

Education Incubator ChatGPT and Higher Education Project: Final Report

Elliott Colyer, Holly Frostwick, Barrie Cooper, Judith Kleine Staarman, Callum Cant

University of Exeter

Abstract

ChatGPT (GPT-3.5) was launched in November of 2022, with the more advanced model GPT-4 following in March of 2023. In light of these new tools, much has been said about the possibilities they may create for academic misconduct or unethical academic practice. As a result, there is a need for higher education institutions to respond to the potential risks and rewards of such technology through clear, informed policy decisions that take account of the views and opinions of all their stakeholders. Throughout April to June of 2023, the Education Incubator's ChatGPT project ran a series of workshops, focus group interviews, and surveys, designed to gather the opinions of the various members of the University of Exeter (henceforth 'the University') community regarding the role of generative AI in higher education. Following this we have produced a series of recommendations for the University concerning learning, assessment, workload, and acceptable use. In particular, we encourage the University to ensure that access to a baseline of generative AI tools is equitable, that everyone is given an opportunity to learn how to use generative AI effectively and ethically, and that staff are given flexibility in ensuring that their teaching and assessment methods are reflective of the current technological climate.

"A bad idea is to ignore it or to try to fight it."



Abstract	1
Introduction	3
Summary of the Recommendations	4
Methodology.....	5
Use Cases	6
Student Use Cases	6
Staff Use Cases	7
Comparing Student and Staff Use Cases	8
Promoting Inclusion and Accessibility	11
Generic Recommendations	12
Access, Cost and Equality.....	12
Staff Training and Workload Implications.....	13
Changing Skills, Course Content and Pedagogy	13
Implications for Assessment	15
Clear Policy	18
Faculty Specific Use Cases.....	18
Faculty of Environment, Science and Economy	19
Business School.....	20
Faculty of Health and Life Sciences.....	21
Faculty of Humanities, Arts and Social Sciences	21
Conclusion.....	22
Appendices.....	24
Defining a Baseline of Generative AI	24
Possible Assessment Changes	25
Novel Use Cases	26
Survey responses	27
Focus groups	39
Looking Ahead	41
Bibliography.....	42

Introduction

Throughout the Spring and Summer of 2023, the University of Exeter's Education Incubator ran a project to gather the experiences and perceptions of the university community about ChatGPT and generative AI more broadly. Our aims for this project were to:

1. Identify the threats and opportunities of AI systems;
2. Understand the ways in which AI will affect teaching, learning and assessment, in addition to broader effects it might have on pedagogy and skills;
3. Allow staff and students to contribute to a long-term vision for how AI systems can be appropriately integrated into teaching to produce a benefit for the university community;
4. Produce relevant outputs to support different parts of the university in responding to ChatGPT and similar AI technologies in the short- and long-term.

To that end, we gathered ideas and opinions from staff and students through workshops, surveys, and focus groups (see below for more details) and the outcomes of this project are presented in this final report. Particularly the data gathered through the workshops we conducted in April and May 2023, have informed the recommendations in this report regarding the support and guidance the university community needs and wants in this area. Our recommendations are thus entirely guided by the university community and those who shared their thoughts, opinions, and concerns with us.

It is also important to note that this report, whilst being the final report of this Education Incubator project, is only one part of the university's response to AI in Higher Education.

We intend to further this work through future research and analysis for academic publications. Moreover, the discussion of university policy, guidance and practice is far from over – this is a nuanced and fast-changing area, and so policy, guidance and practice must be equally flexible. We welcome further discussion on the role of generative AI in higher education and hope this report can be a catalyst for these continuing discussions.

To that end, if you have any comments or feedback about this report, the project as a whole, or the overall topic of generative AI in Higher Education, please get in touch with either the Education Incubator (educationincubator@exeter.ac.uk) or the Academic Development Team (academicdevelopment@exeter.ac.uk).

This report is structured as follows:

We begin with a **Summary of the Recommendations**, which provide an overview of the recommendations we make throughout this report. The **Methodology** then outlines the ways in which we have collected data, after which we present **Use Cases**, which allow us to discuss and compare ways in which staff and students use generative AI and how they think generative AI may be used in future teaching and learning. This section also includes a comparison between the ideas of staff and students and considerations on how generative AI may help or hinder efforts to promote inclusion and accessibility within the university community.

Next, we provide **Generic Recommendations**, and we explain of each of these, their rationale, and any relevant problems or concerns as put forward by staff and students. This section also presents implications for assessment. In addition, the **Faculty Specific Use Cases** present a faculty-by-faculty discussion of unique use cases, opinions, or suggestions that are specific to that faculty. The report ends with a **conclusion** and **appendices**, which provide a deeper insight into the specific data and issues we report on.

Appendices:

- A. **Defining a Baseline of Generative AI:** throughout the report, we refer to a ‘baseline’ of generative AI. The field of generative AI is changing so quickly that recommending any one system will quickly become obsolete. To that end, we define what we mean by a baseline, in terms of the computing power and task capabilities of generative AI tools.
- B. **Possible Assessment Changes:** a brief summary of all the possible new forms of assessment that were suggested throughout our research. These suggestions are drawn from discussions with all of the faculties, and are presented in the hopes that they may inspire discussion and innovation across the university community;
- C. **Novel Use Cases:** a brief summary of the novel use cases raised throughout the workshops, again from all faculties;
- D. **Survey Responses:** a summary of the data that was collected from the responses to the staff and student surveys we ran. The surveys included multiple-choice ranking style questions (for example, “How knowledgeable do you feel about how you can use generative AI in HE/your study?”) as well as longer comment responses;
- E. **Focus Groups:** a summary of the focus groups we ran in June with senior Academic Conduct Officers and international students;
- F. **Looking Ahead:** a description of the ways which generative AI may be developed in the future to the benefit of the university community.

Summary of the Recommendations

Each of our recommendations will be expanded upon in more detail below, but we summarise them here:

1. Access, cost, and equality
 - 1.1. The university must ensure that all staff and students can access a baseline of generative AI tools such as ChatGPT, regardless of their financial ability to do so.
 - 1.2. We recommend that all members of the university community must have equal access to a baseline of generative AI tools and that no course or assignment requires access to a fee-paying generative AI service.
 - 1.3. The university must also ensure that those who object to using a specific generative AI tool for ethical or other reasons will not be unreasonably expected to do so.
2. Staff training and workload implications
 - 2.1. The university must ensure that staff have the time, space, and opportunity to properly engage with this technology and decide how to change their courses and assessments in light of it.
3. Changing skills, course content, and pedagogy
 - 3.1. There is a strong desire among both staff and students for education and training in effective and ethical prompting skills.
 - 3.2. Both staff and students also understood that generative AI such as ChatGPT often produces inaccurate or misleading answers, and so both groups would benefit from education and training in how to spot those inaccuracies and check them.
 - 3.3. There is also an understanding that critical thinking and creative skills will become more important and valuable, both in academia and industry.

- 3.4. These are skills that must be embedded into intended learning outcomes and marking criteria.
4. Implications for assessment
 - 4.1. Staff and students are aware of the availability of developing AI detection software (via Turnitin or AI tools), but instead would prefer to alter assessments to meet the changing environment of higher education.
 - 4.2. Even if some assessments are changed there will still need to be greater provision for invigilated in person assessments, whether that be traditional exams, presentations, or vivas.
 - 4.3. Some faculties which require key competencies to be assessed, or which have externally set and accredited assessments may have limited ability to change assessments.
 - 4.4. The university must further allow flexibility for staff to develop and change course content and assessments more frequently if needed, rather than the current exception-based system for in-year module changes.
 - 4.5. This process will also have significant implications for staff workload models.
5. Clear policy
 - 5.1. Whatever the university's decision on generative AI usage is, the policy must be clear and specific.
 - 5.2. Many students and staff are already confused and concerned about what is and is not considered ethical usage by the university.

Methodology

The main data for this project was collected through a series of staff and student workshops run in a hybrid in-person and online format during April and May of 2023. Each workshop was targeted towards staff or towards students of a specific faculty, to gather faculty-specific use cases, concerns, and suggestions. The workshops for staff and students were held separately to allow, and encourage, students to be open about their usage of generative AI tools without the risk of penalty. Both staff and students were asked to experiment with ChatGPT (with subscriptions to OpenAI's ChatGPT Plus available for those in-person participants), and then discuss their thoughts, comments, and ideas, recording these on an online collaborative whiteboard (Miro, see Fig. 1). Specifically, participants were asked about:

1. Their feedback on the answers they received from ChatGPT;
2. Their suggestions for current use cases among both staff and students at the university;
3. Their beliefs on how many people at the university are using ChatGPT;
4. The ways that access to generative AI may affect learning, skills, and assessment in their field; and
5. Their thoughts on specific suggestions for ways to incorporate AI into higher education.

In addition to the responses from participants collected using the online whiteboard, we also recorded conversations from participants in Teams break-out rooms, as well as audio-recorded conversations of groups who participated in-person. The data was analysed thematically, using NVivo, to identify patterns and themes among the data, and within individual faculties. This allowed us to determine any concerns or suggestions that were unique to any given faculty, as well as those that were widespread among the participants.

More than 250 staff and 60 students engaged across the eight workshops (attendance log numbers are higher, but some online attendees only joined briefly), resulting in eight Miro boards (see Figure 1 below for an example), more than 50 hours of audio recordings and more than 50 ChatGPT chat logs to analyse.

