# London Higher response to the DfE's generative artificial intelligence in education call for evidence

***This submission has been compiled by Darren de Souza, Senior Policy and Projects Officer, London Higher.***

London Higher is the representative body for 50 universities and higher education colleges in the capital, including all 5 London Russell Group universities, 10 world-leading creative arts institutions, 10 small specialist institutions designated world leading by the Office for Students, 5 of the top 10 ranking universities in the REF 2021, and 20 universities teaching healthcare and medicine helping to staff these crucial sectors in London and across the whole of the UK.

London Higher welcomes the opportunity to provide feedback to the Department for Education (DfE) on the subject of generative artificial intelligence (AI) in education. London Higher has responded to each question as set out in this call for evidence, and this submission draws from responses across the London Higher membership and its networks.

If you would like to discuss this submission further or require any further information, please contact: Darren de Souza, Senior Policy and Projects Officer at London Higher at darren.desouza@londonhigher.ac.uk .

**(Theme 1 - Question 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.**

Across responding London Higher members, use of generative AI is varied. For some, it is still in its infancy, and tools such as ChatGPT are being used to familiarise students and staff with the technology. By upskilling and improving their knowledge of available tools, staff and students are consequently more aware of the challenges surrounding the use of AI (discussed further in questions 11 and 15), demonstrating how institutions are promoting ethical and responsible use of a rapidly developing technology.

In other institutions, generative AI has been used by educators and learners in various fields, to assist with tasks such as:

- assessments for generative music and computational arts subjects;
- creating subject handouts, meeting notes, quizzes or strategy documents;
- explaining concepts, drafting essay structures, improving grammar and readability;
- coding blockchain with Python or other programming languages; and
- patient and environmental simulations for under-/postgraduate health and care courses.

Several tools have been mentioned, most frequently OpenAI's ChatGPT. PCS Spark has been used to create health/care simulation environments. Whilst recent discourse has been dominated by ChatGPT, members are aware that other tools with similar capabilities are available, such as Anthropic's Claude 2, Google's Bard and Microsoft's Bing Chat.

**(Theme 1 - Question 11) What were the main challenges you faced in using generative AI and how did you address these?**

The following challenges were identified:

- establishing guidelines for use for acceptable/legitimate use in students' work and assignments;
- internal discussions around banning generative AI and reliable detection;
- designing assessments that allow for generative AI use in a way that assesses a student's critical thinking and engagement with the technology;
- ethical and copyright issues regarding authorship;
- information verification, for example where AI generates academic references that do not exist;
- paywalls and lack of direct access to code and models in more advanced technical projects; and
- understanding of specific tools, their capabilities and limitations, and engineering prompts to produce useful answers.

Distinctions have also been made between free/paid versions of tools, for example ChatGPT 3 (free) and ChatGPT 4 (paid), the latter of which is viewed as more reliable and less likely to 'hallucinate' (create outputs that do not reflect real world data, such as references). There is concern that this could lead to a digital divide, with those who are able to pay subscription costs being placed at an advantage.

It is too early to say whether the above challenges have been satisfactorily addressed, however many institutions have created staff steering groups and forums on AI applications. This has enabled staff to discuss the use of AI in the curriculum and assessments – as well as consider how assessments can identify improper use. Moreover, these forums allow institutions to capture what current student and staff understanding of generative AI is, and what priority areas need to be addressed.

**(Theme 1 - Question 12) What was the result of your use of these tools, including any impacts?**

There is broad consensus that it is too early to definitively convey the impacts of generative AI tool usage, but early positive outcomes have been identified. These include:

- saving staff time on routine tasks (such as marking rubrics or quiz creation);
- increasing student engagement with AI and consideration of ethical questions and real-world applications; and
- opening dialogue within institutions on best practice and next steps.

Our institutions have noted early adoption and clear guidelines have also been appreciated by their students.

Members agree that there is significant potential for AI in education, and integration of tools into curricula and assessments that can mitigate misuse through proper management. It is clear that the use of generative AI brings both improvements and risks to higher education institutions, which are detailed further under the relevant themes.

**(Theme 2 - Question 13) How do you think generative AI could be used to improve education?**

There is broad agreement that generative AI could facilitate a more personalised learning experience for students, through generated materials that are able to account for individual learning styles, using different techniques to convey key concepts and subject areas, and facilitate further research.

