share and develop innovative approaches to assessment for students navigating an AI-rich context, including informed avoidance.

The university emphasised that all assessments, including closed-lane assessments, should be designed around intended learning outcomes. It is not sufficient for an assessment to be ringfenced from AI use; it should also effectively measure the capabilities it is intended to assess.

Chapter 3:
AI literacy and skills – Students and staff



Context

76. Student AI literacy refers to the knowledge, skills and capabilities students need to be effective, responsible users of AI while they study and in their future professional and personal lives. Being AI-literate means having the practical ability to use AI tools to achieve goals, as well as the ability to evaluate outputs critically, understand the limitations of AI tools and make informed judgements about when and how AI should be used.[30] One framework for higher education breaks AI literacy into four areas:

    a. **Technical.** Understanding how AI systems work and their limitations.

    b. **Evaluative.** Critically evaluating AI outputs.

    c. **Practical.** Effectively applying, integrating, and managing AI tools.

    d. **Ethical.** Understanding the ethical and social implications of AI use.[31]

77. Research shows that levels of AI literacy vary according to student characteristics such as sex, domicile, age and socioeconomic status.[32] Students arrive in higher education with a wide range of digital skills shaped by, among other things, variable access to technology, which can persist or worsen throughout their time in higher education, leading some to warn of a 'digital divide'.[33] As AI skills depend on digital skills, AI is likely to exacerbate this divide.

78. The future impact of AI on the labour market is highly uncertain and the extent to which challenges facing recent graduates can be attributed to AI, rather than other economic and social factors, is widely debated. Current evidence suggests that AI has not yet replaced a significant number of jobs, but it is expected to have a significant impact in the long term.[34] Almost all employers surveyed by the Institute of Student Employers expect entry-level tasks and responsibilities to change in the next three years.[35] Employer skill requirements for entry-level roles appear to be changing, with increasing emphasis placed on AI literacy alongside critical thinking, communication, and adaptability. Requirements are likely to continue to shift in the coming years.[36]

---

[30] See Hackl, Veronika et al, 'The AI literacy heptagon: A structured approach to AI literacy in higher education'.

[31] See Kassorla, Michelle et al, 'Defining AI literacy for higher education'.

[32] See HEPI, 'Student generative artificial intelligence survey 2026'; O'Dea, Xianghan et al, 'Factors affecting university students' generative AI literacy: Evidence and evaluation in the UK and Hong Kong contexts'; Celik, Ismail, 'Exploring the determinants of artificial intelligence (AI) literacy: Digital divide, computational thinking, cognitive absorption'.

[33] See Advance HE, 'Pre-arrival questionnaire national pilot wave 1 initial results'; Digital Inclusion Research Hub, 'Digital divides in higher education are an extension of digital divides in earlier years'.

[34] See Stanford Institute for Economic Policy Research, 'What is really happening to jobs? Separating AI hype from reality'.

[35] See Institute of Student Employers (ISE), 'ISE Student Development Report 2026'.

[36] See ISE, 'ISE Student Development Report 2026'; World Economic Forum, 'AI skills are becoming more important than job experience. Here's how job seekers can adapt'.

79. Research on staff attitudes to AI suggests confidence is uneven, with a growing gap between the most and least confident users.[37] Around half of students say their lecturers or teaching staff help them develop AI skills for their career.[38]

Findings

Students' AI literacy

80. The students we spoke to felt strongly about the importance of AI literacy. Many felt that potential employers would expect them to be able to use AI effectively. Attitudes varied by field: students seeking to go into coding and finance, for example, saw AI as essential; students in other fields expressed more diverse views. Students on creative courses were split between those expecting to use AI as a core part of their work and those entirely opposed to using AI.

81. Some students referred to peers' lack of criticality in assessing the quality of AI outputs. Even some students who opposed using AI in their academic work recognised the importance of understanding how AI tools work. On this basis, student participants frequently argued for more AI training from their institution, especially subject-specific training.