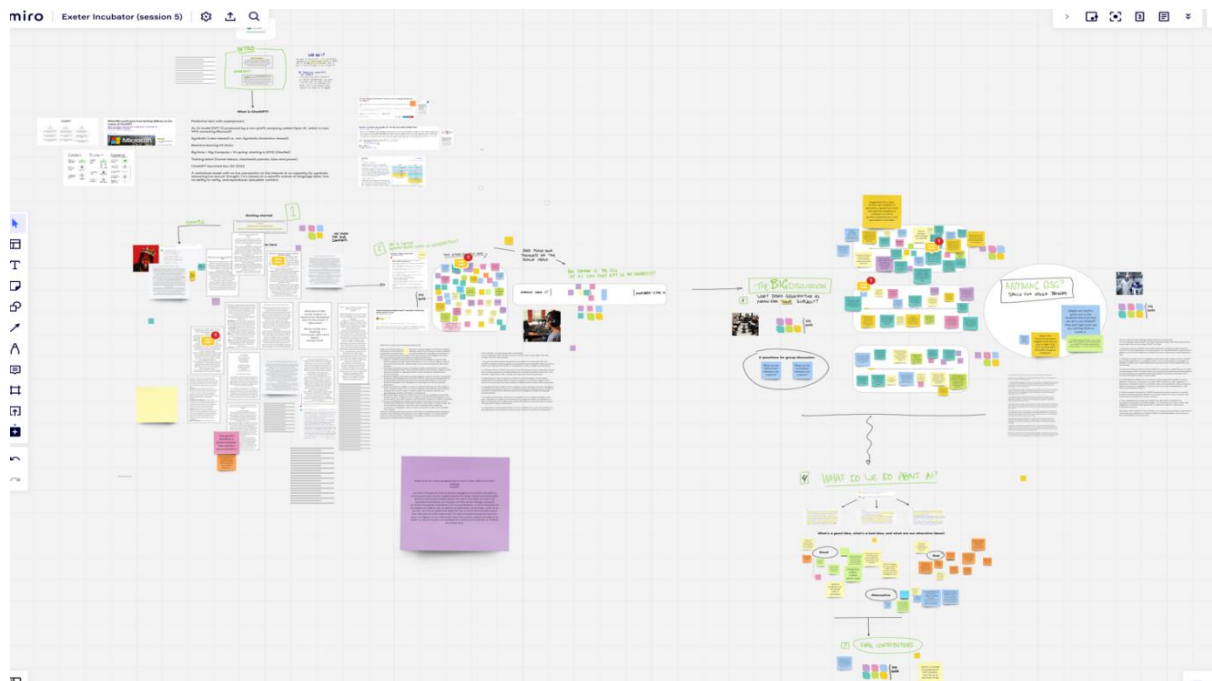


Figure 1: A completed Miro board from a staff workshop.

Subsequently, we carried out separate staff and student surveys to obtain further feedback on generative AI use cases and concerns. We also conducted two focus groups to explore key themes with specific stakeholders.

The main body of the report addresses the findings from the workshops, with supplementary analysis from the surveys and focus groups appearing as appendices.

Use Cases

In this section, we explore the primary use cases for both students and staff, as revealed through workshop discussions. Following this, we draw comparisons between the two groups, spotlighting any shared perspectives or stark contrasts in their attitudes and applications of the technology.

We have included suggestions that were made throughout the workshops on possible use cases for generative AI in an abridged form in the appendix 'Novel Use Cases'.

Student Use Cases

The students who attended the workshops raised a number of current and possible use cases. These were often far more sophisticated and nuanced than the staff we spoke to anticipated, and certainly more developed than the media coverage in this area would suggest. In particular, no student who attended a workshop mentioned using ChatGPT to do their work for them. Instead, they mentioned use cases such as the following:

1. **Coding help**, including creating code, helping to debug existing code, and interpreting complex abstract results.
2. **Accessibility support**, including describing visual figures for visually impaired students, language support for international and English as an Additional Language (EAL) students, reworking written text into clearer forms, and making skills more accessible to neurodiverse students.
3. **Research help**, including helping find sources, summarising difficult readings, and explaining tricky concepts (particularly for inter-disciplinary readings).
4. **Proof-reading and language support**, to help students improve their own original writing, and tidying Teams transcripts.
5. **Tutoring and explanations**, including using ChatGPT as a ‘study-buddy’ style assistant to prepare for upcoming seminars, as a language tutor for conversational practice, and as an interactive textbook or dictionary for querying specific problems. It is particularly worth noting in this last instance that some students expressed that they turned to ChatGPT for help if they felt that their tutor or lecturer would be unsympathetic about ‘trivial’ problems and concerns.
6. **Career help**, including CV advice, cover letter drafting and editing.

As such, the use cases were wide ranging and included using ChatGPT to give *“a very surface level insight into different topics and how they might connect to each other”*, using it to *“interpret my statistics results”* into comprehensible language, and *“summarising the lesson for the week”* to better prepare for seminars and lectures. Students also emphasised the importance of a strong knowledge foundation when using ChatGPT, due to its inaccuracies and hallucinations.

While it is important to recognise that students may not have been comfortable disclosing if they had used ChatGPT in ways that may be perceived as more unethical, we believe that the discussions that were had suggested a more sophisticated understanding and usage than simply using generative AI as an essay-writing tool.

Staff Use Cases

“It’s an accelerator, it speeds up certain workloads.”

During the workshops, staff identified a number of areas where students might use generative AI tools such as ChatGPT, as well as possible use cases for themselves in both teaching and research (see Fig 2 for a word cloud illustration). Staff-specific use cases that we saw appear throughout the workshops, regardless of faculty or subject area, included the following:

1. **Literature reviews**, including finding sources, as well as summarising them briefly to expedite this element of the research process.
2. **Drafting writing**, including producing drafts of emails, funding applications, and abstracts. It is worth noting that most use cases staff presented in this area concerned writing the more formulaic and less important areas, so that they can focus more strongly on the core of their research.
3. **Proof-reading and editing**, including translation.
4. **Career help**, including CV advice, cover letter drafting and editing.

whereas students, insofar as they are using it for submitted work, often mentioned using it to improve their own original work. Table 1 below highlights a number of use cases mentioned by both staff and students, and states whether these are uses already performed (retrospective), or possible future uses (prospective). It also highlights whether these uses are intended to occur at the beginning of the research and writing process (foundational), or after work has already been produced by the user, with an aim to improve that work further (developmental):

Table 1: Selected use cases by staff and students.

Staff/ Student	Retrospective/ prospective	Summary	Quote	Foundational/ developmental
Staff	Retrospective	Foundational work	<i>"It gives you the basis for you to build on"</i>	Foundational
Staff	Retrospective	Starting point	<i>"It was a good starting point"</i>	Foundational
Student	Retrospective	Statistical analysis	<i>"I had it interpret my statistics results from R"</i>	Developmental
Student	Retrospective	EAL assistance	<i>"As an international student, I have very limited writing skills, and then ChatGPT can help me to paraphrase my sentence"</i>	Developmental
Student	Prospective	EAL assistance	<i>"I think what ChatGPT could help with is, at least for most of the folks who don't have English as a first language, it helps with the barrier"</i>	Developmental

Indeed, students seemed to use it as a revision tool (an interactive textbook or study partner), and for improving existing work, far more than they did for generating the content itself. However, staff show more concern about the impact on traditional skills and discovery processes in learning and research.

It is also worth noting that during the student workshops and focus group, participants mentioned other generative AI tools they were using (Grammarly, Quillbot, Perplexity AI). These were often more specialised tools for their task, however staff discussions had very little mention of other AI tools.

Finally, it is interesting to explore the differences between the use cases reported by students and the imagined use cases of students by staff (see Figures 3 and 4). While there are a lot of similarities, it is interesting to note that staff mentioned that they think students are using ChatGPT to paraphrase and copy texts ('write short summaries', 'paraphrase texts', 'conduct literature reviews') and to generate content for assignments ('copy/paste essay answers', 'writing dissertations', 'complete assignments', 'write parts of essay'). Yet, students did not mention these kinds of uses. Instead, they emphasised its use to help rewrite their own ideas ('rewrite for conciseness', 'communicate ideas as EAL', 'rewriting academically') and to help them with their understanding of course content ('suggest further reading', 'guidance on tasks', 'signpost key authors', 'summarising teaching notes'). While we would like to reiterate again that we fully recognise that students may not have wanted to disclose certain 'less ethical' ways of using ChatGPT, this difference may indicate that students are, at the very least,

Promoting Inclusion and Accessibility

The workshop discussions highlighted the potential of ChatGPT in addressing accessibility and inclusion concerns, particularly for EAL learners and those from diverse educational cultures. Participants viewed the AI tool as a means to help these students navigate language barriers and unfamiliar academic norms.

While there were limited direct discussions on how generative AI could impact Equality, Diversity, and Inclusion (EDI), several points indirectly alluded to EDI considerations. For instance, ChatGPT could serve as an interactive tutor or learning aid, potentially assisting students with learning difficulties, EAL learners, and possibly visually impaired individuals, when combined with other assistive technology. A visually impaired student we spoke to explained that ChatGPT was particularly helpful for understanding the context of pictures—why they were used, their placement, and their content. They noted that such nuances aren't captured by the alt-text field and typical text readers lack interactivity.

ChatGPT opened up a new level of detail, functioning like an ever-available, patient personal assistant. The ability to ask follow-up questions for clarity, to inquire about a picture's positioning or specific features, proved incredibly beneficial. The student noted, *"It definitely helps to feel much more included, and it is now much easier to have a common frame of reference about things to talk with my friends"*. They added that *"ChatGPT can patiently explain what seems very trivial information for others."*

Our workshops were primarily focused on ChatGPT usage, as this was the most widely known and accessible tool for our participants. However, this is obviously not the only generative AI tool that is available, and many students mentioned using multiple AI tools in conjunction with each other to better aid their access to academic content. In particular, the visually impaired student we spoke to mentioned using AI-powered image description tools to first produce a detailed description of each image, which can then be given to ChatGPT in addition to the image's surrounding context in order to better aid their understanding.

Additionally, EAL individuals mentioned using multiple different AI tools to help improve their English writing, as well as their understanding of difficult texts. This proliferation of different tools may be particularly useful for EAL students to aid their learning, as ChatGPT may be less sophisticated in languages other than English. Indeed, the training data for ChatGPT-3.5 was a filtered version of the Common Crawl dataset (Brown *et al.*, 2020, p. 8), which is roughly 45% in English, with the next most common languages constituting only around 5-6% of the data (Common Crawl, 2023). As a result, the university must be aware of the possible uses of these tools.

It is also important that with a potential shift towards alternative assessment methods, such as viva-style exams or practical assignments, EDI is appropriately considered. These diverse forms of assessment may accommodate varying learning styles and abilities, fostering inclusivity. However, they must be implemented with care to avoid unintentionally disadvantaging certain student groups, such as those with anxiety disorders, as one participant pointed out.

Generic Recommendations

Five main recommendations emerged from the workshop analysis. Each of these is motivated and discussed in more detail below.

- Establish a baseline of generative AI technology freely available to all staff and students.
- Ensure the university allocates sufficient time and resources for staff to engage and explore this technology and consider necessary adjustments to their courses and assessments.
- Implement a universal university module to enhance students' competency and key skills in utilising generative AI.
- Staff must consider the capabilities of generative AI when setting assignments. The current exemption based in-year module change system will need to be revisited.
- Implement a clear, user-friendly policy that undergoes regular reviews to stay in step with advancements in generative AI technology.