A number of institutions noted that AI tools can be used by students (including neurodiverse students and alongside assistive technology) to help with spelling, grammar, and structuring/planning written work. Integrating the use of AI in study support and teaching pedagogies could reduce the time needed for students to process and learn information and allow students to flourish. Moreover, students being encouraged and enabled to use generative AI responsibly will upskill and improve their employability, as they will become increasingly familiar with tools that are becoming more commonplace in workplaces.

Efficiency and capability improvements in teaching could be achieved through the application of generative AI technology in areas including:

- lesson plans and outcomes;
- quizzes and formative tests at varying levels of difficulty/knowledge bases;
- group activities tied to specific learning objectives;
- handouts and subject summaries; and
- slide decks and other visual material.

- For assessment purposes, generative AI is viewed as a potentially powerful tool that could facilitate:
- more authentic, real-world assignments;
- analysis of student submissions against a canon of work and provide instantaneous feedback to students; and
- marking of multiple-choice questions tests (MCQs).

Generative AI also has the capacity to assist professional services and administrative staff with routine and repetitive tasks, such as:

- management of credit transfer, degree validation;
- report generation and timetabling;
- real-time smart campus updates, such as changes to lecture/seminar locations; and
- attendance management systems where registration, recording and absences could be directly uploaded to a student's file.

However, successful and sustained use of generative AI for these tasks and processes is predicated on the development of the technology to become more reliable, its ethical use by staff and students, and critical analysis of outputs. A more detailed explanation of concerns over the use of generative AI in education is given in response to question 15.

**(Theme 2 - Question 14) What subjects or areas of education do you believe could benefit most from generative AI tools?**

Due to the vast array of generative AI available, there are several subjects/areas that could benefit from use of these tools. These include:

- technology-related subjects;
- creative arts that use virtual and augmented reality;
- language learning, for interactive and real-time exercises with language prompts and scenario-based speaking, writing and listening tasks, as well as instant feedback on grammar and vocabulary usage;
- computer science, for example using AI to build and de-bug code or using AI to code blockchain in different programming languages;
- medicine and health-related subjects, with the codification of medicine ontologies, modelling of data sets, and experiment simulations: the long-standing use of generative AI in personalised care has guided diagnoses, drug discovery and operational efficiencies;

- Humanities and historical disciplines, through the provision of contextual information, testing student understanding of complex events and facilitating the critical analysis of sources;
- adaptive learning systems: using AI algorithms to personalise the learning experience for individual students (and student groups); and
- data analytics: analysing large volumes of educational data to gain insights and improve learning outcomes. For example: automating marking/grading, providing intelligent tutoring, and generating personalised recommendations.

As noted in response to Question 13, encouraging and enabling students to use generative AI responsibly will upskill and improve their employability, as they will be becoming increasingly familiar with tools that are becoming commonplace in workplaces. Fields such as law, architecture, engineering, business, life sciences, and social sciences, have been identified by Goldman Sachs as having high shares of industry employment (>30%) exposed to AI automation. It is of paramount importance, therefore, that learners in these fields are adequately equipped to understand the impacts of generative AI and prepared for use of these tools in related professions.

**(Theme 3 - Question 15) What are your main concerns about using generative AI in educational settings?**

Several concerns have been raised regarding the use and nature of generative AI in education; successful integration and utilisation of this technology is contingent upon these concerns being addressed. Specific areas for concerns include:

- **Datasets and sources:** large language models (LLMs; chatbots) are trained on publicly available written material, processing information from news articles, scientific literature, books, and even blog posts or reviews that can contain personal or sensitive data.
- **Training practices**: questions have been raised over the labour practices involved in training LLMs and AI platforms. For example, supervised learning, (a human-led process training practice through which large datasets are labelled in order for AI to then recognise patterns) has been accused of being exploitative, with some content moderators for ChatGPT describing 'psychological trauma of reading and viewing graphic content, low pay and abrupt dismissals'.
- **Inaccurate outputs and misleading results**: generative AI can produce results that are misleading or inaccurate due to limited datasets, and there are concerns that users may engage uncritically with the generated material. For image generators, there are concerns about the generation of deep fakes and potential for these to be weaponised for misinformation.

