143 - Generative AI - whos the expert - 17 June 2024 small

ROB MOORE: Hello everyone, and welcome to Student Hub Live. You're here with me, Rob Moore, tonight, and we've got a really interesting topic. We're going to be talking about generative AI. Tonight's going to be a little bit different. We want you to really drive the discussion. So we're going to be coming to you for a lot of questions as we go through the evening. And I've got a couple of experts that are going to help us get to the bottom of this whole new topic of artificial intelligence.

We've got Jay and we've got Mary, who are sitting there in the chat box. They'll be looking after you tonight, answering your questions. If you've got some specific questions, pop them in the chat and they'll do their best to answer those and also share some key links with you and some places you can go to get more information. And of course, we've got Heidi with us. And Heidi is going to be representing you tonight. She's going to be talking about the things that you're putting in the chat, the answers that you're giving to the widgets.

So say hello, Heidi. And tell us a bit about how tonight is going to work in terms of the widgets and the chat.

HEIDI: Sure. Well, hello, everyone. Just to say, I'm not an expert when it comes to AI at all. I've started dabbling with AI quite recently. I did a previous show around AI with Rob and found it absolutely fascinating. So I'm really, really looking forward to tonight's show.

So as Rob said, this is a really, really interactive show. So we're going to have a number of widgets on the screen. So we're going to ask you to get involved and to give us your thoughts and your opinions. And then I'll be feeding those back throughout the show.

You might find that the chat goes quite quickly. So on the top-- I think it's that side, if you're looking at it on the screen-- there's a little pin. If you click on that, it will slow down the chat for you. And that's always really helpful for me because the chat does tend to move really quickly.

If you haven't met me before-- I know it's the end of the year, so I probably know lots of you and you've probably been to lots of Student Hub Live sessions before-- but my name is Heidi, as Rob said, and I did all of my undergraduate study with The Open University. I absolutely love the institution. I went off and I did my master's elsewhere and I'm now coming back. So I work at the OU, but I'm also starting a PhD in October.

So I have lots of understanding about what you're going through as an OU student. I know all of the hurdles and that enormous sense of achievement, as well, once you get the TMAs in and the EMAs and you finish the year and you can have a nice holiday and it's all over and you can just pat yourself on the back.

So I want to say a few hellos to those that are joining us so far this evening. We've got tonnes and tonnes of students joining us from Scotland, which is great. We normally get a really good number of students joining us from Scotland. So I'm going to say some hellos. So Jacqueline is in Glasgow, and Jacqueline is studying English literature. That's what I studied, Jacqueline. Good choice.

Dee, I'm near Glasgow and I'm studying a BSC cyber security.

Carol is in the very wet Scottish Borders at the moment. And that's disappointing to hear. It's really sunny here where I am. I'm down near Milton Keynes. I'm near The Open University's headquarters.
So, Fiona, I'm based in Edinburgh and I'm studying English literature and creative writing. And Hannah is in Angus and also studying English. So we've got lots of people studying humanities.
Lorraine is in Inverness and studying accounting and finance. And we've got Alan, from Perth, who is studying cyber security.
Moving down south now to England, we've got Sally, in Canvey Island, Essex. And then, in Northern Ireland, we've got Kieran, who's in Derry. And Kieran is due to start a postgrad in education later this year.
And then for our international students, we've got Leah, who's joining us from sunny Switzerland-- is it Lee or Leah-- I do apologise-- if you could let me know in the chat so that if you comment again I can get your name correct; so I do apologise if I've said that wrong-- studying BSC data science.
And then we've got Sue, who's joining us from Normandy. And finally, we've got Wilfried, who is joining us from Germany this evening. So I'm going to test out a bit of my German. Are you ready, Rob. Guten abend, Wilfried. Wie geht's? There we are. That's about as much as I know. I can also tell you I've got a Guinea pig, but I don't think that's very useful.
ROB MOORE: Absolutely. We might come back and find out about your Guinea pig later. So thank you for that.
So we want to see as many comments from you as possible in the chat box today. Just remember that this is an open forum, so don't share any personal information. Don't put anything in there that you don't want to appear in the public domain. Apart from that, we're really interested in what you've got to say.
So there will be a ticker question coming across the bottom of the screen. And a ticker question is just something that prompts you to put something in the chat. So the first one is, have you used generative AI before? And if so, how? So we want you to put that in the chat.
And there will also be the first two of our widgets. And the widgets today-- the little things you fill in-- they're all Yes or No questions. So the first two are, is using generative AI cheating, yes or no? We're not giving you a middle option. None of that sitting on the fence today. And then, when using generative AI, is it OK to upload some of your module materials and ask it to generate a summary? So we just want you to vote on those and then put some comments in the chat.
And while you're doing that, I'm going to introduce our two guests tonight. We've got Lorna and we've got Mike. So, Lorna, would you like to introduce yourself and tell us a bit about what you do with the OU?
LORNA SIBBETT: Well, I'm Lorna Sibbett, and I'm what's termed as staff tutor in the Open Programme, which means I'm responsible for the student experience in our Open Box modules. But I've recently become chair of the Open Board of Studies, which means I'm responsible for all things quality if you're doing an Open degree or any degree that's within the Open Programme.
For my sins, I'm also an academic conduct officer, which we might touch upon this evening. And in terms of locale, I'm just down the road from Kieran, it seems, because I'm down in mid-Ulster, which is a wee bit south of Derry, where he is. So I'm with Kieran in that Derry versus Belfast sort of fight. If he's wanting to enter into that, he'll understand what it means.
I've lived half my adult life in Scotland as well, so I'm feeling great affinity for the Celtic nations that are well represented here this evening. Pet-wise, there could be two dogs appear at some point, but my husband is looking after them at the minute so it's quiet behind me just now.