Access, Cost and Equality

Establish a baseline of generative AI technology freely available to all staff and students.

During this project, staff, departments, and services made it clear that they wish to have access to ChatGPT-4, both through discussions in the workshop and by contacting us directly. Further, the institutional statement (Education Board, 2023) clearly expresses an expectation that students will need to be supported to learn and use these tools. A university-wide policy is needed around the provision of a baseline of generative AI access for all. Further details of what a baseline of AI tools could constitute is included in the appendix 'Defining a Baseline of Generative AI'.

Significant concern was raised by staff and students about the possibility for generative AI to exacerbate existing socio-economic division among the student population. Given that OpenAI's ChatGPT Plus model, which gives access to the newer ChatGPT-4 model, was shown to "outperform existing LMs on all benchmarks" (OpenAI, 2023, p. 7) including the free-to-access ChatGPT-3.5, there is clearly a concern regarding accessibility. A monthly ChatGPT Plus subscription currently costs \$20 USD (£16.02 at time of writing), for a total \$240 USD a year (£192.20). Yet the recent Student Living Index 2022 found that "35% of students have found themselves running out of money by the end of term" (NatWest, 2022). Hence, this additional financial burden would be inaccessible to students already struggling financially.

It is therefore our strong recommendation that, if the resulting policy regarding generative AI would be to allow or encourage its usage, then effort is made to ensure that a baseline of access is available for all students. This should be free to all students and staff.

There were students and staff who expressed concern about the ethical issues surrounding generative AI tools, and ChatGPT in particular. This includes concerns regarding bias encoded within large language model training data (Bender *et al.*, 2021), the working conditions of those enabling moderation of existing models (Perrigo, 2023), and the environmental impacts of AI research (Strubell *et al.*, 2019). There will therefore be active and passive resisters to using specific generative AI tools, for these ethical and other reasons. As such, the university needs to understand and make provision for staff and students who hold such views. Namely, no member of the university should be unreasonably required to use any specific generative AI tool.

Staff Training and Workload Implications

Ensure the university allocates sufficient time and resources for staff to engage with and explore this technology and consider necessary adjustments to their courses and assessments.

The participants who attended our workshops presented the full spectrum of confidence using AI tools, from those who had never used it before to those who had used it in many different areas of their professional and academic lives. Both staff and students expressed a desire for support, training, and education on how to use ChatGPT (and generative AI more broadly) effectively, usefully, and ethically. Participants who were the most worried about the implications ChatGPT may have on academic integrity and ethics became significantly more confident having had the opportunity to use ChatGPT and understand its capabilities and limitations.

“Or nobody’s engaging with it because they are too scared? I’m definitely too scared.”

Staff further expressed a desire for the time, space, and support to understand ChatGPT (and other generative AI services) and to develop assessments and teaching approaches that take account of its abilities. Yet there are significant workload implications that come with this, which must be considered. However, without this opportunity there will also be significant workload implications if staff were to continue to assess students as they traditionally have.

“We’re going to need to rethink how we frame our assessments; we can’t just roll out the same way that we’ve been doing it year after year.”

Traditional coursework assessment methods risk obsolescence without considering generative AI, but it is impractical to scrutinise each assignment for false or inaccurate AI-generated references. As such, without appropriate support in this area, staff will likely resort to invigilated closed-book examinations to manage the situation and assure the credibility of their grades. This is despite the expressed desire of many students and staff to avoid this scenario. A prevailing, but by no means uniform, view was that this type of assessment is inappropriate and unreflective of the modern world.

“Our response has been to make them exams again, which I don’t think was the right move.”

Therefore, it is vital that the University factors in the workload implications of adapting to generative AI to achieve optimal results. Overlooking this aspect may result in less favourable outcomes and constrain staff from fully leveraging the benefits of generative AI. This concern is especially critical due to the potential of establishing a negative ‘path dependency’. Initial, suboptimal adaptations to generative AI could set a precedent, leading to future instances where staff continue to grapple with maximizing the potential of generative AI.

Changing Skills, Course Content and Pedagogy

Implement a universal university module to enhance students’ competency and key skills in utilizing generative AI.

“To see how this might change the landscape in sense of what employers want, I think that’s critical.”

“Could we include it [generative AI] as a key module or component of a course for a humanities student?”

A central theme in the workshops of both staff and students was how skills, learning, and assessment could and should change in light of an increasingly AI-enabled world. Staff and students recognised that significant changes to required skills would emerge, in particular a greater focus on critical thinking and creativity. With the ability to offload much of the writing and recall process to AI tools, there was a strong sense that these more abstract skills would become increasingly valuable, both in academia and industry. During the workshops, participants frequently mentioned the potential of using ChatGPT (or generative AI more broadly) for research assistance – specifically, to locate pertinent literature and subsequently provide succinct summaries of these materials. Given this ability, a student’s ability to ask questions about readings, and hence inspire further research, will become more relevant than simple reading comprehension and summarisation.

There was also significant discussion about the growing importance of developing sophisticated prompting skills. This was a suggestion that came up particularly with those workshop participants that had more experience using generative AI; the more one uses these tools, the more one learns how to get the best out of them. Students also wished to see training in the use of generative AI, considering the role it will play in industry. This training should include prompting skills as well as how to critically evaluate AI generated content, given the clear cases of AI hallucination and inaccuracies.

Students recognised the importance of a solid foundation of knowledge when using generative AI, so it is imperative to ensure that their education includes a strong focus on these foundations, to ensure that students are not led astray by AI hallucinations. Some students also raised that ChatGPT can be very easily confused (*“I’ve had an experience where it will actually go back on correct answers and give you completely wrong answers for very simple things”*), and so it is important that the training includes how to fact check ChatGPT’s (and other generative AIs) answers, as well as recognise when they might be being misled.

“Verifying the validity of information and recognising when to engage directly with original sources will be even more important.”

Many university members expressed a wish for this generative AI training to be a generalised module available for everyone within the university to develop these skills in both prompting and fact-checking.

“[There should be a] faculty wide elective, non-assessed module based on teaching students effective prompts and uses of ChatGPT.”

It is important that these new generative AI skills are embedded into intended learning outcomes, and that marking criteria are revised to enable the evaluation of these skills. Ensuring that students are developing these skills, and being meaningfully evaluated on them, will provide them greater success and desirability as employees as they move into their careers.

A common staff concern was that, if we do embrace ChatGPT and generative AI fully, then students may be encouraged to use it to offload much of their thinking, and hence fail to develop strong critical analysis skills.

“I worry about the effect that it’ll have... students are getting more and more hoop-jumping, where they don’t care about that critical thinking element of how to improve.”

While concerned with this issue, staff also recognised that this technology is likely going to be a large part of students’ careers as industry increasingly embraces generative AI tools. As such, there was a

strong recognition that every member of the University must learn to understand and use this technology.

“From an employment context, this is going to be used throughout business.”

We must also stress the possible sources of conflict in this area, as staff and students responded to these issues in distinct ways, and these distinctions will need to be considered and resolved, where possible. Staff and students had differing opinions about the importance of the following skills:

1. **Writing skills:** students felt that, in an age where the bulk of the writing can be performed by generative AI or similar tools, their own writing skills will become less important. By contrast, staff focused heavily on the importance of writing skills, and in particular the need to develop a unique writing style and voice.
2. **Automation and productivity:** there was a discrepancy between staff and student opinion on the value of AI's increasing ability to automate certain tasks to increase efficiency and manage students' workloads. Students found this as a welcome change, feeling that it would help them save time and work more productively, however staff were very concerned that doing so would lead to students cutting certain corners and, as a result, losing out on valuable opportunities for critical engagement and analysis.
3. **Literature engagement:** as we saw above, one of the most common use cases we heard from students was to use ChatGPT to summarise complex ideas and academic writings. They emphasised that doing so enabled them to better understand texts that might have been difficult to understand due to their writing style or assumed level of knowledge (this is something we particularly heard from those students studying subjects that often had crossover with other disciplines, such as medical students grappling with complex biochemical theory). Staff, however, were concerned that doing so would lead to students failing to engage deeply enough with the literature, and thus missing out on valuable skills and insights.

Generally, staff may worry about the potential devaluation of important academic skills, such as detailed research and academic writing, while students might view the evolution of these skills as a natural response to technological advancements.

Implications for Assessment

Staff must consider the capabilities of generative AI when setting assignments and how they can best engage with this. The current exemption based in year module change system will need to be revisited.

While tools such as Turnitin are currently developing AI-detection software, and certain AI tools claim that they are capable of detecting AI-generated content, the workshops we ran suggested that neither staff nor students would value this as the solution. There were many concerns raised about the accuracy of these tools (and, indeed, recent cases suggest that this is a warranted concern (Marcus, 2023)), as well as the relevance assuming other changes to assessment were made. Indeed, participants expressed the sentiment that assessments have not been the best and most helpful they can be for some time now, and staff and students would welcome the opportunity for that to be changed.

“In lots of cases, exams should have been reconsidered already years ago.”

“We have been relying quite a lot on the academic essay, as I think it’s a matter of heritage. ... maybe this is the technology that will allow us to, it will compel us to rethink the academic essay.”

Generative AI poses problems for the current methods of assessment such as essays due to its powerful capabilities in text generation. It is important that assessments give students the opportunity to demonstrate their knowledge and skills. Most students and staff held positive views about the use of generative AI in higher education. However, the one consistent reservation among participants concerned students handing in text that had been AI generated without modification or acknowledgement.

“I have no problem using it for revision. I have no problem with them using it to look at past papers and work out good answers and points they might like to make. I don’t see as a problem there, but I think there is a problem in terms of actually handing in work that’s been generated by a chat bot.”

“[If a student] asked a question in there [ChatGPT] and then just copied and pasted it and submitted that as a piece of work ... it is an authorship problem ... but if they are not doing that [instead using it as a tool] ... is it actually wrong?”

Whilst there was a great deal of support and enthusiasm for exploring new modes of assessment that embrace the use of ChatGPT (and generative AI more broadly), there was also recognition that there may need to be an increase in invigilated forms of assessment in conjunction with these AI-enabled assessments. This was particularly a concern raised in the Faculties of Environment, Science and Economy, as well as Health and Life Sciences. There was a view within these faculties that certain key competencies do need to be assessed in a more traditional form, where access to ChatGPT (and generative AI more broadly) is removed. The specific concerns of these faculties will be explored more below.

