



Generative Artificial Intelligence in Education - QAA's response to DfE call for evidence

This is QAA's response to the Department of Education's (DfE's) consultation on generative artificial intelligence in education. The consultation sought views on how generative artificial intelligence is being used across education in England and the opportunities and risks it presents.

Experience with generative AI

10. Have you or your institution used generative AI in an educational setting? If so, could you briefly describe the ways it was used and the specific tools used.

QAA does not deliver education directly, though we do have experience of using generative AI. Our evidence draws on the deep expertise in the higher education (HE) sector that we have collated in formulating our guidance, and on our wider discussions with HE providers and others.

Opportunities and benefits of AI

13. How do you think generative AI could be used to improve education?

Generative AI can be used to improve education in a number of ways:

- Integrated into curriculum and assessment to deepen learning, and build cohorts of AI confident and literate learners to meet the needs of a rapidly evolving labour market.
- Used by staff to speed up processes and drive efficiencies where appropriate.

AI literacy is already a desirable skill in the labour market. Generative AI tools can be integrated into the curriculum and in assessment, both as tools to deepen learning and to build up AI literacy that will be needed in the labour market after graduation.

It is crucial that the education system provides consistent training and a consistent approach to AI literacy throughout a student's academic journey - from primary, through secondary, to higher education. While generative AI tools can be used to deepen learning and improve education, their effectiveness depends on the quality of staff support in teaching the best use of these tools. Education institutions must therefore train their staff to use generative AI effectively.

One use for generative AI in education is to support students in reflecting on their academic skills. For instance, students from non-traditional backgrounds or those for whom English is not a first language can use generative AI tools like ChatGPT, and other large language models (LLMs), to improve the style and tone of their writing. Then, they can write a critical reflection on what they learned about their academic tone and style.

Similarly, with adequate staff training, generative AI tools have the potential to improve the quality of provision and student experience by making it easier and quicker for lecturers to easily personalise their feedback on assessments, enhancing a student's ability to learn how to improve academically.

The rise of generative AI also provides a further incentive to re-imagine the purpose and methods of assessment in education. Authentic assessment - that is, assessment design which properly prepares students for life after graduation - is an important part of developing highly skilled graduates who go on to contribute to society economically, culturally and socially. With generative AI posing challenges to existing assessment methods, this is an opportunity to interrogate why and how we assess students to ensure that they are adequately prepared for the labour market. [QAA's recent guidance](#) explores this topic further.

Crucially, it is only a matter of months before generative AI will be fully integrated into software regularly used within educational settings, such as Microsoft Word. Such software is sometimes provided for students by the institution itself. The need to rethink the way we use AI in education will be reinforced as the technology becomes deeply embedded in the tools that students already use to generate outputs on which their learning is assessed.

Concerns and risks of AI

15. What are your main concerns about using generative AI in educational settings?

- Academic integrity
- Potential grade inflation
- Automation/depersonalisation of feedback in assessment
- Digital divide
- Discrepancy of approach leading to differentials in graduate skills

There are a number of concerns associated with using generative AI in education settings. The first is academic integrity. The rigour and value of education qualifications relies on being able to trust that a student's work is their own and truly representative of their ability. With open access to generative AI tools, it is increasingly difficult to have that trust. This is compounded by the fact that AI detection tools are unreliable at best, and can also incorrectly flag students' original work as AI generated, especially if they are a non-native English speaker. AI detection software like that developed by Turnitin has seen significant pushback in the UK higher education sector. It is currently being used inconsistently both within institutions and across the sector as a whole. The software provides an estimate of the percentage of an assignment that is AI generated. Some institutions use the AI generated score as a blunt threshold to refer someone into an academic misconduct policy. Some use it as one of many data points to triangulate a judgement. Some have chosen not to use it at all.

Detecting AI-generated content is challenging, especially with generative AI already being fully integrated into Microsoft Word through its new CoPilot product. Therefore, it's increasingly important to actively work with students to help them understand good academic practices and what's expected of them, to preserve academic integrity and maintain the trust and value of degrees.

One of the key issues here is assessment strategies within educational settings, and the extent to which they genuinely test a student's ability. AI provides a much-needed catalyst to

re-think how and why we assess within education. It is important that this takes place across the whole educational lifecycle, from primary all the way up to tertiary education, in order to support students' transition between each stage of education and prevent disparities embedding early on in the system.