> 'Personally, I'd rather not use it at all, but I understand that I do need to know about it. I do need to understand how it works.'
>
> (Student participant, online roundtable)
>
> 'How do we ensure that students use AI ethically? How do we ensure they feel confident and comfortable using it ethically? Because if we don't, I think we're setting them up to fail. There are a lot of employers now verbalising quite openly that they want their staff to have appropriate skills in the use of AI.'
>
> (Staff participant, online roundtable)

82. Many staff participants also described AI literacy as a key capability that institutions should help their students to develop. Participants said that equipping students for future employment is a core purpose of higher education and adoption of AI tools by employers, including graduate employers, is widespread and expected to increase significantly.

83. Many staff raised concerns that a lack of AI literacy prevented some students from engaging critically with AI outputs and led them to overlook the integration of AI in the tools they use regularly.

[37] See Jisc, 'Staff perceptions of AI 2025'.

[38] See HEPI, 'Student generative artificial intelligence survey 2026'.

Developing student AI literacy

84. Participants described a range of approaches to develop students' AI literacy, including:

    a. **Generalist training.** This is sometimes embedded into study skills training and covers practical skills, responsible use, ethical considerations and sustainability. Training may be delivered through online, self-directed modules or in-person workshops and is generally not compulsory.

    b. **AI capability frameworks.** These help students understand the knowledge and skills required to reach progressively higher levels of AI literacy.

    c. **AI-integrated curricula.** This involves reworking courses in ways that reflect industry practice and embedding responsible AI use into learning activities and assessments in discipline-specific ways.

    d. **Training on subject-specific tools and applications.** This might include, for example, teaching students to use GenAI and machine learning in chemistry, helping students critically evaluate AI-generated outputs against scientific literature in biological sciences, or equipping students to use AI-enabled software in film studies.

    e. **Activities run by students' unions.** These include workshops for students, staff training resources and AI working groups, providing a mechanism for student representation and voice in institutional responses to AI.

> 'We're going through a curriculum transformation process. All of our courses are being redesigned over the next three years. And one element that will be brought in, whether it's a dedicated module or whether it's elements within different modules [...] is around Digital and AI. It's about preparing graduates for the future workplace that they'll be in […] and using AI in a responsible, critical way.'
>
> (Staff participant, online roundtable)

85. Many staff and students said their institution provided limited or no AI training or resources for students. Others described training as uneven and at an early stage of development and rollout.

86. Some participants described drawing on employer and industry input to support curriculum design. Overall, we saw limited evidence of systematic, strategic engagement with employers to identify emerging AI skills needs and reflect them across programme and curriculum development. Significant variation in the level and depth of AI adoption among employers, and uncertainty about future skills requirements, were cited as barriers to preparing students for an AI-enabled workplace.

> ‘We have had no teaching on AI. At all. I think this is a huge part of the problem. Helping students understand how to use AI correctly and efficiently will significantly lower incorrect uses of AI.’
>
> (Student respondent, online survey)

Staff AI literacy and skills

87. Staff participants widely considered improved AI literacy among staff a critical enabler of effective AI integration and the appropriate mitigation of risks. Some suggested a lack of understanding of AI tools among staff had created a fear of inappropriate use, deterring experimentation. As AI use varies significantly between subjects, staff frequently argued that training should be subject-specific.

88. Levels of AI literacy among staff vary significantly. Many staff understand AI and its application to their fields and professional work well, while others do not or only have limited experience of using AI tools. Participants considered generic AI training useful for less experienced staff but of limited value to those with stronger AI literacy, who they felt would be more likely to benefit from discipline-specific training and workshops.

89. As with student training, opportunities for staff to develop their AI literacy vary widely across the sector. Approaches adopted by institutions include:

    a. **Providing AI workshops, training and online modules**, focused on staff literacy or subject-specific skills. These are generally optional.

    b. **Appointing AI ‘champions’ or ‘leads’**, often within programmes or departments, who help other staff develop their capabilities and support structured thinking about how AI affects teaching, learning and assessment.

    c. **Undertaking practice-based approaches**, such as ‘AI playgrounds’, which allow staff to test and experiment with AI tools.

    d. **Joining communities of practice**, including the Jisc AI community.