Certain staff members also raised the concern that a number of their assessments are accredited or externally-set, and so a wholesale embrace of generative AI may not be possible in these areas. This was a particular concern in the Faculty of Health and Life Sciences, and in some areas of the Business School.

A particularly popular suggestion was an increased use of presentations and viva-style assessments (*“I think the best way is the oral exam”*). There was also a lot of support for assessments which critique generative AI answers; that is, an answer would be generated for students to analyse and critically assess.

“A critical use is asking students explicitly to use it and then reflecting on the limitations of it and what went wrong in the use of it.”

However, a complete overhaul of all current assessment techniques is not necessarily required. Instead, current methods could be slightly altered, for example asking students to provide a copy of the generative AI prompts they used during the assessment. Perhaps more radically some module leaders may not feel the need to change their current assessments even after considering the capabilities of generative AI.

“[What if] when they submitted their assignment, they had to put in what they asked GPT and then also the what the additional prompts were.”

“What if somebody asked the exact specific questions? Like really good questions from ChatGPT and gets an absolute brilliant answer? Like [another participant] said earlier, what’s wrong with that?”

Indeed, these suggestions, which were made before the university released guidance on how to properly reference generative AI in academic work, are nonetheless in accordance with this guidance (University of Exeter Library, 2023).

During the workshops, staff expressed concerns about maintaining meaningful student engagement in their modules and assessments. They highlighted the importance of fostering an understanding of the value of assignments in students’ academic and personal growth, as well as enhancing their overall learning experience. With the advent of generative AI, there is a growing worry that students might complete their degrees without truly engaging with the course material. This could further exacerbate a perceived trend in student mentality that treats university as a series of checkboxes to be ticked, rather than a learning journey. Therefore, when considering new forms of assessment and teaching practices in general, it is imperative that students can clearly see their purpose in light of an AI-empowered world.

“I suppose the real question is how to make assessment meaningful then ... the writing up of the assessment isn’t the important thing... it’s actually what they did, what they learned from doing the assessment rather than just the writing up of it.”

“My real concern is the fact that we might end up with lots and lots of students coming out and barely engaged because they haven’t actually intellectually grappled with anything.”

However, as mentioned above, there are both staff and students who will choose to resist the use of specific generative AI tools, and so any changes to assessments to incorporate generative AI should be mindful of this. That is, no assignment should require the use of a specific generative AI tool to complete.

There is also particular staff concern about the speed at which generative AI is developing, and the comparative inertia of module and assessment changes. ChatGPT-3.5 was released on the 30th of November 2022, with ChatGPT-4 rolling out in March of 2023, and the plugins and web-browsing capacity of ChatGPT-4 in May of this year. Each successive improvement to the ChatGPT models has revolutionised certain use cases and areas, as the technology becomes increasingly capable and more accurate. Given this context, staff are concerned that the current exception-based system around in-year module changes will be insufficient, if the quality of education and qualification wish to be preserved. The university must therefore offer staff flexibility, recognising that this entire situation is unprecedented.

We must also raise that all these assessment changes will have significant implications for staff workloads, as we have discussed above.

Clear Policy

Implement a clear, user-friendly policy that undergoes regular reviews to stay in step with advancements in generative AI technology.

There is a strong desire for any resulting policy to be clear and specific as to which actions are and are not allowed. There is a large concern among students that, without this kind of clear-cut policy in place, students will be unfairly penalised for ChatGPT use (and generative AI more broadly) that they were unaware was not allowed.

“I don’t want to end up in a situation where I’ve, for example, used it to get feedback on an original draft that I’ve written... only to get kicked off of my course because I’ve not used it in an appropriate way that hasn’t been clear to me.”

For example, there was much discussion in the student workshops about the difference between using ChatGPT (and generative AI more broadly) for research help, or for feedback on an essay, compared to generating written work wholesale from ChatGPT (and generative AI more broadly).

Faculty Specific Use Cases

Throughout the course of our workshops, we identified a number of use cases that were common across all staff, all students, and even of both groups. A selection of these were identified above. However, it is also important to note the areas where use cases and concerns are specific to certain disciplines and faculties. With that in mind, we now turn our attention to the specific faculties and the areas of concern that they need to tackle directly.

It is important to stress that a trend throughout our workshops was that some staff expressed a view that their subject area was (perhaps uniquely) immune to the threats that were identified in the rise of generative AI.

“I’ve tried it several times but all the code it suggested never works out.”
(Business School)

“So far, I’ve tried it to solve a few prompts with GPT-3, and it just flat out fails...”
(FESE)

“For my field, it’s not so much as a problem.” (HASS)

“At this point, it doesn’t quite have the capability to really be able to analyse and synergise lots of information and come up with higher-level understanding of what that information means.” (HLS)

However, staff were always able to identify possible risks for other subjects and areas. While it is the case that specialists will be able to have a better understanding of their own area and thus whether there are certain ‘AI-proof’ elements to it, it is important that we take care to avoid thinking that any area is immune to this challenge.

During the workshops, numerous staff members discovered that collaborative discussions about the challenges and opportunities posed by generative AI sparked more ideas than when they contemplated these issues individually. In fact, some workshop groups decided to maintain contact and form dedicated groups for ongoing idea exchange. Therefore, faculties should consider

establishing internal discussion groups and meetings dedicated to the exploration of generative AI-related challenges and potential solutions.

To facilitate further discussion among members of the university community, we have included the proposed new forms of assessment from all faculties in the appendix 'Possible Assessment Changes'.

Faculty of Environment, Science and Economy

During the workshops for the Faculty of Environment, Science and Economy, possible use cases often surrounded the interconnectedness of disciplines and subjects within this faculty. Given the considerable overlap among various scientific disciplines, it was suggested that ChatGPT, and generative AI more broadly, could be employed to help students or specialists in one area broaden their understanding of a problem's full scope. To support this use case, it is important that educators emphasise the importance of teaching students to validate the information they receive, indicating a shift towards fostering critical thinking and digital literacy.

"I think it really important that we teach our students how to use it now."

Similarly, students raised the possibility of using ChatGPT to analyse and assess multiple different methods for solving the same problems more efficiently and quickly.

This is also (perhaps unsurprisingly) the faculty that had the greatest focus on the use of ChatGPT for coding help. Participants suggested both that the rise of generative AI would make the need for coding skills less relevant, since tools like ChatGPT can produce code when asked, but also that a firm understanding of the principles underpinning programming will become more important, in order to optimise or correct AI-generated code. Furthermore, participants suggested that ChatGPT could be used as part of the learning process for coding due to its unique and time saving abilities.

"It is some kind of learning ... There was a library which I didn't know about, but I learned to use it with ChatGPT If I wanted to do it with Google it may take days but with the ChatGPT three or five hours."

There was also a view that, given ChatGPT's coding abilities, it will become more important for students to interpret statistical results in a meaningful way, rather than needing to understand how to produce them in the first place.

"They [students] could probably only benefit from ChatGPT, in terms of help with coding, and then they can focus more on understanding the ideas."

We also noticed that this was the faculty with the least concerns about their assessments, perhaps because many subjects within FESE have already returned to in-person invigilated exams after the Covid-19 pandemic. However, it is interesting to note that many staff within this faculty believed that invigilated exams were one of the only methods able to accurately assess students in light of generative AI. Given these considerations, then, much of the discussions in the FESE workshops surrounded using generative AI to write code, as this is the main area that is rarely (if ever) assessed in invigilated exams in this faculty. Staff were generally optimistic about generative AI use in this area. Many staff expressed that they felt that using ChatGPT to generate code is different than using it to develop, for example, an essay. That is, staff already expect students to turn to the internet for help with programming assignments, and so ChatGPT merely represents one more tool in their arsenal.

“If they use AI to generate [code] and it works, I think that’s absolutely fine... I already see lots of copy-pasting from Stack Overflow etc... and you can still see if [students] understand what they’re doing, so I think I might not change anything.”

Business School

Staff members in the Business School seemed to hold some of the strongest views of any faculty, whether that be a desire to embrace this new technology, or a wish to prevent its use entirely as it is only a tool to “cheat.” Given this stark divide, we would recommend that the Business School in particular (and the Faculty of Environment, Science and Economy more generally) be prepared for the likelihood of there being resisters to this technology, as we outlined above.

Those members of the Business School that encouraged embracing the use of ChatGPT acknowledged that traditional forms of assessments would need to be re-evaluated due to AI’s ability to generate plausible answers. There was an emphasis on shifting towards assessments that measure the application of knowledge, competence, and critical thinking, rather than rote memorization or the ability to generate well-written essays. Moreover, the potential of AI to transform the workplace and the resulting shift in skills requirements were acknowledged. Skills such as critical thinking, creativity, and the ability to work effectively with AI were highlighted as increasingly important and therefore essential to assess through new assessments.

“What this is [going to] make us do is actually teaching them competence and skills rather than knowledge. And I completely agree.”

“The use of AI and AI generated data is going to be the future of the world of work. Then maybe that is one of the key skills we require of a management student.”

The suggestions for new assessment methods were wide ranging, such as adding reflective assessments like an individual learning log or assessing students on a portfolio of work. Other ideas revolved around students submitting their ChatGPT prompts alongside their final piece of work. This could help alleviate student concerns about what uses of ChatGPT are allowed as highlighted earlier in the report and promote critical and thoughtful use of generative AI tools.

“[What if] when they submitted their assignment, they had to put in what they asked GPT and then also the what the additional prompts were.”

There were also many discussions about the possibility of oral assessment formats, such as interviews or vivas. However, as when this came up in the other faculties, there were significant concerns about staff workload and the realities of marking these assessments.

To support any changes which incorporate AI use into the curriculum and assessments, staff highlighted the need for training to understand and adapt to these new tools. They also recognised the importance of teaching students to use AI effectively while understanding its limitations and verifying outputs. Trust in AI systems was a major concern of some members within the Business School.

However, as we have mentioned, there was also a marked disparity between those in the Business School who wished to incorporate ChatGPT and generative AI into their courses, and those who felt

any usage constituted unethical practice or academic misconduct, or that students would only be using AI for those reasons.

“I think it is more negative than positive... Just adopt exam format as assessment. Coursework would just be easily defeated... To make sure that students actually learn something and don’t just cheat their way out.”

Given this view among staff in the Business School, it is important to ensure that provisions for invigilated forms of assessment are prepared and available for the coming academic year. From what we heard in workshops, it seems likely that there are staff members in this faculty who would wish to return to traditional exam formats to ensure that the key competencies and knowledge they wish to assess can be done so in a way they are most comfortable with.