Over time, it is likely that human/AI writing will be considered the norm. Even Microsoft Word currently attempts to predict the next word an individual types. If we are moving towards a world where human/AI hybrid writing is the norm, then the educational sector will need to reconsider what it means by 'plagiarism' and associated academic misconduct. The work of researcher Sarah Eaton in her 2021 book 'Plagiarism in Higher Education: Tackling tough topics in academic integrity' is valuable in this regard, outlining factors of a 'post plagiarism era' and what it might mean for education. In the higher education sector, we would warn against any knee-jerk reactions to the re-design of assessments that see a return to in-person, invigilated exams as the predominant form of assessment, as it has been shown to be a poor reflection of student ability, inaccessible for many with disabilities and additional needs, and not adequately preparing students for the workplace.

Secondly, there is a concern about the effect of generative AI tools on grade inflation, particularly within higher education. If generative AI - even the legitimate, authorised use of it within assessments - leads to students being more capable of achieving higher grades, we could see a rise in grade inflation over the next few years, which could compromise the integrity and value of awards. If human/AI hybrid writing is going to be the norm in future, especially in the workplace which educational settings should be preparing students for, then it might be necessary to re-think or re-establish the baseline level of achievements students are capable of in order to maintain the integrity of the grade classification system.

Thirdly, while generative AI presents opportunities for staff to automate and speed up processes that currently use up time and resource, it is possible it could lead to the automation and depersonalisation of assessment feedback. This could lead to a reduction in the quality of feedback and impact student learning and achievement. Educational institutions will need to decide the extent to which they use generative AI for processes like this and make those expectations clear to students.

Fourthly, the existence of generative AI tools could deepen the existing digital divide that was exposed so clearly in the pandemic. Already, the more sophisticated generative AI tools are behind a pay wall. This digital divide could deepen discrepancies in learning outcomes and achievements. Therefore, educational institutions must ensure that all students have equal access to generative AI tools that might be permitted for particular assessments. (The digital divide also applies to improper use: tools that are more likely to evade detection software are more expensive, and contract cheating outfits (like essay mills) are already charging a premium for writing assessments that can evade detection tools.)

Finally, it is important that there is some level of consistency of approach to integrating generative AI in education settings because discrepancies could lead to differentials in graduate skills. If some educational institutions do not adequately support and educate students on how to use AI tools ethically, critically and appropriately, their graduates may lack the skills required to work in an environment where generative AI is prevalent. This will make them less competitive as a potential employee than a student who has been taught critical AI literacy and can use AI effectively in a job to delivery higher productivity. Therefore, it is important to ensure that all educational institutions adopt a consistent approach to using generative AI in education to ensure that all graduates have the necessary skills to succeed in the future job market, contributing to the economy and wider society. Institutional autonomy allows institutions to take different approaches, but institutions, faculties, courses and modules will have to be incredibly explicit about what constitutes 'unauthorised use' at each stage so as not to penalise students unnecessarily.

16. If at all, have these concerns impacted your use of generative AI? Please explain how.

These concerns have influenced the focus of our work as a quality agency to support the higher education sector to respond to the rise of generative AI. As of August 2023, QAA has released three pieces of advice and guidance to the sector: 'How to approach ChatGPT'; 'Maintaining quality and standards in the ChatGPT era: QAA advice on the opportunities and challenges posed by Generative Artificial Intelligence'; and 'Reconsidering assessment for the ChatGPT era: QAA advice on developing sustainable assessment strategies'.

The guidance provides advice on how to manage the rapidly increasing use of generative AI tools in higher education, while also maintaining academic standards. It covers topics such as:

- The potential benefits and risks of using generative AI tools
- How to support academic integrity in a world of generative AI tools
- How to design assessments that are fair, authentic and accessible and reliably measure student achievement.

The guidance also includes a number of practical tips for providers, such as:

- Be clear about your institution's policy on the use of generative AI tools
- Stay agile and responsive to new developments in software
- Train staff in critical AI literacy
- Approach plagiarism detection software with caution
- Partner with students in approaching generative AI integration and use.

In addition to the guidance, the QAA has also hosted a number of webinars and events on generative AI for the HE sector. These events have provided opportunities for providers to share best practices and learn from each other. The webinars, and [all our work on ChatGPT](#) can be found on our website.

17. Are there specific subjects or areas of education where you believe generative AI should not be used? Why?

It is difficult to specify certain areas of education where AI should not be used, simply because the available tools are so wide-ranging across subject areas and skill, and their outputs are incredibly difficult to detect as AI generated.