90. Participants said that training sessions are mostly optional and take-up is often low, with only the most enthusiastic members of staff electing to participate. To increase take-up, some institutions have made training mandatory, but this approach has not been widely adopted because of concern about existing pressures on staff time and workload, the challenges of adapting training to staff with very different base levels of AI literacy, and the possibility of resistance among those opposed to AI on ideological, ethical, environmental or employment grounds.

> ‘Colleagues should model criticality and media literacy – these are the ultimate skills required for being productive citizens in contemporary society.’
>
> (Staff respondent, online survey)

> ‘One of the problems we have is the range of abilities of staff when it comes to AI, because some of them are way ahead of anything we could put on. […] And then there are other people who are right at the beginning.’
>
> (Staff participant, online roundtable)

Opportunities and risks

Opportunities

91. By providing meaningful opportunities to develop AI literacy, institutions can help students to reap the pedagogical benefits of AI tools while also supporting them to navigate AI’s growing role in the wider world.

92. AI literacy is becoming increasingly important to professional practice across many disciplines. There is an opportunity for institutions to help students develop useful AI skills that will equip them to compete in the workplace, including in specialist fields. This will be particularly beneficial for those seeking to enter sectors most exposed to AI disruption. Such skills should supplement, rather than displace, the wide range of other capabilities, such as critical thinking and communication, that employers continue to value highly.

93. Higher levels of AI literacy among staff will allow institutions to take advantage of the opportunities discussed elsewhere in this report, including facilitating the integration of AI into their teaching, designing effective assessments and making efficiencies in their work, as well as delivering more effective AI literacy training for students.

Risks

94. AI literacy is an increasingly important student capability in education, employment and everyday life. Currently, not all students are well equipped to consider critically the ethical implications of AI technology and use it responsibly in their studies and future personal and professional lives. It is important for institutions to develop this capability in students while accounting for differences in students’ prior experience and understanding of AI.

95. Given the speed of AI development, the pace of curriculum development and approval processes at some institutions may be too slow to ensure teaching of AI literacy reflects the latest technology and practice. Institutions will need to balance the flexibility required to keep curricula current with the need to provide prospective students with clear and accurate information about course content.

96. Low levels of AI literacy among staff may make the risks discussed elsewhere in this report more likely to materialise. For example, staff are unlikely to be able to design effective assessments which are not secure unless they have a strong understanding of the capabilities of AI tools. While some staff may be opposed to using AI in the process of teaching their subject, it is important for institutions to ensure staff have sufficient AI literacy and understanding of how AI tools are likely to be used by students to mitigate the risks associated with AI and deliver excellent outcomes for students. As discussed in Chapter 4, low levels of AI literacy also increase risks to data security when working with AI tools and of insufficient scrutiny of AI outputs.

97. There is a further risk that training opportunities are misaligned with staff capabilities. While generic training on staff literacy may be useful for bringing everyone to a common baseline, the significant variations in the implications and applications of AI by discipline mean the most effective training for many staff is likely to be subject-specific.

Case study: AI platform procurement, University of Kent

In March 2026, the University of Kent introduced an institutionally licensed ChatGPT Edu platform. Noting that many staff and students were already using AI tools independently, the university aimed to reduce the associated risks by providing a secure and supported environment for AI use.

Before accessing the platform, students must complete training modules on AI literacy, responsible use, data protection, prompting techniques and academic integrity. The training is intended to develop their AI literacy and workplace skills.

Use of the platform is optional, but students who choose not to sign up are still encouraged to understand the technology and its implications for their discipline. The training does not require the use of AI tools, ensuring it is inclusive of students and staff who may prefer not to use them. AI content has also been incorporated into existing data protection training for staff.

Chapter 4:
Strategy and operations



Context

98. Research on AI uses in higher education has focused primarily on teaching, learning and assessment practice, with comparatively little attention paid to strategy and operations.[39] Some research has argued that, by accelerating the processing and analysis of data, automating repetitive tasks and testing possible scenarios, AI tools can drive operational efficiency and improve strategic planning.[40] AI has also stimulated broader strategic conversations about the purpose and priorities of higher education, including its role in an AI-enabled society, the relative importance of knowledge, skills and human development, and how institutional culture, governance and values should shape decisions about approaches to AI. These important issues are beyond the scope of this report.