ROB MOORE: We always like to see a bit of photo bombing from the pets, absolutely.
Mike, if you'd like to introduce yourself and tell us what you get up to.
MIKE RICHARDS: Hello, Rob. Hello, everyone. I don't have any pets, I'm afraid. So I'm feeling rather left out tonight. There is a LEGO parrot behind me, though. So that's the closest you're going to get.
I'm a senior lecturer in computing, and I'm very glad to hear that some of the people tuning in tonight are doing one or more of our programmes. I've worked mostly on our level 1 curriculum. I've also contributed to our cyber security Programme and postgraduate. So sooner or later, if you're in computing, you're going to encounter some of my stuff. I hope you like it.
In the last year or so, I've been working an awful lot with generative AI, particularly on how it influences how we teach and how you do assessments. So I have thoughts on this issue and I'll be interested to see what you all think about it. So let's have a great conversation. There's a lot to learn on all our parts.
ROB MOORE: Absolutely. Thank you, Mike. And as I say, keep putting your comments in. Anything you'd like to say about the topic, pop it into the chat. Or if you've got any questions as we go through, Heidi is going to be there to bring your questions to the team.
So just before we start looking at your answers-- so to Lorna, when we talk about generative AI, what do we mean by that and what is the OU currently doing about it?
LORNA SIBBETT: I think Mike might be better placed on that-- what do we mean by gen AI point a few-- being the computer scientist. So what do we mean? The use of artificial intelligence is the AI bit of it. The generative bit means it's making something new from it. It's generating. And my understanding-- that Mike can fill in a bit more on-- is that it's generating from being educated through data, initial sources that have been presented to the AI tool. So it's learning from those and moving forward within that.
You see what we're looking at nowadays are these large language models. I know we're not supposed to get technical, but it's just a good, wee piece of guesswork of, you've said such and such, so I think the next word that follows should be this. OK? So it's just building through like that. So generative AI--
ROB MOORE: So should we ask Mike, then? Would you be happier if I asked Mike what we're doing?
LORNA SIBBETT: Yes.
ROB MOORE: So, Mike--
MIKE RICHARDS: Yeah, I think you--
ROB MOORE: --what work is the OU doing--
MIKE RICHARDS: I think you're really unfair.
ROB MOORE: --on these at the moment? I know. But everybody who comes to the show, knows I'm unfair. I don't pretend to be anything else.
MIKE RICHARDS: So yeah, I think Lorna made a really good stab at it. And generative AI is a subset of a much larger area, called artificial intelligence. It's any form of artificial intelligence that generates new data from what we call training sets, which is huge amounts of text or images or music that's been previously given to these computer programmes. They learn, in a sense, what the patterns are in this data and they can create new data.
We all got kind of caught out a little bit when, in late 2022, OpenAI, which is a large corporation in the United States backed by Microsoft, released a Programme called ChatGPT, which basically put a nice, friendly front end on a large language model called GPT. And it blew everyone away because, for the first time, all you needed to do was type something into your computer and you got something back that was more or less like English. And since then there's been an explosion of these chatbots. So if you use