- **Copyright and intellectual property:** some generative AI models use datasets that draw from copyrighted material, raising copyright infringement concerns and potential cause for litigation. For example, a University of Sydney academic prompted ChatGPT to produce passages from a copyrighted text, which it did. Where AI is used to refine work, there are questions of intellectual property and ownership of the created content.
- **Ethical and bias issues:** there are concerns about bias in AI source data, and as institutions still do not possess clear understanding of provenance, there is a risk that generated content could contain biases. There is a particular concern related to how potentially biased content relates to protected characteristics, which is a significant challenge for educational and healthcare settings.
- **Data privacy:** in addition to questions over sources that AI platforms are trained on, it is also unclear how user-submitted data is then used by generative AI technology. There are concerns that users could input sensitive personal information, potentially unknowingly, which is then used as source training data, posing risks to privacy. For example, there have been instances where students are reported to have used chatbots as proxies for mental health therapy and advice, raising major ethical concerns.
- **Corporate ownership and transparency**: ownership of AI tools by large corporations with varying degrees of transparency is a cause for concern, as these will influence the direction and impact of AI use, especially if some tools become integral to institutional/academic practice.
- **Digital divide:** Increasing use of AI hinges on student access to such technology, and those who are able to afford subscription fees, such as for ChatGPT-4's premium model, will be able to benefit from more refined tools. For example, ChatGPT-4 allows users to input both text and visual prompts, and has been trained to reduce its tendency to hallucinate compared to ChatGPT-3.5. If courses are developed to reflect the advancements in AI, and subsequently integrate specific tools, which are then paywalled, this would pose challenges to inclusive learning environments.
- **Critical analysis skills:** With uncertainty over the accuracy of AI-generated content, it is crucial that users critically engage with this content, rather than taking outputs at face value. This is due to unreliable or opaque references, which can also be falsified. AI should help enhance (and not replace) critical analysis skills. Increased human intervention and interaction with AI technology is vital to refine these systems and improve reliability.
- **Academic misconduct:** Concerns have been raised regarding students potentially submitting AI-generated work as their own, or uncritically engaging with the generated material, as stated above. This is significant as 'content from generative AI is a nonrecoverable source as it can't be retrieved or linked' in order to be referenced. Workarounds such as referencing AI-generated content as personal communications, including full transcripts (with prompts) or screenshots do not address the issues of non-recoverability.

- **Difficulties with detecting AI-generated content:** London Higher has previously produced a briefing paper on academic integrity, and notes that tools such as the Turnitin bolt-on Authorship for Investigators scan document metadata and analyse factors such as languages levels across a student's assessment portfolio in order to identify possible instances of academic misconduct. However, as generative AI technology is continuously developing and as chatbots are refined to produce more natural language outputs, current tools may have difficulty detecting what exactly has been AI-generated.
- **Assessment design:** Significant work is required by institutions to ensure that assessments are designed (or redesigned) in a way that acknowledges and allows for pro-active and responsible use of AI where appropriate via full risk management processes. Many London Higher members offer subjects which include Professional, Statutory and Regulatory Bodies (PSRBs) and other accrediting bodies that act as regulators for certain professions and qualifications. Modernised assessment in these fields would foster responsible and ethical use of these tools in the related professions, and would encourage learners to engage critically with outputs, stating where and how they used AI.

The overarching concern across institutions is the impact of AI on academic integrity and the need for risk management processes to mitigate negative impacts. Addressing these concerns requires a multi-pronged approach involving technological solutions, ethical and academic guidelines, and pedagogical strategies to ensure that the benefits of generative AI can be harnessed whilst acknowledging and mitigating risk.

Outside the integration and use of generative AI technology in education, there are clear concerns over corporate ownership, transparency, accountability and labour practices, and wider unease over accuracy of information, data privacy and security. While these remain unaddressed, there will understandably be hesitation to fully embrace AI within educational settings. Further detail over views regarding ethics, data privacy and security when using generative AI can be found in response to Question 18.

**(Theme 3 - Question 16) If at all, have these concerns impacted your use of generative AI? Please explain how.**

It is too early to say how concerns have impacted institutional practices surrounding the use of AI. Currently, policies and procedures advise caution and make clear any potential risks associated with AI-generated content (and improper use thereof) in work. Moving forward, the successful and well-managed integration of AI in educational delivery will be guided by updated:

- institutional policies (such as for academic misconduct);

- ethical guidelines and codes of conduct; and
- best practice guides.