99. Evidence suggests the potential for significant efficiency gains through AI adoption. Research by the Boston Consulting Group suggests that across the global economy, 50 per cent of companies currently generate value from AI initiatives, up from 5 per cent last year.[41] Systemic adoption in the higher education sector appears to be in its early stages globally. A UNESCO survey in 2025 found that 19 per cent of responding institutions had a formal AI policy and a further 42 per cent have AI guidance in development.[42] Jisc's research suggests that in the UK, institutions are much more likely to use AI in teaching and learning than in operational processes such as planning or reporting.[43]

100. To help institutions manage the implementation of AI technologies in the context of existing digital infrastructures, Jisc has published a tool intended to help them assess their digital maturity.[44]

Findings

Strategy

101. At most institutions we spoke to, approaches to AI have been shaped primarily by innovation and experimentation within departments and by individual members of staff, rather than by a coordinated institution-wide strategy. Some participants raised concerns that this had led to missed opportunities to leverage the benefits of AI, variation in students' experiences and a lack of coordination across education, operations and research.

[39] See Zawacki-Richter, Olaf et al, 'Systematic review of research on artificial intelligence applications in higher education: Where are the educators?'; Jiang, Xinyi, and Abdullah, Zuraidah, 'AI-enabled governance in higher education: A systematic review of applications, outcomes, and emerging implications'.

[40] See HEPI, 'AI and the future of universities', Chapter 3.

[41] See Boston Consulting Group, 'Are you generating value from AI? The widening gap'.

[42] See UNESCO, 'UNESCO survey: Two-thirds of higher education institutions have or are developing guidance on AI use'.

[43] See Jisc, 'AI and Jisc's Leadership Survey 2025'.

[44] See Jisc, 'Launching our digital maturity assessment to support transformation across UK higher education'.

102. Some institutions have established dedicated governance and leadership structures to coordinate their AI response, including steering groups and taskforces, working groups, and cross-functional networks. Some institutions have launched transformation programmes to embed AI in course curricula, assessment and institutional operations and to train staff. Others have hired dedicated AI leads or champions, or appointed members of senior leadership teams to coordinate their institution's AI response.

103. The extent to which these structures and roles operate as part of a coherent whole appears to vary. Participants highlighted considerable differences in how arrangements were implemented and embedded across institutions, with some reporting difficulties maintaining engagement or operating effectively in practice. Many highlighted gaps between institutional policies and implementation by staff. Some staff reported challenges producing guidance that is specific enough to be useful, yet general enough to adapt to local requirements and remain relevant as AI tools develop.

104. Participants were asked whether their institution has a strategic approach to AI that considers the implications of AI for its institutional mission, the educational experience it intends to provide and the capabilities and attributes expected of its graduates. Only a small number of participants said their institution had an AI strategy or that AI was embedded in its strategic approach, and few described their approaches being shaped by them, suggesting strategic approaches to AI have not been widely adopted or embedded throughout the sector.

> 'We have produced a lot of staff training and guidance but recognise that there is still much room for improvement here. It is very difficult to get a handle on the situation, some staff are very keen and confident with its use where others are more reluctant.'
>
> (Staff respondent, online survey)
>
> 'If you let a thousand AI flowers bloom without any quality control then things are going to get messy. Institutions need to have a process whereby projects are registered, checked, selected for scale and supported as they scale.'
>
> (Staff respondent, online survey)

Efficiency and operations

105. We heard about a range of AI uses intended to make processes faster and more efficient. Staff told us they and their colleagues use AI to create marketing and operational content, input data and review curriculum materials at scale. Some said that staff morale had benefited from the outsourcing of dull or repetitive tasks.