Faculty of Health and Life Sciences

A point raised in the HLS workshops was the problem of accredited assessments and examinations; certain courses and modules are required to assess students in a very particular way that is outside the control of the university.

“Our accrediting bodies specify how some of our assessment must look (e.g. a theoretical essay) so we may come up with creative ideas to work with generative AI, but these might not fulfil what our accrediting bodies expect of us and there will be a delay in them catching up!”

This is a particular example of the necessity of the university allowing greater flexibility for in-year module and assessment changes than it has in the past. If these accrediting bodies are delayed in publishing what they believe is ‘best practice’ in light of generative AI, staff should be able to change their own modules as soon as these bodies reach a new consensus, which may well require a great deal of flexibility within the academic year.

There was also much discussion in this faculty about the necessity of students memorising the information they are given, as a medical professional will need that knowledge immediately to hand. As such, there was support within this faculty to ensure that this foundational knowledge is adequately assessed through traditional invigilated assessments.

“For an F1 Doctor, if you’re training a medic, they really need to have certain things at the tip of their fingers, so that they’re using their heuristics to diagnose somebody, or to act in a very stressful situation. You need them to be tested on their knowledge.”

Faculty of Humanities, Arts and Social Sciences

As this faculty is likely to have the most essay-based assessments, it is unsurprising that it was often these subjects that became the most obvious areas of concern throughout workshops. The potentially dubious validity of essay-based assessment in light of technology that can be used to write entire essays is a deeply concerning one. Indeed, staff members in Student Cases and Academic Conduct offices who attended workshops expressed that they had already seen cases where students wrote entire assessments with ChatGPT.

However, despite this risk, it is important to stress again that the students we spoke to expressed many use cases for ChatGPT, none of which were asking it to complete their assignments for them. Instead, these students focused on using it to help understand and revise the necessary content

through summarisation and tutoring, and to improve their existing work through grammar and proof-reading checks.

In addition, some staff members expressed that their discipline and their existing assessments were already structured such that ChatGPT (and generative AI more broadly) is only mildly useful, if at all. In particular, staff emphasised the role that critical thinking and creativity played in a number of their assessments, with ChatGPT (and generative AI more broadly) currently limited in its ability to replicate these qualities. Staff also expressed that ChatGPT (and generative AI more broadly) may be a great benefit here – if it can be used to provide the factual information, then more teaching and writing time can be dedicated towards the analysis of theories. However, it should be highlighted to students that generative AI is susceptible to hallucinations.

“For my field, [the immediate access of information] is not so much of a problem... For me, given that they can access very easily... say, what theories of emotion there are, I can tell them ‘Look you can’t get [that] wrong in your essays.’”

A unique use case discussed within the faculty was generative AI’s use in education, such as to aid the process of creating lesson plans, assessments, and disclaimers or content warnings.

“I’ve also used ChatGPT to figure out certain things like ‘how do I approach certain topics with certain individuals’. Especially when teaching, it’s important to talk about disclaimers, so what are the easiest way to provide disclaimers and so on?”

“It’s very specific to the teaching profession but it creates entire lesson plans and assessments.”

While there were concerns about a potential increase in educators’ workload due to changes in assessment formats, faculty members recognised AI tools as beneficial for planning and structuring assignments, as well as identifying student cohorts’ areas of weakness. This faculty was particularly progressive in its outlook, embracing the potential of generative AI to foster critical student engagement and enhance their synthesis skills.

There was strong support for restructuring assignments to mirror real-world situations where AI tools are common, and the potential for personalizing assignments was also discussed. Yet, it was acknowledged that it is essential to educate students about the limitations and inherent biases of generative AI. Equipping them with critical literacy skills is crucial to prevent over-reliance on AI, which could hinder the development of important academic skills.

“So perhaps what we need is tasks that are sufficiently real world like that it’s fine to use ChatGPT as a support tool when you’re doing that, because that’s what you’d be doing in the real world.”

Conclusion

Throughout our workshops, we have seen a wide variety of possible ways that ChatGPT (and generative AI more broadly) can be incorporated into higher education, both for staff and students. While staff were often very concerned that students would be using it as a simple ‘plagiarism-machine’, the discussions had with students suggested that this was far from the case. Instead,

students were using it in far more sophisticated ways in their academic and professional work, most often letting ChatGPT (and generative AI more broadly) take the role of a research assistant for the student's work and a final proof-reader.

It is important that we recognise, however, that these results are inherently limited: student attendance to these workshops was much lower than staff attendance, and it is possible that those students who do use ChatGPT (and generative AI) to answer their assignments for them did not attend our workshops, and that if those who attended did use it for less ethical purposes, they may have been reluctant to discuss them. Both staff and students engaging with the project are necessarily self-selecting and ultimately the participants represent a small percentage of the overall student and staff population.

Moreover, the underpinning issues, including the technology itself, can be polarising and divisive and, as such, a consensus does not exist. We have tried to highlight these divergences, whilst also striving to identify general trends and themes. Related to this are risks of confirmation bias, e.g. respondents argued both that ChatGPT made exams essential and irrelevant (and both could be true in context).

The effective and ethical use of generative AI is a complex topic and rising to the challenges it poses will require good leadership, community and collegiality (including co-created responses), bravery to challenge ourselves and re-conceptualise what we do and why, and many other attributes we champion as an institution.

In summary, we recommend that the University and its constituent Faculties: increase provision for staff and student training on generative AI; seek to find a viable method to ensure that students have equal access to a baseline of generative AI; plan for the likelihood of an increase in invigilated forms of assessment, while also giving staff the time and space to reflect on their assessments and to design forms of assessment that they are comfortable giving their students; and to create a clear and specific policy to avoid confusion and accidental academic misconduct.

Appendices

Defining a Baseline of Generative AI

As we have highlighted in this report, ChatGPT-4 offers significant advantages to users over the free to access ChatGPT-3.5 and other generative AI models. Indeed, ChatGPT-4 is shown to “outperform existing LMs on all benchmarks” (OpenAI, 2023, p. 7). However, we also highlighted how the additional financial burden would be inaccessible to university members already struggling financially.

The ideal situation in the current landscape of generative AI tools would then be to give all university members access to ChatGPT-4. This could be achieved by holding discussions with OpenAI about providing university wide access. Their website states that “we are actively exploring options for lower-cost plans, business plans, and data packs for more availability” (OpenAI, 2023) so this represents a possible avenue for investigation. It could also be achieved by working directly with the OpenAI API and building a university wide generative AI service; however, this could also have significant financial implications to set-up, run and maintain. There could also be problems surrounding rate limiting on the API making the service slow and hard to use. Grants or other funding sources could also be considered that could subsidise the cost of access for university members.

However, considering the rapid advancement in generative AI technology, future tools surpassing ChatGPT-4 are inevitable. These future tools might not be developed by OpenAI and could be accessible only through a paywall or subscription service. Therefore, it is likely not feasible for the university to continuously provide access to the current top-tier generative AI technology for all university members. Additionally, the leading tool in one subject area may not be as effective in another, adding further complexity to the situation. Moreover, it is possible that in the future, all generative AI tools, including those currently free to access like ChatGPT-3.5, might only be available via a paywall.

Nevertheless, given the transformative effects already demonstrated by generative AI, it is critical that university members can freely interact with and utilise this technology without financial hindrance. This access is vital to adequately prepare them for the workforce and to maximise their productivity and quality of work.

These complexities must be taken into account when determining the baseline set of generative AI tools that should be freely and equally accessible. This collection of tools should be usable on all university members’ personal devices; no one should be required to use an on-site university computer for access, ensuring fair and equal availability. Our recommendation is that these baseline tools should offer performance at least on par with the currently freely available tools, and efforts should be made to secure free access to ChatGPT-4. Furthermore, given the rapid rate of AI development, the definition of a ‘baseline’ tool should be regularly reassessed.

In addition to the established baseline of freely accessible generative AI tools, the university should also consider offering all members opportunities to access more premium tools, such as the current ChatGPT-4. This could involve subscribing to paywalled generative AI tools, which could then be made available via on-campus university computers, for example. The access to these premium tools could be regulated in various fair ways, such as booking slots to use a computer with a premium generative AI tool enabled, similar to the existing model for study rooms. Providing this access would be extremely beneficial for students and would also enable the university to monitor how AI tools are used, ensuring their use aligns with university policy.

Possible Assessment Changes

Throughout the workshops, a number of possible assessment formats and provisions were raised by both staff and students. Many of these were only mentioned by members of specific faculties, and so have been included within the faculty-specific sections of this report. So that all these suggestions can be considered by the whole university community, we therefore include them in brief here. We hope that by doing so they can inspire discussion and innovation across the University.

1. **Invigilated forms of assessment**, including traditional in-person assessments, presentations, or interview or viva-style assessments. The latter suggestions would allow staff to interrogate students on their understanding of a topic, thereby ensuring that the use of generative AI in creating the presentation will not be an unfair advantage.
2. **Critical analysis of AI-generated answers**. This form of assessment could require the staff member to generate a generative AI answer to a subject or question beforehand, for the students to then analyse and critique. Students could also generate the answer themselves however there needs to be consideration students may have ethical or other objections to using specific generative AI tools.
3. **Individual, reflective learning logs**. This may include reflections on the development of students' understanding and opinions, or a critical analysis of specific discussions they had or contributions they made during classes.
4. **Portfolio assessments**. Given the rather generic and uncontroversial tone that appears in AI-generated text, many staff foresaw that an individual voice and writing style would become more important for students to develop. An assessed portfolio was therefore suggested to evaluate a student's development and continuity through their assignments.
5. **Inclusion of all AI-generated content** (such as generated images or chatlogs) to be included with an acknowledgement of the use of generative AI in assignments. This is the current University suggestion, with a guide being produced by the library describing how to accurately and ethically acknowledge AI usage (University of Exeter Library, 2023). This would promote critical and thoughtful use of AI.
6. **Group projects**. An increase in the amount of group work was proposed throughout the workshops. It was suggested that getting students to collaborate and work together required them to implement skills that could not be replicated by AI, namely organising, assigning responsibilities, and working collaboratively.

Novel Use Cases


During the workshops, several novel use cases of generative AI were brought up. These are listed below in the hope that they may be able to help you improve your productivity and assist your academic work.