Microsoft Bing to search, there's an option there to involve a Programme called Copilot, which is also included into Microsoft Office.
If you're into graphics, there are programmes such as DALL-E and Midjourney, which can create new pictures just by typing something in. And there are hundreds of other of these programmes out there.
Now, the OU is doing research into using large language models that is also how you can use them safely, but also whether we can actually have some of our own large language models inside the university, which you can use safely on our course materials.
Now, they're not ready for everyday use yet. These are research projects. These are prototypes, and there's lots of questions we haven't answered yet. So it's in the future. Hopefully, you will be able to use a large language model at the Open University to explore your module materials, perhaps have it rephrased if you're struggling with some of the language, perhaps have something like a personal assistant that's there at 2 o'clock in the morning, when your tutor has actually gone to bed, but can answer your questions immediately.
They're coming, but don't expect them yet. But we will be showing some of them in various forums as prototypes, and there's lots to be excited about. But we are very aware that there's lots of risks with using this technology-- some of the topics we're going to discuss later tonight.
ROB MOORE: OK, thank you, Mike. Excellent.
So we're going to start by having a look at your answers. So the first question we asked you is, is using generative AI cheating? So what did we get? So 41% of you said no-- I'm sorry-- 41 of you said no, and 15 of you said yes. Right. And, panel, what would you say to this question? Do you agree that it's not cheating? And of course, I know what the answer is going to be because it's a general tutor answer. So to start with you first, then, Lorna, what's your gut reaction? Is it cheating?
LORNA SIBBETT: It's not cheating if you're open about how you've used it. So I put the qualifier on it. So that's a wee bit of a cheat because we didn't allow you to do that whenever we posed the question. So I'll come down on the side of no, it's not cheating, but don't use it without letting us know that you've used it. And tell us which one you have used and how you've used it. So that's me with my academic conduct officer hat on potentially.
ROB MOORE: Thank you.
Heidi, have we got any expansions on that answer? So did any of the students give us the general tutor, "It depends"?
HEIDI: We've got lots of conversations going in the chat and lots of things that people are talking about. So when it comes to actually cheating, there are people that are kind of on the fence. So Carol says, waiting for more information before I decide on whether it's good or bad. I think people have been putting their responses really directly into the widget rather than talking specifically about cheating. But lots of really, really great comments that are coming out in the chat from people who are talking about how they use the AI.
And I really like what Mike said there about the idea of having a personal assistant at 2 o'clock in the morning-- you know, somebody that if your tutor is not awake that you can then just ask somebody the question. So yeah, it'd be interesting to see. So let us know in the chat whether you think that will be useful for you in your studies, for sure. But yeah, lots of conversations going on about how people are engaging with it so far, Rob.

ROB MOORE: OK, lovely. So we're going to go on to widget 2, and we're going to come to Mike in a minute on this one. So when using generative AI, is it OK to upload some of your module materials and asking for it to generate a summary? And 43 of you said no, 10 of you said yes.
So this is a slightly more technical answer, then, Mike. Is it OK to upload module materials and ask for it to be summarised in your view?
MIKE RICHARDS: OK. Well, I mean, sorry, the default answer is, it depends. We have to do a technical conversation here about something called copyright. Now, copyright is a legal concept that, when you create something, you have ownership over it. The Open University is a copyright holder. So everything we publish-- now that depends-- module materials, in particular, our assessment-- is copyright, The Open University, unless we say otherwise. So stuff on OpenLearn, for instance, is usually not copyright. It's held under a different licence.
So the university has the right to publish its material, but you don't. So that includes sending it to another service. So if you take a chunk of Open University material and publish it on your own website or you send it to a large language model, you have broken The Open University's copyright.
Now, that's a bad thing to do. Don't do it. It helps pay our salaries, for a start. When you send material to a large language model and you ask it to summarise it or something like that, what it actually does is it goes through the data and it tries to produce an answer from it. But what it can also do is it can upload the copyright data into its training set from which it learns and gets better. And that means Open University material could end up being republished by the large language model later on.
There are certain constraints, depending on which large language model you use, about whether data is uploaded into the training set or not. But generally, please don't do it unless we say, here's a chunk of text you can upload and use for your assessment.
What you can do, obviously, is use a large language model and say, hmm, in their assessment, they're asking me about-- let's choose a topic; I don't know-- the pyramids of Egypt. Tell me about the pyramids of Egypt. And then have a conversation with the large language model about the pyramids of Egypt, in this case, but don't actually upload Open University material.
ROB MOORE: OK, thank you. I've noticed we've got a comment come through about it's useful-- something about feedback-- direction from the feedback that you're given. So very much what you've just said there, Mike, which is you can have the conversation, see if it's got any suggestions you might want to follow, but ultimately it needs to be your own work.
And the other thing I'd pop into this conversation as well is that the comments and feedback from your tutor should also not be uploaded. That's actually copyrighted as well. So it's not just the module materials. It's the feedback and the comments from your tutors. So thank you for that.
MIKE RICHARDS: Copyright also applies to you. Anything you create is also copyright, you. So your own words are copyright. So you should think about ownership of material, particularly if you're thinking about going into publishing stuff or writing books or stuff like that. Copyright is really important because, at the end of the day, you might be making money out of it as an author or a creator in anything. The only thing really commonly you can't copyright a recipes-- for a very strange reason. But yeah, treat copyright seriously. It is a thing. There are consequences for everyone by breaching copyright. We do worry about it an awful lot.
ROB MOORE: OK. So widgets 3 and 4 should have been on the screen, and hopefully you've been completing these. So we're going to move on to the next question, which is, when using Gen AI, is it OK