106. However, most staff participants felt the implementation of AI tools had not yet led to significant efficiency gains in their work. They said the high volume of errors currently made by AI tools necessitates extensive human oversight, meaning little time is saved. Many felt their institutions had not yet properly explored the opportunities for efficiencies, because of inadequate staff capacity and the absence of a coordinated, institution-wide approach to identifying and prioritising initiatives. Staff also pointed to cultural challenges in finding AI-

enabled efficiencies, observing that financial pressures had heightened staff sensitivities to changes perceived as putting jobs at risk.

107. Staff reported several other barriers to the effective adoption of AI. These included immature underlying data infrastructures, which inhibit effective, organisation-wide information sharing, and inefficient processes for approving and implementing new technologies. Long institutional policy and strategy development cycles were cited as barriers to timely adoption.

108. Of the approaches to AI integration we discussed, AI-enabled chatbots were among the most visible to research participants. These are intended to support students by quickly triaging them to relevant information. Students had mixed views about the value of chatbots: while some considered them a useful means of accessing information quickly, others found them annoying or unhelpful. Concerns related to the effectiveness of the technology rather than its use in principle.

109. Participants reported risks to data security and governance, including concerns that some AI services process or store data outside the UK, and that content entered into some consumer AI tools may be used to improve or train future models. Some participants reported poor staff understanding leading to sensitive data being entered into free chatbots. Others said that staff were sometimes not sufficiently critical of AI outputs, leading to errors getting through. Participants also raised concern about the integration of AI tools that are currently available for free or at low cost, given that future price increases would expose institutions to substantial additional costs.

> ‘In 2025-26, we have run a number of pilots testing AI use cases. These include for enquiry management and academic administration processes. However, these have reinforced that AI is expensive to deploy and requires considerable human oversight.’
>
> (Staff respondent, online survey)
>
> ‘There are massive opportunities around automation of routine administrative tasks: timetabling queries, drafting communications, summarising committee papers, processing feedback. […] Research data and unpublished results are commercially and academically sensitive; leakage into external models is a serious risk.’
>
> (Staff respondent, online survey)

Opportunities and risks

Opportunities

110. There is an opportunity for institutions to use AI strategies to set guiding principles and drive institutional culture on engagement with AI. This may include providing active and sustained support for staff to navigate changes to professional practice, and involving staff and students in developing, testing and reviewing institutional approaches.

111. By developing and implementing strategic approaches to AI, senior leaders have the opportunity to shape institutional culture by setting out key principles for staff engagement with AI. Training and clear guidance from senior leaders may help staff navigate changes to professional practice, academic judgement and established ways of teaching and creating knowledge. AI strategy development can be strengthened by involving staff and students in developing, testing and reviewing institutional approaches.

112. AI tools present an opportunity for greater efficiency in the student interest. In a resource-constrained environment, automating administrative and routine work may create opportunities for academic and support staff to spend more time with students and undertake other high-value activities that make full use of their specialist knowledge and skills.

Risks

113. To date, AI uses aimed at improving institutional operations have primarily been driven by individual staff, teams and departments. Greater central oversight and coordination is likely to be necessary for institutions to realise efficiency gains from AI tools at scale, though this should not come at the expense of innovative individual practice. This will frequently depend on having a mature digital infrastructure that facilitates effective information sharing and data management.

114. Many of the challenges identified in this report, such as digital transformation, staff training and organisational efficiency, require a strategic, senior-level response. Flexible, principles-based approaches, whether through a dedicated AI strategy or integrating AI into wider institutional strategies, may help departments adapt approaches to their local context while remaining responsive to technological change.

115. While streamlining processes is important to ensure institutions can respond rapidly to developments in AI, this may sometimes be in tension with the need for effective guardrails that protect data security and privacy. Careful consideration of staff requirements is needed to ensure these responses put a burden on staff time in proportion to the risks they are intended to mitigate.

116. The complexity of institutions' integration of digital services may make it difficult to change supplier, driving so-called 'vendor lock-in'. As a result, a potential increase in the price of providing AI tools poses a significant risk to sector finances. To help address this issue, Jisc is developing a new digital assessment feedback framework to support the procurement of assessment solutions.[45]

[45] See Jisc, 'Why higher education needs a new approach to digital assessment'.