1. **Generating multiple choice questions.** Generative AI can produce multiple choice questions based on a broad range of subjects. The staff who have used this have found it very useful, stating that, *"It makes me massively more productive in that area"* and *"It's given a really spot on question, and these questions ... take us about half an hour to an hour to write one question"*. The questions it produces would obviously need to be fact-checked and verified, however this would still constitute a large time saving in this area.
2. **Various teaching assistance methods,** including checking lecture or seminar plans to make sure no content is missing, using it to write entire lesson plans, and creating model answers, especially to new questions.
3. **Language learning tool.** Participants stated that it was a very impressive language learning tool for all levels of knowledge, from beginner to expert.
4. **Interview preparation.** Workshop participants described leveraging generative AI to aid in interview preparations. They would input job descriptions, company information, and personal applications into the AI system, prompting it to generate potential interview questions. Furthermore, they could then feed their answers back into the AI to receive suggestions for improvement.
5. **Tweet writing to summarise work,** *"I've seen colleagues using it for writing tweets, like short summaries of messages they wanted to prepare, and they were happy with the outcomes"*. This could be used to increase social media presence of members of the university and their work, generating more engagement for the university and its ongoing research.
6. **Checking work against marking criteria.** Students described employing generative AI to evaluate their completed assignments against the marking criteria. This process helped them ensure they had addressed all relevant points and demonstrated the skills targeted in the assessment. They could then utilize this feedback to enhance their assignments.
7. **Promote student engagement.** Generative AI can provide answers to questions in seminars to get a discussion started. *"Everybody comes offering up the answer that it's given. And therefore, you're not offering up your own opinion, you're just offering up a different opinion that can then lead to a discussion about the subject. And then nobody needs to be worried that people are judging them."*
8. **Translation between programming languages:** *"one of my team has been using it to translate between programming languages. ... So, she found a tool written in a programming language she wasn't familiar with, but she wanted to use the function. So, she got chat GPT to rewrite it into a program that she did know, and then she could incorporate it into her work."*
9. **Assisting visually impaired students.** One visually impaired student found generative AI can be very useful for describing visuals and figures in research. Moreover, they found that it can produce an answer far faster and more conveniently than traditional methods. They use a separate tool for describing the figure, and then copy the description into ChatGPT for explanation.

Survey responses

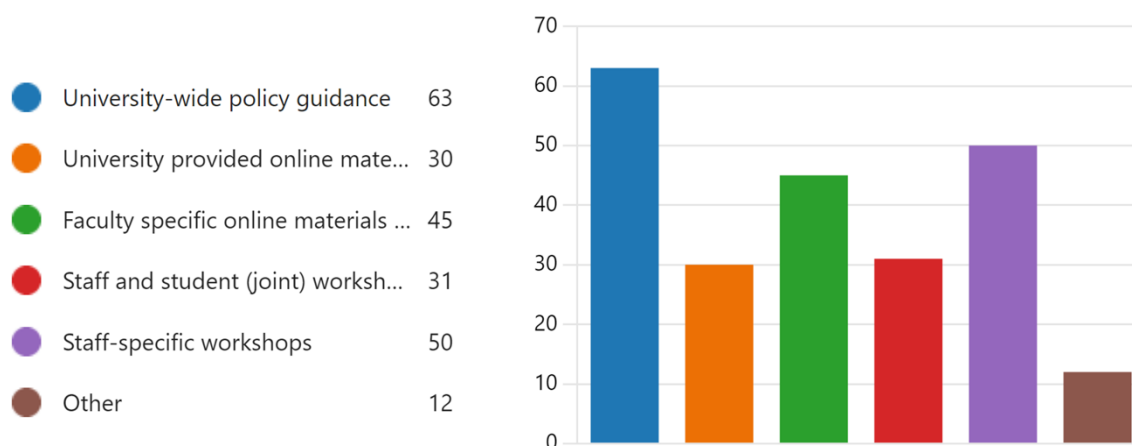
To receive further feedback regarding use cases and concerns highlighted during the workshops, we released a survey to gather the views of the wider university community. At the time of writing the survey has received 81 staff and 76 student responses. Below is a summary and brief discussion of the responses to selected questions.

How knowledgeable do you feel about how you can use generative AI in HE/your study?




Of particular note is that more than half of staff respondents feel that they are “Not knowledgeable” or “Slightly knowledgeable” about the uses of generative AI, compared with one-third of students. In both cases, these contingents are large and will require significant support if they are to use generative AI effectively.

What support would you like to receive around generative AI in HE (staff)



What support would you like to receive around generative AI in HE (students)




Both staff and students want to see university-wide policy guidance as a priority, as well as faculty-specific online materials and guidance, and support targeted specifically at their group. Staff have a preference for workshops, whilst students indicate that they would prefer online materials.

Non-accredited modules addressing generative AI are significantly further down the list of support that students would like to see. Interpreting this in the light of the workshop discussions is difficult and further work might need to be done to understand whether such modules would need to be credit-bearing, faculty/department-specific, integrated into existing modules and programmes, or if the idea was popular with those staff and students who engaged in the workshops, but has little traction in the wider student community.


Suggested uses of ChatGPT in HE (staff)

Strongly disagree Disagree Neutral Agree Strongly agree



Suggested uses of ChatGPT in HE (students)


Strongly disagree Disagree Neutral Agree Strongly agree



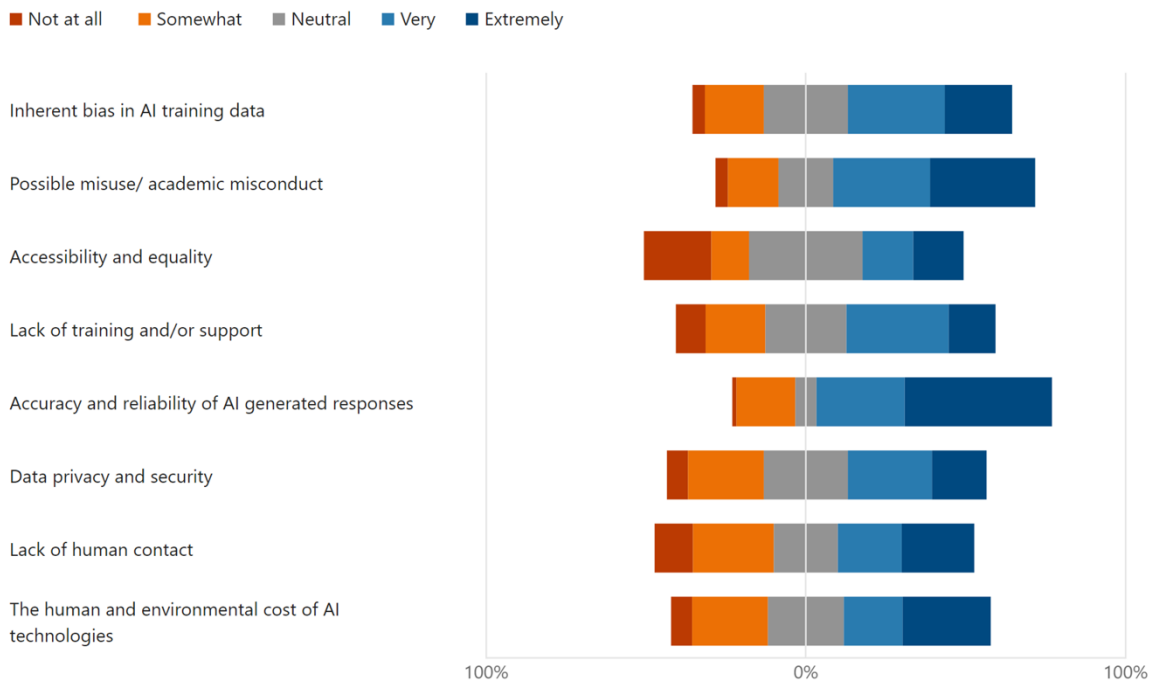
Of particular note here are the discrepancies between staff and student responses. Students are significantly more positive than staff about the potential uses of ChatGPT for general support throughout their studies. For all suggestions, there is a wide spectrum of views indicating the polarising nature of such potential use cases.

Possible concerns around ChatGPT (staff)

Not at all Somewhat Neutral Very Extremely



Possible concerns around ChatGPT (students)



Again, we see significant differences between the student and staff responses regarding concerns over the use of generative AI. Staff indicate more concern about each of the indicated issues than students. The major concerns of both groups are “Accuracy and reliability of AI generated responses” and “Possible misuse/academic misconduct”. A potential concern is that students tended not to feel that “Lack of human conduct” is a concern (although there is a full spectrum of views), whereas staff are very concerned about this. Particular care may be needed to address these divergent views, especially in the light of changing engagement patterns and expectations following the pandemic.

Use of ChatGPT for specific course-related tasks

We asked participants for their views as to whether ChatGPT could assist with the following tasks:

- How useful do you think ChatGPT can be as a virtual assistant, providing responses to queries about course materials, assignments and deadlines?
- How useful do you think ChatGPT can be for personalised learning support, providing explanations, clarifications and resources?
- How useful do you think ChatGPT can be for feedback on your writing and proofreading?
- How useful do you think ChatGPT can be in facilitating and assisting group work?
- How useful do you think ChatGPT can be for exam practice?
- How useful do you think ChatGPT can be for module selection and academic advice?
- How useful do you think ChatGPT can be for summarising literature?
- How useful do you think ChatGPT can be for marking and providing feedback on assignments?

The responses to all these suggestions were overwhelmingly negative from both staff and students, with the exception of “summarising literature”, which received an almost equal split across positive, neutral and negative responses from students.

What other concerns or limitations do you foresee in the use of ChatGPT for higher education?

We gave survey respondents an opportunity to provide free text responses to the above question around concerns and limitations. Some notable comments are quoted below and illustrate the divergence of views and highlight the challenge institutions face when considering these technologies. The comments are sorted into similar themes to help the reader, but many comments address multiple issues.

Please note, due to the nature of the question posed, the responses provided largely express scepticism towards the use of generative AI in higher education. While these concerns are indeed legitimate and demand attention, we should remember that the overarching sentiment expressed in the workshops was generally positive and receptive to generative AI.

(In)equality, access and ethics

Many responses received to the surveys expressed ethical concern around the problems of inequality and accessibility that are inherently present in the training and use of generative AI. As we mentioned above, some members of the university community may have ethical concerns with utilising specific generative AI tools, due to reasons such as the environmental costs of training and maintaining this technology, the working conditions of those involved in the moderation and training of these models, and the inherent lack of representation in training data scraped from the internet (which as a result over-represents generally affluent, generally Western voices).