The aim of these management approaches is to minimise risk and maximise benefits for learning experiences. Responsible and ethical use of generative AI is the crux of the issue; students are not yet fully able to understand limitations, inaccuracies, legal or ethical problems and the extent to which these could have an impact. They will need clear guidance from their institutions, who will in turn need time to discuss and update relevant frameworks to ensure continued academic robustness.

**(Theme 3 - Question 17) Are there specific subjects or areas of education where you believe generative AI should not be used? Why?**

Questions have been raised regarding the suitability of using generative AI for subjects that teach concrete skills that are then automated by AI, for example, the use of chatbots (large language models) when teaching writing, or image generation models (such as Midjourney) when teaching illustration. It is noted that in these instances, AI could still be used when the learning is focused on conceptual elements rather than the skills.

In general, the range of tools on offer is extremely broad, and it is likely that any discipline or subject area may be able to find a use for aspects of generative AI. As the potential of this technology is not yet fully realised, it is too early to suggest whether there are areas in which use should be prohibited or discouraged. For the time being, institutions are focused on ensuring that the benefits, risks, and boundaries for use are as clearly communicated to students as possible, so that any negative impacts may be managed.

**(Theme 4 - Question 18) If any, what are your views regarding ethics, data privacy and security when using generative AI in education?**

Ethics, data privacy and security are key issues surrounding generative AI in higher education. There are several specific issues that are, as yet, unclear or unresolved, such as:

- misuse of copyright material;
- corporate ownership and lack of transparency;
- generation of inaccurate information and misleading content (deepfakes, false references);
- data collection, consent and informed participation;
- use of private data submitted as prompts to large language models;
- training and source data, including bias;
- academic misconduct;
- concerns over labour practices involved in the training of AI;

- ownership of content and intellectual property; and
- fairness and accountability.

Policies, practices and partnerships will prove key to addressing/remediating these issues, for example:

- policies must be updated on an iterative basis, taking into account new developments and capabilities of generative AI technologies;
- practices must be guided by sector-wide and institutional policies, with ethical, responsible and fair use at the core. Moreover, dialogue within institutions and involving students in developing practice is necessary. This can increase trust, buy-in and adherence to academic conduct guidelines, recognising that students may have expertise with these technologies that teaching staff may not, and highlights co-operation as key to advancing understanding and fair use; and
- partnerships with trusted suppliers or developers of generative AI will be crucial to alleviating several concerns outlined above and throughout this submission.

For professional services' use of AI, for example, the incorporation of student data (including demographics, academic records, learning preferences, attendance and behavioural data), potential consequences like identity theft, fraud, academic malpractice, discrimination and harassment are of concern in the event of any data breaches or misuse.

Although initial reactions to generative AI included an outright ban, this is no longer considered feasible. With such a wide variety of tools on offer, various educational areas that could be positively impacted, and real-world applications of technology that students will encounter beyond university, it is crucial that higher education provides a strong steer for ethical and responsible usage. Generative AI must be approached critically, and staff and students alike must understand ethical implications, GDPR, data privacy and security issues so that they can be as informed as possible when using these tools.

**(Theme 5 - Question 19) How do you see the role of generative AI in education evolving in the future?**

AI is a rapidly developing technology – ChatGPT, which has prompted much of the increased recent discourse, was only released in November 2022. There is therefore a need for the education sector to move quickly and in the right way, and it must be supported to do so; as AI evolves, so too will its use in education. As covered in more detail in Question 13, there are several ways in which generative AI could be used in educational settings by several different stakeholders.

For students, there is potential for learning experiences to be personalised through generated materials that account for individual learning styles. With assessments, generative AI could potentially facilitate analysis and provide instantaneous feedback to students.

For academic and teaching staff, generative AI technology can be used to assist in the creation of pedagogical materials and course planning. For professional services and administrative staff, generative AI may increasingly be able to assist with routine and repetitive tasks.

The discussion surrounding the impact of AI in assessments is likely to gain traction, as institutions will have to:

- revise and update academic integrity guidelines to keep pace with developments in technology;
- reflect on the human interaction element of AI generated content and the development of staff competencies in the assessment process. In particular, institutions must recognise the limitations of AI detection tools, and avoid overreliance on these. This is, in part, due to a time lag, as design and development of detection tools follows advancements in AI technology itself; and
- redesign assessments to incorporate fair and ethical use of AI whilst deterring students from improper or unethical use.