Case study: AI oversight at Anglia Ruskin

Anglia Ruskin University has established an AI oversight group to coordinate the university's response to AI. The group is chaired by the deputy vice chancellor for education and includes senior staff with responsibilities for research, strategic planning and digital, as well as external members with expertise in AI and higher education.

Taking a responsible AI adoption approach, the institution has implemented mandatory training for staff on key issues such as academic integrity, alongside support to use AI in their teaching. This is differentiated by staff confidence with AI, with introductory training given to less confident or more sceptical staff and further training given to enthusiastic adopters to ensure consistency of practice across the institution.

The university has also worked to create a safe environment for debate on the topic by holding forums for discussion, funding initiatives where there is the opportunity to make improvements in practice, and working to accommodate the ideas and concerns of staff.

Chapter 5: Regulatory and policy environment



Context

117. Prior to undertaking this research, the OfS engaged with the sector on AI through listening sessions and a student debrief. Informed by this engagement, the OfS set out its initial position on AI in June 2025.[46]

118. Advance HE has supported the higher education sector in building a shared understanding of generative AI by delivering bespoke AI-related consultancy, member benefit projects and an annual AI symposium. It also funds collaborative development projects aimed at promoting innovative, ethical and inclusive use of AI.

119. A number of other UK sector bodies have played an important role in supporting institutions to develop their approaches to AI, including the following:

    a. The Russell Group's principles on the use of generative AI in education, which set out how its institutions would respond to AI, were among the first interventions in the sector on this issue.[47]

    b. Jisc has published a range of resources, including a digital capability framework, maturity toolkit and strategic framework for AI adoption in colleges.[48] Its AI community, including groups focused on teaching and learning, research and professional services, facilitates the sharing of knowledge and best practice among educators. It has also published research on student and staff attitudes to AI.[49]

    c. The Quality Assurance Agency for Higher Education (QAA) has published guidance with a particular focus on maintaining quality and standards.[50] It has published research exploring institutions' approaches to assessment in the context of AI, and recently announced a programme of work focused on assuring the standard of awards in the context of AI.[51]

The impact of the regulatory environment

120. Participants were divided on the impact of the regulatory environment on their work:

    a. Some participants felt the OfS had facilitated effective practice by setting out the principles within which institutions could experiment and innovate.

    b. Some participants felt regulators had not provided clear or sufficiently specific guidance about how institutions should respond to AI.

[46] See OfS, 'Embracing innovation in higher education: Our approach to artificial intelligence'.

[47] See the Russell Group,'Principles on the use of generative AI tools in education'.

[48] See Jisc, 'Artificial intelligence'.

[49] See Jisc, 'Student perceptions of AI 2025' and 'Staff perceptions of AI 2025'.

[50] See QAA, 'QAA advice and resources on Generative AI'.

[51] See QAA, 'The perfect storm: AI, assessment and a sector under pressure' and 'Assuring the standard of UK awards in the age of GenAI'.

c. Some participants said the regulatory environment had not significantly shaped their approach. Instead, their institution's approach to AI was driven by institutional philosophy and subject area need.

121. Others described the regulatory environment constraining their ability to respond effectively to the risks and opportunities presented by AI:

    a. Some participants felt a lack of clear direction from the OfS and other sector bodies may have contributed to an uncoordinated response.

    b. Some participants felt the requirements imposed by professional, statutory and regulatory bodies, which can include specifying assessment types, may have restricted the ability of some institutions to adapt their curricula and assessment models on some accredited courses. The need to ensure compliance with such requirements may also slow down the process of changing and updating courses. However, the high standards set by some of these bodies may also have helped ensure the efficacy of assessment methods.