“My main concern is that AI entrenches already existing inequalities and this needs to be absolutely fought against by not [sic] making critical questioning and alternate source use a key part of the application of AI whether in teaching, assessment or research.”

(Staff, HASS)

These ethical comments and concerns broadly covered the following key areas. We have included some of the comments verbatim to illustrate the spectrum of views offered, but these are by no means exhaustive.

1. **Inequalities and Bias:** Many expressed a fear that AI could reinforce existing inequalities and biases, due to reliance on Western and colonised sources. This was particularly related to university efforts to decolonise curricula, which may be hampered by an embrace of technology which is so heavily trained on Western data sources.

“Decolonise the curriculum – can this ever align with those efforts?”

(Staff, ESE)

“I’d like to think in these discussions, we talk about how to utilize AI as a tool for the benefit of both staff and students in ways that promote what the institution stands for (its core values) while not losing sight of them in a techno-chauvinist

haze. [...] It is also worrying that AI draws from heavily Western and colonized literature and history.”

(Staff, HASS)

2. **Ethics and marketing:** There were concerns about the limitations of ChatGPT, how it is being marketed, and its non-compliance with GDPR.

“[...] There’s so little recognition of how it’s [ChatGPT] merely an algorithmic relation between words and cannot provide factually correct information, and how it’s built upon exploited human labour (never mind the environmental cost). I’m disgusted with how little recognition there is of this, and how much the university have gobbled up the marketing hype around this tool (which by the way, isn’t even GDPR compliant!). Great way to dodge addressing the actual problem of university marketisation!”

(Staff, ESE)

3. **Legal Framework:** Respondents expressed a need to wait for a global legal framework to regulate AI before embracing it fully in higher education context. While efforts to do so are underway in the form of the EU’s AI Act, these frameworks are far from complete.

“There needs to be a global legal framework around AI before it is accepted in HE.”

(Staff, HASS)

4. **Reputation and IP Risks:** There were concerns from both staff and students over what impacts AI usage may have on authorship rights, intellectual property, and academic reputation.

“Loss of IP [Intellectual property] – as an academic, my value is in the sale of my IP. While I’d be happy for some of this to be disrupted (e.g., publishing industry) there is a risk of other activities (e.g., consultancy) being undermined and de-valued.

“Misapplication of knowledge – information stripped of context and the original intentions, may be used to misrepresent the authors intent, or result in dangerous/damaging decisions.

“Reputational damage – e.g., my department has lots of links with people working in creative industries that are threatened by AI, would they still work with us if we were seen as endorsing it?”

(Staff, ESE)

“[...] ethical concerns about where they are getting data from (scraping people’s work ... without permission) and how it’s getting cleaned (unethical use of underpaid workers who are not trained properly/forced to see horrible things to clean dataset). Finally concerns about how my data is being used – i.e. if I entered a report or essay for feedback, ChatGPT now had my report which can be used to

further train my model and of relative writing in my voice. And if they then choose to paywall that..."

(Student, ESE)

5. **Job Loss:** Anxiety about AI replacing jobs, particularly among administrative staff, as well as pastoral care and wellbeing staff.

"...I am also concerned about ChatGPT being used as an excuse to make certain job roles redundant: it is absolutely not a good substitute for pastoral care, academic guidance, wellbeing support, social interaction, etc."

(Staff, HASS)

"Loss of staff – lots of these questions seem centred around the potential to replace administrative staff, which is something we've seen elsewhere."

(Staff, ESE)

Skills and Learning

Significant concerns were also raised surrounding skills and learning. Respondents to the survey expressed fears that generative AI may foster a culture where only answers matter, not the learning process and lead to a reduction in the level of skills held by university members. In particular, a loss of independent thinking was highlighted as a key concern along with fears that generative AI would hinder students' ability to critically engage with material.

A summary of the opinions expressed are provided below:

1. **Promotion of Critical Thinking:** AI should not just provide answers but encourage critical thinking. The concern raised was that a failure to promote critical thinking with AI use may foster a culture where only answers matter, not the learning process.

"It risks removing the learning process - and encourages a culture of 'answers are all that matters'. This is a university – students need to learn to think independently, and look to pursue questions for which there are not yet 'answers'. ChatGPT will further reinforce a growing culture of 'will this be in the exam?' and students who can secure 1st class degrees but not hold a conversation about their field"

(Staff, HLS)

2. **Risks to Independent Thinking:** Concerns were raised that ChatGPT might discourage independent thought, reinforce a culture of exam-focused learning, and hinder students' ability to discuss their field critically. More generally, there were also concerns that the use of AI could shift students towards passive consumption of knowledge rather than active engagement and critical thinking.

“The utilization of ChatGPT raises concerns about the long-term impact on human involvement in knowledge production. With individuals relying on language generation algorithms, there is a potential shift towards passive consumption of knowledge rather than active engagement in its creation. This shift arises from the algorithms’ reliance on existing knowledge, essentially recycling and regurgitating information without fostering original thought or critical thinking.”

(Student, HASS)

- 3. Issues with False Citations:** There are worries about the risk of non-existent citations in works generated by ChatGPT, exacerbating the ‘fake news’ problem and potentially misleading students.

“I am concerned about generative AI as there are numerous documented examples of non-existent publications cited in essays created by ChatGPT. We live in an era of fake news as it is, so students potentially referencing or trying to access non-existent scholarly material is an even more serious issue. AI could be a very useful learning tool for students, but the fact that there is little awareness amongst students of any of the context or issues surrounding AI or the ethics of its use means that we have a long way to go within the University before it can be successfully adopted as a matter of course.”

(Staff, Education and Academic Services)

- 4. Erosion of Academic Skills and the Risk of Overuse and Replacement:** Overreliance on AI could lead to the erosion of critical academic skills, such as research, writing, and creative thinking, making academic degrees less meaningful. Participants also voiced concerns about overusing AI to replace human roles such as marking assessments, cautioning against using ChatGPT as a substitute for reading research articles or other academic activities.

“Erosion of human critical thinking, writing and research skills. If students use AI to summarise their readings or literature, generate their essays and provide feedback on their writing, they will become lazy and not develop the skills that have defined academic learning and research for generations: creative thinking, intelligent connections, good writing, etc. There will be no individual expression, no original thought, instead everyone will sound the same and degrees will become meaningless because they have been earned by a computer, not the mind and creative imagination of a thinking and skilled human.”

(Student, HASS)

Assessment

The concerns raised about assessments in response to the survey question mirror many points discussed in the workshops. There is a view that assessments will need to be reconsidered due to the rise of generative AI technologies. Conflicts around the role of invigilated exams were again present

within the responses as well as suggestions that other assessment methods will need to be modified. A summary of the responses to the survey question under this theme are detailed below:

1. **Reassessment of Exams:** Some members advocate for invigilated exams, viewing them as fair, cost-effective, and less prone to academic misconduct. However, others raised contrasting views that there was a need to modify exams to reflect the integration of tools like ChatGPT, and to better prepare students for real-world applications of such tools.

“Not directly a concern, but I hope consideration of AI kills once and for all the idea that there is something wrong with invigilated exams. Exams are not the perfect assessment tool, but they are the best we have. Exams are fair, reasonably cost effective, force the students to think and they are (almost) immune to academic misconduct. “

(Staff, ESE)

“[T]he use of ChatGPT may be considered cheating but if the question/task is completed correctly with ChatGPT maybe it is better to modify the way we do exams. For example, with [subject] exams maybe it can be in-person but everyone has a computer with ChatGPT and are asked to answer more questions and put it into context for given scenarios. Although very complex, it is better than ignoring ChatGPT entirely when it is very likely it will just be used by the [subject] students in their future employment, one person I know said it is similar to asking mathematicians to complete a test without a calculator when in real life they will always have a calculator, so is what they are being tested for actually necessary for today’s work environment. Therefore, I worry exams being put back in person with no access to computers may simply not reflect someone’s true capabilities in a real working environment which I believe may have more value.”

(Student, HLS)

2. **Alignment and Adaptation:** Concerns were raised about the time needed for faculty to adapt to AI integration, critically assess their own assessments, and ensure alignment with key skills and knowledge.

“Time is needed to understand constructive alignment, the assignments need to clearly measure KSBs that are applicable to the module/ subject area. I have [a] lack of confidence that all faculty will take time to critically assess how they assess. This is a game changer for HE and time is needed to assimilate and adapt. Time is a commodity that is lacking in an ever expanding university.”

(Staff, ESE)

3. **Reimagining Assessment:** Respondents expressed that the opportunity presented by AI should encourage a shift towards more creative and process-oriented assessments. Which

better evaluate students' thinking skills, critical reasoning, and knowledge, as conventional methods may be undermined by AI tools.

*“Not taking the opportunity it presents to get people to reimagine assessment
- to be more process-oriented, more creative and imaginative.”*

(Staff, Education and Academic Services)

*“I strongly believe that the university may need to rethink the testing methods.
And find ways to test thinking skills, critical reasoning and knowledge of students.
Essays and written exams can no longer be trusted.”*

(Student, Business School)

4. **Plagiarism Detection:** There is a worry about detecting when students use AI, like ChatGPT, for their coursework, specifically when students alter AI-generated code or content to avoid direct copying.

*“Whether or not there will be any way for detection if students use it for their
coursework (e.g. they use ChatGPT to create code instead of doing it themselves
and then they change certain bits to ensure that it is not directly copied and
pasted)”*

(Student, ESE)

5. **Academic Conduct:** Concerns are raised about students using AI to produce essays, potentially leading to a disparity in effort and performance between those who use AI and those who don't. This underscores, once more, the necessity for a well-defined policy that ensures no one refrains from using generative AI tools, which could inadvertently lead to a disparity in effort.

*“Cheating. I have witnessed people just copy an essay question into the box and
base their entire argument from what the AI outputs. People sometimes do not
copy it word for word or anything which makes it more difficult to detect, but
pretty much all of the key points and pieces of evidence have been pulled from
what the bot has said. This is extremely worrying for me as I really struggle to
formulate essay responses, and I believe that essay structure and evidence choice
is really crucial in writing a good essay, and is what really distinguishes a bad one
from a good one. Therefore, it is worrying that I can spend days and days trying to
plan a good one, when a peer can just press enter in their keyboard with no
thought and output an essay that ultimately ends up being better than mine.”*

(Student, HLS)

Student support and guidance

There was a concern that without integrating AI in education there may be more adverse effects than if it was incorporated, due to the ability to guide student usage and promote critical engagement. However, concerns exist over AI's limited potential to replace human support, given its lack of empathy and the importance of human expertise and intent was raised. Moreover, the necessity for clear guidelines on AI use was raised to prevent unfair advantages and preserve essential skillsets.