to upload your own notes and generate flashcards for revision? And the results that we get for this, 15-17. So if you haven't voted, there's still time. You'll actually see your votes going on if you're very quick.
So yeah, most people are saying yes on this. So I'm going to come back to you-- well, I'm going to go to-- either-- [INAUDIBLE]-- I've got the power to pick who I want. I'm going to go to Heidi first. Any comments on this one, Heidi? Mainly because I'm making Andrew rush round in the background to try and work out who I'm going to pick next.
[LAUGHTER]
HEIDI: So, again, we just got people voting on this one at the moment. I will give people a few minutes to type out their thoughts so they can expand on that. So I can see-- so we've got 29 yes at the moment, haven't we, about flashcards, and 13 saying no-- oh, 31 now. Yes, that's really interesting.
Yeah, getting some really good comments. You just picked up on something there, Rob, in the chat. It was actually Oudone. I hope I pronounced your name correctly. And Oudone said, it's about generating ideas, not generating answers. Generating answers that you use directly is cheating, but using the ideas and direction from the feedback is extremely useful. And I thought that was a really, really good summary of that there.
And what Mike was saying then-- as I said earlier when I just did the introduction there earlier on, I'm not an expert at all, I'm new to this-- new to AI-- and I had no idea that, essentially, that central repository is keeping everything. So whatever I'm feeding in might be kept and then used elsewhere. And I find that fascinating. I had genuinely no idea about that. So I think it's really great that, yeah, we were making people aware of that. I think it's really important.
ROB MOORE: Absolutely. And as you say, it's the fact you don't think that what you put in there, you're actually providing data to the model so they can use whatever you put in there.
So, Lorna, that question to you then, is it OK, from an academic conduct point of view, to put your own notes and ask it to create flashcards?
LORNA SIBBETT: From an academic conduct perspective, yeah, that one's actually OK. But as Mike said, be aware that you might be protective of your own copyright for your notes and not wish to share it beyond the confines of your own desktop. So that one's OK for me. And it could be a really, really useful way of generating that sort of study buddy-- so something that will generate quizzes for you, something that will help you find key points that you'd have within a flashcard-- so to digest down maybe your comprehensive notes that you've taken across a whole month of study and produce some of the key points, organised in a nice way, to remind you of-- so you might appreciate someone to do that with you, someone to talk things through. So I'm OK with that one if you're OK sharing.
But I think what we keep coming up with is that idea that you really need to know how these tools are operating. So are they just something that's sitting there in your browser and your data is not going any further? Or is it an agreement through using that particular tool that the data that you're giving it can be shared back somewhere? Are you secure about what you're agreeing to do by sharing the material there?
ROB MOORE: Yeah. So you're giving away your hard work. And my personal view on this is flash cards are great. But actually the creation of the flash cards is part of the learning process. So why would you want to give that to a Programme to do? You're actually missing out on some really good interaction yourself.

A better question for me for this one is get it to ask you the questions based on the flash cards. Upload your flash cards. Then get it to ask you the questions because then it's taking on a step that is something you perhaps wouldn't do.

The fourth question, then, is it OK to upload your TMA-- the TMA that you've written-- and ask it for suggestions before you submit it? So let's have a look at this one. Yeah, 48% of you say no. Sorry, I say percent-- 48 of you say no, 4 of you said yes.

I'm going to come back to you, Lorna, because you're our academic conduct expert. Would you think it's OK to ask for suggestions?