122. Staff participants raised concerns that consumer protection law prevents providers from changing their method of assessment in response to the use of AI by students. However, this is not explicitly prohibited in the Competition and Markets Authority's guidance on consumer protection law.[52] While providers cannot rely on broad, unrestricted contract terms that allow them to alter significant aspects of a course, they are nevertheless likely to be able to make changes where fair and transparent terms explain when and why these may be necessary. This may include for example, making changes that are necessary and proportionate to ensure courses are up-to-date, or protect the integrity and value of qualifications to students or to meet other regulatory conditions. Where a provider has committed to assessing a module in a particular way prior to a student signing their contract, the provider may need to consult with students on any changes. Giving students the option to switch to another suitable module may also help ensure this is fair for students. Whether or not the provider has made a specific commitment to students, any changes should be accompanied by appropriate support for students who may be negatively affected as a result of the change, including the provision of reasonable adjustments and a route for students to raise concerns. If no equivalent courses are available or adjustments are not possible, it may be necessary to consider allowing affected students to cancel their contracts without penalty. However, providers should take their own legal advice on how these requirements apply in their circumstances.

Suggestions for regulatory and other sector bodies

123. Many participants called on the OfS, Advance HE and other bodies to adopt a collaborative approach in response to AI. There was agreement that the sector response to AI would involve close working between sector bodies and collaboration between institutions. Participants emphasised the value of sharing best practice in relation to shared challenges, particularly in areas such as data governance and security, academic integrity, and the responsible, ethical use of AI in professional contexts to help students prepare for the

[52] See CMA, 'UK higher education providers – advice on consumer protection law'.

workplace. Because of the rapid pace of AI development, participants stressed the need for collaborative work to be ongoing and iterative.

124. Participants generally said that further regulation from the OfS, such as new conditions of registration, would be unhelpful. There was a concern that further regulation would stifle innovation, reduce the space to make mistakes, and further slow institutions' responses. A small number of participants disagreed and called on the OfS to introduce regulation restricting AI use or slowing institutions' adoption of AI tools.

125. Participants argued widely for more regulatory clarity from the OfS, particularly its B conditions (which regulate quality) and E conditions (which regulate governance), so they could be confident they are compliant with regulatory requirements. Some suggested the OfS should develop worked examples setting out how its regulation would apply in particular contexts. Some participants also called on the OfS and Advance HE to be transparent about their own use of AI and consider ways of using AI to become more efficient themselves.

Conclusion

126. When used effectively, AI can support learning. It may have particular benefits for some student groups in reducing barriers to learning and improving their wider student experience. There are many examples of innovation across the sector, but there is headroom for institutions to go further in harnessing the opportunities AI presents to make improvements to teaching practice.

127. Using AI uncritically in place of important learning develops neither students' AI skills nor the other capabilities they will need to succeed beyond higher education. Institutions should therefore make deliberate choices about where and how students are expected to use AI.

128. AI poses significant risks to student learning and assessment integrity, including the risk of awarding marks that do not appropriately reflect students' knowledge and skills. Institutions should develop accessible and robust approaches to assessment that effectively assure student achievement, while working to create opportunities for authentic assessment and the responsible integration of AI.

129. In designing their approaches to assessments, institutions should distinguish unauthorised AI assistance from the legitimate use of assistive technologies and reasonable adjustments for disabled students.[53] Tools provided through reasonable adjustments processes should effectively address the assessed disadvantage without providing a further advantage.

130. It is increasingly important for students to have AI skills for their future employment and participation in modern society. Institutions should provide structured opportunities for students to develop the knowledge and judgement needed to use it well. AI literacy should not be assumed to arise solely through the personal use of AI tools.

131. To date, innovation and experimentation with AI has generally been led by individual departments and academic staff. Institutions should now demonstrate leadership on AI by adopting a strategic approach that recognises the technology's complex and wide-ranging implications across teaching, professional services and beyond. This requires visible strategic leadership that empowers staff to make decisions about the technology's use in the interests of students, alongside driving the meaningful implementation of any AI strategy and ensuring staff are AI-literate.

[53] See Equality and Human Rights Commission, 'Advice note for the higher education sector from the legal case of University of Bristol vs Abrahart'.

Effective institutional practice

Study and teaching

132. Effective institutional practice in study and teaching is likely to require:

a. Considering carefully whether and how the integration of AI tools across different courses can support teaching and learning.

b. Assuring the quality of AI-supported teaching materials and ensuring that AI-supported feedback effectively supports reflection and improvement.

c. Ensuring staff use of AI in teaching and feedback is clearly communicated to students.