1. **AI Integration:** Incorporating AI into educational practices could help guide students' usage, mitigate potential dependency, and avoid adverse effects on student attendance.

"If we do not integrate it into our practice then students will use it without our guidance. We still need to find ways of assessing student knowledge in the context of access to AI. Student attendance may worsen if they think they can find all the answers on something like ChatGPT, which is even more reason to embrace it and fully integrate it so that this doesn't happen."

(Staff, HASS)

2. **Human Support vs AI:** Concerns arise about the potential replacement of human support systems with AI, emphasizing AI's limitations in empathy and the risk of losing the original authorial voice in summaries. Moreover, respondents mentioned the irreplaceable expertise of human specialists, the importance of human interaction, and concerns about the reliability of AI.

"I am extremely concerned that the questions put forward in this survey seem to suggest that the university would use AI chat as a way to circumnavigate around human support systems for students, AI cannot have empathy (yet!) and the suggestion that the university may use it for student support is in my opinion abhorrent. I believe that AI chat is only useful for summarising already existing work, however, even then it partly eliminates the personal written voice of the author, so should only be used with care and understanding and how you can use the AI generated text and reformulate so that it keeps its original voice."

(Staff, Education and Academic Services)

"Please please please DO NOT use AI. Subjects taught at the university go into so much depth and detail that need a SPECIALIST of the subject, not an AI. There is much benefit from interaction with humans – i.e., lecturers, other students, support systems at the uni. There is no guarantee that AI will be a reliable source, and no amount of AI training will be able to match the human mind."

(Student, HLS)

3. **Policy and Guidance:** The main issues include the lack of guidance or policy regarding the use of AI programs and the potential advantages for some students. Also, there is a concern about the replacement of certain skills, such as coding, by AI. Responses also highlighted the need for a strong foundation of knowledge to best use generative AI.

“The current main issues are guidance/policy around using these AI programmes. Some students may already be at an advantage by utilising it where others may be hesitant due to the perception of cheating. Then the replacement of certain skill sets that may be useful – such as coding from scratch, but as a student I’m unsure what in real terms is useful and what is replaced by AI across the board in universities and work places. The current usage for me just acts like an accelerated google search, making getting a result faster, but it still requires understanding and correct application to my assignment.”

(Student, ESE)

Workload and training

Some responses in this area were well discussed in the prior workshops such as addressing the challenges of workload and training in AI implementation, and the urgent need to enhance educators’ knowledge of AI to ensure its effective use in teaching. However some concerns here were more novel, such as concerns about a lack of understanding around the tool’s scope, and the potential impact on career progression for staff who focus time and energy towards the challenges and opportunities presented by generative AI.

1. **Tool Misunderstanding:** Concerns were raised about a lack of understanding about ChatGPT’s scope. The respondent suggested generative AI is often misunderstood and its potential uses are overlooked due to concerns about its faults.

“Lack of understanding of the scope of the tool. People that don’t take/have the time to explore chatgpt misunderstand the purposes, and stop at its faults ... It’s like a knife. People are worried about the stabbings, and slipping and cutting yourself. I’m interested in chopping, slicing, dicing, carving wood, or foam, building, tinkering, shaping....”

(Staff, Education and Academic Services)

2. **Staff Workload and Career Progression:** There are concerns about increased staff workload and potential risks to career progression for those tackling new challenges, akin to the situation during the pandemic-induced transition to online teaching.

“Staff workload and the risk to career progression for those investing their time in tackling new challenges (similar to online teaching during the pandemic).”

(Staff, ESE)

3. **Educator Knowledge and Upskilling:** There’s an urgent need to upskill educators to teach students how to use AI effectively, which could lead to increased workload demands.

“As educators our knowledge is limited in effective use, and we need to be urgently upskilled so as to teach students how to use it effectively. This will have workload issues for already hugely stretched university workforce.”

(Staff, HLS)

Focus groups

Following the workshops, we were keen to explore some of the emerging issues with specific stakeholders who might be able to provide more insight. We held two focus groups: one with seven Senior Academic Conduct Officers to explore their experiences, particularly around misconduct cases; and another focus group with seven international students, including a visually-impaired student, to further understand the benefits of generative AI for these students.

Senior Academic Conduct Officers

The key points that emerged from this focus group are:

1. Generative AI provides a catalyst to rethink what we are assessing (is fluency/coherence redundant?)
2. What is *authorship* and *plagiarism* in respect of these tools?
3. Generative AI is a potential leveller for international students, dyslexic students, and other university members who have difficulty writing and summarising, due to its capabilities in these areas.
4. When is it useful to “*accelerate*” learning and when does learning need to be *slow* to encourage critical thinking, reflection, opinion-forming and other key learning skills?
5. Fewer cases of unethical ChatGPT use have been seen than expected (but they note it is possible we are only detecting the obvious cases).
6. Concerns that if we don’t publish a clear and representative policy regarding the use of generative AI, staff will disengage from formal processes in line with this policy and apply their own sanctions.
7. It is important to address the fundamental challenges generative AI presents – this requires resourcing and a fundamental rethink of what it means to learn, teach and assess.

“Can we have some bravery, please? We are a Russell Group University ... [we need to] make a statement about the impact of this, the complexity and the difficulty of dealing with this. And actually there’s not an awful lot of people in the sector willing to do that because if we start to admit what’s going on we’re scared there’s going to be an avalanche and a crisis, and we’re going to be inviting all kinds of negative press. So we’re ignoring something which is tunnelling into the foundation of the thing that we’re supposed to be doing. If we can’t be brave about this, I don’t really know what we’re doing.”

“[We need to] take this more seriously as the danger to our academic integrity that it actually is. Maybe we’ve had this coming for a long time. Maybe we’ve refused to look this in the eye for a very long time.”

International students

The key points that emerged from this focus group are:

1. International students are using ChatGPT as a tool to help them with grammar, fluency and vocabulary.

“sometimes when I write the emails [to my supervisor], I struggle with how to phrase something that is acceptable. Culturally you know. I ask it literally ‘is it polite to ask this?’”

2. It can be used to help understand course content by reading, or summarising content before a session such as a lecture or seminar.
3. They voiced concerns around what is acceptable use of generative AI, due to conflicting messages, uncertainty around academic conduct, and how ChatGPT works.

“I don’t put like, to be honest, big chunks [to check my own writing] because I’m afraid that it stores the data that I gave because, like that I’m plagiarising, so I’m afraid. So just, only the things that I was really struggling with.”

4. Concerns around the reliability of generative AI due to biases in the training data and cases of hallucination.

One of the participants in this focus group is visually-impaired and gave us express permission to discuss the use of ChatGPT in their context with respect to inclusivity:

1. ChatGPT is an always available, infinitely patient personal assistant
2. ChatGPT allows further questioning and prompting, which helps to explain things in detail and to provide context which is not obvious to explain for abled people. This is not provided by traditional text-speech software or screen readers.
3. Again, they expressed hesitation about using generative AI more broadly, due to uncertainty about acceptable and unacceptable use
4. Ideally, there would be an integration of generative AI to support disabled students.

“It definitely helps to feel much more included, and it is now much easier to have a common frame of reference about things to talk with my friends.”

“ChatGPT can patiently explain what seems very trivial information for others.”

Looking Ahead

Throughout this project, we were mostly focused on the current capabilities and uses of ChatGPT in the context of higher education. Through our workshops, however, both staff and students presented a number of suggestions for where this technology could take us, and ways future developments could improve both teaching and learning at the University.

The university could consider using generative AI to improve the services it delivers to students and staff. For example, by allowing for question answering over university content. The major suggestion in this area was that, in developing a scalable service for university stakeholders, we incorporate this service into the existing University intranet. If it were sufficiently integrated in this way, it could then be used as a tool to query specific lecture notes, module descriptors, or records.

If such a possibility could be brought to bear, this may also alleviate some staff and student concerns regarding AI hallucination. Since such a model would be trained on our internal data, it should produce those answers which are already present in that data, in addition to all its training current data as well. As a result, if a student were using it when studying, they could query the lecture notes that were already available via the intranet. In essence such a tool would simply be a very powerful search function for existing lecture notes. Such a tool would also help staff to walk the fine line between embracing novel technology, while also ensuring students engage meaningfully with the content of a module.

There are major practical and ethical issues around such an implementation, including the extent to which we would want any AI tool and provider to have access to potentially sensitive and commercially valuable institutional data. However, there are significant opportunities to improve the experience of staff and students by integrating AI into our systems, processes and data should this be something we wish to explore as an institution.

Bibliography

Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. Virtual Event Canada: ACM doi:10.1145/3442188.3445922

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., et al. (2020). *Language Models are Few-Shot Learners*. ArXiv doi: arXiv:2005.14165

Common Crawl (2023). *Languages*. Statistics of Common Crawl Monthly Archives. Retrieved 27th June 27, 2023, from <https://commoncrawl.github.io/cc-crawl-statistics/plots/languages>

Education Board (March 16, 2023). *Institutional Response to ChatGPT*. University of Exeter. https://www.exeter.ac.uk/media/universityofexeter/secureall/pdfs/Education_Board-ChatGPT_Statement_Response-Institutional.pdf

Marcus, J. (May 18, 2023). *Professor misuses ChatGPT and fails most of class for AI plagiarism*. The Independent. <https://www.independent.co.uk/news/world/americas/chatgpt-ai-plagiarism-texas-a-and-m-b2341238.html>

NatWest (2022). Student Living Index 2022. *University Living Costs*. Retrieved June 9, 2023, from <https://www.natwest.com/life-moments/students-and-graduates/student-living-index.html>

OpenAI (March 27, 2023). *GPT-4 Technical Report*. arXiv doi: arXiv:2303.08774

OpenAI (February 1, 2023). *Introducing ChatGPT Plus*. <https://openai.com/blog/chatgpt-plus>

Perrigo, B. (January 18, 2023) *Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic*. Time. <https://time.com/6247678/openai-chatgpt-kenya-workers>

Strubell, E., Ganesh, A., and McCallum, A. (2019) Energy and Policy Considerations for Deep Learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics. doi:10.18653/v1/P19-1355.

University of Exeter Library (2023) *LibGuides: Referencing: Using generative Artificial Intelligence (AI) tools such as ChatGPT in academic work*. Retrieved June 19, 2023 from <https://libguides.exeter.ac.uk/referencing/generativeai>