LORNA SIBBETT: OK, this is where it really depends, and also because the landscape will change and it will continue to change. So as it's worded there at the moment, OK to upload your TMA? Two dimensions underline that. One is be sure-- you cannot upload the TMA questions or the TMA prompts. So anything that is module material, which Mike has already talked about, you should not be sharing because you do not own the copyright for that. And I know that sometimes what you'll do in constructing your TMA is retain the key questions within it and then pop your response in. OK?

So that you cannot share because you do not own-- or sorry, you cannot-- you can, but it would be naughty-- so do not share that because you will be sharing our copyrighted material if you did that.

As regards uploading your TMA itself for refinement, you need to be extraordinarily cautious about that. You need to check and see what our current OU guidance is. So I think that the link to our current OU guidance has been shared with everybody at some point this evening. Also, check and see what your module is guiding you to do. Are you allowed to do this? Are you not allowed? Indeed, does the TMA question maybe ask you to do something with generative AI? It might be where we're going in the future-- is where we're trying to go in the future, to be honest. So check your module guidance.

After the module guidance, check the TMA guidance itself-- you've got to do a lot of reading here-- the fine detail of each TMA and make sure and see if it's allowed in that particular one. And always check with your tutor if you're in any sort of doubt with any of this. So it's sort of not a straightforward one there, I'm afraid.

ROB MOORE: And I'm assuming, if you've used it in any way, you make it clear somewhere in your assignment that you've used it?

LORNA SIBBETT: Yeah, so long as you're allowed to use it. Do check that you are allowed to use that. And if you're allowed to use it, then what we really ask is that you indicate which one you've used, that you indicate to us the prompts that you have used, and that you share that as part and parcel of the TMA. And that way you're being entirely open and transparent. It's not that you're hiding anything that you've done. And certainly, as an academic conduct officer, I would look quite generously and leniently if someone who had been through a little bit of misguidance-- sort of stepped a bit wrong, but I could see that they had tried to do everything right-- so that makes a huge, huge difference.

I think, as much as possible, you just try to be transparent. You're trying to understand what these tools are. You're trying to understand what our current guidance is, and make it clear to us what you have actually been doing and how you have made use of them.

ROB MOORE: So just to remind people, questions 5 and 6 should now be up on the screen. So you can fill those out. I'm going to come to those in a moment.

I'm going to ask Mike a question now. And this is about Grammarly. Am I right in thinking that Grammarly is a form of generative AI?

MIKE RICHARDS: It has a form of generative AI built into the newer versions. Now, I think it's only in the paid tiers of Grammarly. So we as an institution do say to students, if you struggle with English for any reason, there are tools, such as Grammarly, which can help you with grammar and structure of sentences, paragraphs, and things like that.

Grammarly is a business and they have been working and investing quite heavily in generative AI. So one of the options in Grammarly is, essentially, here's a block of text, rewrite it for me, and Grammarly will go away and rewrite it according to its own understanding of English.

It's pretty good at it. However, there is, like most large language models, something slightly off about the language. For those of you who've used ChatGPT or Copilot or something like that, you'll notice it writes very fluent English, but it doesn't seem to feel quite like a human has written it, which it hasn't. And if you use Grammarly to restructure entire paragraphs, it does look a little odd. So your writing style will change. And your tutors are really good at spotting when you haven't written something.

Traditionally, it's been things such as people who've accidentally quoted material in their TMA answer and have forgotten to say, this is a quote. Our tutors now are spotting that some paragraphs of text do look like they've been written by machine, and Grammarly can do this. So if you're a Grammarly user, carry on using it, but be very wary of using the higher levels of intervention to ask it to restructure entire paragraphs.

And it's not just Grammarly that can do it. Microsoft has the Copilot Programme, which can be included in some of the Microsoft Office offerings. I think it's available for all student accounts in the UK now. And one of the things it can do is also restructure text, condense text, turn it into bullet points, expand a block of text. So be very wary about using these tools. If you do, or are tempted to use it, have a conversation with your tutor. Is this acceptable? And they might have to refer it upwards to the module team or to academic conduct officers to actually get a determination.

And if you do use these tools, just write one or two sentences somewhere in your TMA, and say, I used Grammarly to do the following things. As Lorna just said, being honest and open about what you do won't get you into trouble. My experience is that tutors want to have the conversation with you. They'll be incredibly supportive. They want to help you. They don't want to get you in trouble. No one wants to get you in trouble. It creates a lot of paperwork.

ROB MOORE: Definitely. We are not hunting, trying to track down people that have broken the rules.

MIKE RICHARDS: No.