On a practical note, Microsoft Copilot integrates large language models (LLMs) with organisational data with an aim to increase productivity, and Google's Bard will be similarly integrated into Google Suite for the same purpose. Therefore, higher education institutions that use these platforms could see generative AI increasingly embedded and accessible for use in their work through their digital infrastructure. For teaching and learning, this can facilitate learning design, feedback, assessment support, and data analysis.

Generative AI has the potential to facilitate more effective application of knowledge, however as it develops, premium/specialised models could potentially move from open access to subscription based. As highlighted in Question 15, this could disadvantage students who cannot afford to access these paid models. Institutions will need to be aware of the implication this could have in courses and assessments, for example, if a course is designed to integrate a certain tool which is then paywalled.

**(Theme 5 - Question 20) What support do education staff, pupils, parents or other stakeholders need to be able to benefit from this technology?**

Strong leadership is needed to increase staff and student understanding of the potential, limitations, benefits and risks of generative AI technology. Effective integration of tools in teaching and learning, as in professional services, will require training and best practice guidelines. The use of AI could prove to be a game changer in education, and institutions will need adequate support in order to maximise benefits and flourish.

Experts on several aspects of AI (including but not limited to ethics, change management and risk management) available for consultation would also facilitate transitions to more AI-ready educational and professional environments. In addition, financial and strategic support for infrastructure developments will be needed. It is imperative that educators develop foundational knowledge of generative AI technology, in order to effectively implement, integrate and manage this technology in higher education. Upskilling of employees, much like during other technological advances (such as the internet), is fundamental. Up-to-date and continuously developed training is necessary in order to ensure that educators have sufficient knowledge to use and teach generative AI and its applications.

**(Theme 5 - Question 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?**

The Department for Education has an important role to play in providing leadership and guidance to the education sector. This could include:

- offering training and best practice guides for the use of generative AI and ethics, GDPR, privacy and safety, that could be adapted to fit higher education settings;
- working with institutions to develop further training for staff and students, therefore enabling responsible and safe use of technology in higher education settings;
- financial support for the transition to AI-appropriate infrastructure (noting similar funding programmes such as the Office for Students' capital funding allocations for campus improvements);
- funding research into generative AI development and applications of this technology in education;
- establishing clear guidelines for research, development and innovation of these technologies by and in institutions;
- establishing ethical guidelines for the owners of generative AI platforms, addressing the concerns raised by institutions surrounding security, privacy, ethics, transparency, copyright, and labour practices;
- establishing communities of practice or expert networks that can inform the development of policy and regulation of AI; and
- recognising the importance of AI and technology in the training of future educators and assessing digital literacy as a core competency as part of the award for qualified teacher status and/or PGCE, as well as encouraging academic teaching staff in higher education to receive support for integrating AI into their teaching through trusted and appropriate sector bodies, such as Advance HE.

As mentioned in response to Question 15, many London Higher members offer subjects which include Professional, Statutory and Regulatory Bodies (PSRBs) and other accrediting bodies that act as regulators for certain professions and qualifications. DfE could coordinate these stakeholders to develop a shared understanding of generative AI and ensure that assessment practices under their remit are likewise updated to reflect use of this technology. Limiting approaches (such as proctoring) can be circumvented and may only encourage improper use of AI technology. Modernised assessment would therefore foster responsible and ethical use of these tools, encourage learners to engage critically with outputs and state where and how this technology was used.

**(Theme 5 - Question 22) Is there anything else you would like to add on the topic of generative AI in education?**

This submission has demonstrated that the use of AI – generative or otherwise – is a complex issue within higher education. There are clear benefits to its use, such as the personalisation of learning experiences, automation of routine tasks and the streamlining of administration. However, significant concerns have been raised around ethics, security, potential for misuse and academic misconduct, inaccurate information and more. Comprehensive risk management processes, institutional approaches and sector-wide guidance will be required for the effective integration and utilisation of AI. Moreover, leadership from, and dialogue with, bodies such as the DfE and OfS will be critical in establishing common principles and frameworks that allow institutions, staff and students to use AI responsibly. As technology develops, the sector must adapt and its infrastructure must be resilient in order to mitigate negative impacts.

London's higher education sector welcomes the opportunity for further dialogue with the DfE and other key stakeholders on the use of generative AI technology: this should be seen as a priority area with implications that could represent a step change in the delivery of teaching and learning throughout the entire education sector.