That said, clearly there are subjects where the stakes are particularly high when it comes to trusting that a student has acquired the knowledge and skills to complete their course and progress in their profession without immediate recourse to generative AI tools. Medicine, Dentistry and allied health professions are an example of this – it is vital that graduates have the requisite knowledge for their clinical practice. It is still possible to integrate generative AI tools into their learning to aid it, and indeed some medical professions will increasingly have AI-supported tools in the workplace which students will need to know how to use critically and well, but it's important they are able to meet the threshold standards of skill and knowledge on their own merit alone. This also applies to academic courses that lead to professional registration, where competency to enter a profession is assessed through the course.

Ethical and legal considerations

18. If any, what are your views regarding ethics, data privacy and security when using generative AI in education?

Generative AI tools are trained on datasets that use people's information without permissions or accreditation, and that have in-built biases in them that exacerbate existing knowledge bias in academic disciplines. For example, ChatGPT will list 10 male philosophers if asked to list the most influential philosophers. Similarly, photo generating tools use the likeness of photos of real people on the internet without their consent, and the artwork of artists without proper credit. In order to properly train educational staff on critical AI literacy to support students, it will be necessary to upskill them on issues relating to bias, intellectual property rights and GDPR, given how tools like LLMs accumulate their information and reflect in-built biases.

Future predictions and enabling use

19. How do you see the role of generative AI in education evolving in the future?

Over the long term, AI will have transformational effects on education. Firstly, the mass integration of AI tools into software will make it exceedingly difficult to 'regulate' use in educational settings. As aforementioned, if human/AI writing becomes the norm, not least because it will be commonplace in the workplace and education will need to prepare students for that, we need to consider what that means for how we understand plagiarism and academic misconduct, in order to maintain the integrity of educational qualifications and their value in society and its economy. What seem like progressive approaches today are likely to seem quaint or even foolish in a matter of months. For example, asking students to declare their use of AI by submitting their LLM transcript with their assessment to show how they used it critically in completing an assignment, might shortly feel as futile as asking a student to declare the use of google in an assignment.

There is a need for a mass strategic shift within educational institutions to educate and train staff in order to support students effectively to get to grips with AI and how they can (and can't) use it in their learning. This should cover initial teacher training, and continued professional development for existing educational professionals.

20. What support do education staff, pupils, parents or other stakeholders need to be able to benefit from this technology?

Stakeholders need to know what tools are available, their formats (for example, plugins vs standalone web tools, vs integration into software like MS Word), a basic understanding of how they work in order to understand their risks and weaknesses, and how they can be used to deepen learning and be used as a force for good.

21. What activities would you like to see the Department for Education undertaking to support generative AI tools being used safely and effectively in education?

DfE can utilise the expertise within the educational sector to support its own recommendations and response. QAA has been leading the response within the higher education sector, with our work referenced by bodies like UUK, the Russell Group and international bodies like the Global Academic Integrity Network (with representation from quality and regulatory agencies globally), and the European Association of Quality Assurance in Higher Education. QAA would be delighted to support DfE more closely on its work around AI, particularly within higher education, including through contributing our expertise to its AI taskforce, or in regular engagement with relevant officials.

It might be useful for DfE to play a role in investigating interventions to support staff upskilling by providers (in all sectors) so they can effectively support learners as an investment in the UK workforce. This is a significant challenge that will take a consistent, concerted effort over a number of years. Thus, it would benefit from DfE's leadership.

QAA believes there is need for supporting a more consistent and standardised approach to the integration of AI in higher education, which could feasibly be done by DfE with support from organisations like QAA, Jisc and AdvanceHE. We are working with those sister agencies on this but would welcome greater coordination and alignment within DfE to support a consistent approach to avoid discrepancies arising across the sector. QAA also engages regularly with Professional, Statutory and Regulatory Bodies (PSRBs) in the UK which need to be key players in these discussions and we would be happy to facilitate engagement with them in this area. This could look to bring together bodies from across secondary and tertiary, such as Ofqual regulated bodies and IfATE to enable conversation around consistent and standardised approaches to generative AI and assessment to help manage expectations and student transitions.

No doubt DfE is also in regular conversation with other government departments to ensure a cohesive approach, as HE/Education is one of many sectors wrestling with the questions around generative AI and the upholding of standards and integrity.

22. Is there anything else you would like to add on the topic of generative AI in education?

As mentioned, QAA would be delighted to work more closely with DfE on a long-term basis in tackling the issue of AI in education, particularly given our wealth of expertise in this space and our leading response within the higher education sector that is supporting institutions on this journey across the UK.

Published - 7 September 2023

© The Quality Assurance Agency for Higher Education 2023
Registered charity numbers 1062746 and SC037786
www.qaa.ac.uk