Assessment

133. Effective institutional practice in assessment is likely to require:

a. Adopting secure assessments which are designed to prevent students from using AI or, where non-secure approaches are used, ensuring they designed in a way that makes them valid and effective measures of students' knowledge and skills.

b. Ensuring expectations for permitted AI use are clearly and effectively communicated to students, including whether alternative options are available for students who do not wish to use AI.

c. Ensuring academic misconduct processes are fair and reliable, and do not depend excessively on the outputs of AI detection tools.

AI literacy and skills

134. Effective institutional practice in AI literacy and skills is likely to require:

a. Developing structured programmes to develop students' AI literacy, including the practical and critical capabilities they need to use AI effectively and responsibly, taking into account variation in students' AI literacy on arriving in higher education.

b. Being clear about expectations for AI use on a particular course, as far as possible given the rapid pace of AI development, so that prospective students can make an informed choice about whether the course is right for them.

c. Mitigating risks that differential access to and experiences of AI tools and training exacerbate disparities in students' experiences and outcomes.

d. Ensuring staff have the knowledge and skills needed to respond effectively to AI, and addressing significant gaps through appropriate training and support.

Strategy and operations

135. Effective institutional practice in relation to strategy and operations is likely to require:

    a. Establishing effective mechanisms to identify and address weaknesses in institutional capabilities, systems and processes that could create vulnerabilities or limit the institution's ability to realise the benefits of AI.

    b. Exploring opportunities for efficiencies systematically while encouraging innovations in individual staff practice.

    c. Adopting a strategic approach to AI, either through a dedicated AI strategy or by embedding AI considerations in relevant institutional strategies. This may include ensuring AI is discussed regularly at executive and board level, setting out key principles for staff engagement with AI and providing support to help staff navigate changes in institutional approaches.

136. Finally, the rapid evolution of AI will continue to create new opportunities and challenges for institutions to navigate. Advances in AI capabilities, including agentic AI, new applications and changing approaches to the governance and detection of AI tools, are likely to reshape higher education in ways that may be difficult to anticipate, creating a complex and rapidly changing landscape. The OfS and Advance HE will continue to engage collaboratively with the sector, providing further support and updates as the technology evolves.

Acknowledgements

We are grateful to everyone who contributed to the research for this report for their time and valuable insights:

- 91 students and 244 staff who participated through our live events and online prompts.
- The following sector experts for their support with the research design and drafting of the final report: Professor Chie Adachi (Queen Mary University of London), Michal Bobula (University of the West of England), Professor Parama Chaudhury (UCL), Professor Michael Grove (University of Birmingham), Professor Rose Luckin (UCL Institute of Education), Phil Maull and Rebecca Robinson (QAA), Professor Alan Spivey (Imperial), Michael Webb (Jisc) and Trevor Woods (Proudly Human).
- Staff from Imperial, the University of Bath, the University of Kent and Anglia Ruskin University for their support developing the case studies.
- Staff from the following sector bodies and mission groups: GuildHE, Independent Higher Education, London Higher, the Russell Group, Universities UK and University Alliance.
- Staff from the following technology companies: Ellucian, Faculty AI, Kortext, Studiosity and Turnitin.
- Staff from other organisations: the General Medical Council, the Institute of Student Employers and the Nursing and Midwifery Council.

Annex A: Participating institutions by tariff

This table shows the institutions represented in our research by tariff, according to the OfS typology of higher education providers.[54]

Tariff	Number of institutions
High tariff	23
Medium tariff	26
Low or unknown tariff	32
Levels 4 and 5	10
Specialist	15
Postgraduate	3
Other	18
Total	**127**

[54] See OfS, 'Provider typologies 2022: Methodology for grouping OfS-registered providers'.

Office for Students

In partnership with

Advance HE

© 2026 Office for Students and Advance HE. All rights reserved.

This publication is available under the Open Government Licence 3.0 except where it indicates that the copyright for images or text is owned elsewhere.

www.nationalarchives.gov.uk/doc/open-government-licence/version/3/