ROB MOORE: But am I right in thinking that our plagiarism software can actually pick up things that have been uploaded to the large language model? Even if you put it up there and didn't use it, there's a danger that you could be flagged up as plagiarising because somebody else could then end up using parts of your assignment. Is that a realistic concern? That's a concern I have.

MIKE RICHARDS: Right. There are a number of tools out there-- and probably I should put a warning in here about using them personally; I'll come back to that-- remind me if I forget-- that claim to be able to look at the statistical patterns inside the language generated by large language models. And they award a score of how much they think the text has been generated by artificial intelligence.

The Open University has a range of tools we've used for a number of years for plagiarism and things like that, and all scripts go through them. One of the tools is called Turnitin. It's widely used in academia. And Turnitin has an AI detection tool built into it. The OU has experimented with using this feature to award scores, but it hasn't currently used it in academic conduct cases.

There's a large discussion in academia about whether this tool is reliable and how it should be used, if it's used at all. This may change in the future. Currently, the OU does not use it in academic conduct cases, but that may change.
Coming back to my first point, there are a number of websites online which say, upload your script here and we will score it to see whether it's been generated by artificial intelligence. We recommend not using them for a range of reasons. We don't know how they work and how they score your thing. So you might get a very high score for something you've generated yourself. And also, again, the data you upload to them they might put into their training sets and use without your control at that point. So you lose copyright and you might be scored badly against it.
So things like Chat GPTZero is the most famous one, but there's a range of them. I recommend don't use them.
ROB MOORE: OK. Thank you for that, Mike.
And, Heidi, have we got any comments? What's being said in the chat at the moment? Are people helpful?
HEIDI: Yeah, we've got quite a few questions coming in, actually, Rob. And I know we're quite tight on time, but I've got a couple of questions. So Kate has asked, when you say, check your module guidance, does that mean that some modules allow the use of AI and others don't? If so, I think that could be quite confusing for students. So if we could respond to that one, that would be great.
And then Declan has followed up, saying, how does copyright with essays work? We're the creators, but is our copyright then retained when we submit it to the OU, or is that copyright transferred to the OU at that point? And I thought that's a brilliant question, Declan. So yeah, two questions there if we've got time, Rob.
ROB MOORE: We'll hear Lorna laugh now because those are two for you there, Lorna. So the first question was about the guidance. And it will be different on different modules because it will go by learning outcomes, I would guess.
LORNA SIBBETT: Yeah. So it will be different. And at the moment-- and Mike has already alluded to it-- we've got some experimentation happening within the OU. And we will continue to have for a while. And people responsible for your modules are always innovating. They're always rewriting things. We're always investigating the new tools that are out there. They want to make sure that you are being prepared for the life that you lead with the OU. So that means, where there's new technical tools that come in that might be relevant to you, that we have a responsibility to make sure that you've got experience of them, that you can apply them and put them into practise, use them ethically and appropriately.
So we have an absolute necessity to make sure that we begin to encourage students to use Gen AI-- says an academic conduct officer-- we do have that. We need to be doing that. And that requires you to stop and think about why we're actually assessing you in the first place maybe, because we're not assessing out of any degree of cruelty. We're assessing because a well-designed assessment, should help you to learn something. It's maybe helping you to consolidate a piece of knowledge. It's maybe getting you to practise a skill.
Yes, we want you to demonstrate those skills and knowledge. And that's how, often, your feedback is worded. So you'll see a shift in the things we're asking you to do in an assessment-- if you're going to be with us for the next few years-- where we're going to be asking you to make your construction of your response quite visible. So like writing an essay, we'll actually want to maybe see the essay plan. Or we

might suggest, try to use this Gen AI tool to write your essay plan, then critique it, say how you would do it differently, and take it from there. So integrating the Gen AI usefully into what you can do.
So a very long response to, yes, it's going to be different in the different modules. And we're not making this an excuse about that. It's because we need to investigate the capabilities and what's right for you. And that's not going to happen overnight.
ROB MOORE: And while I've got you on the screen with me, Lorna, the next question is about, when using generative AI, is it OK to upload an abstract and ask for the word count to be reduced? And I know you've got a very personal example of this, so I'll let you share that before we look at the votes.
LORNA SIBBETT: I don't know if I can without announcing someone else in trouble. OK, there's not a them and us here in higher education. There shouldn't be. We're all learners. So the staff members I was just talking about within the module there that are writing new materials, they're learning how to make use of these tools. They're trying to learn the capabilities of the tools and develop new things alongside it. Mike and myself, in other aspects of our life, what we're trying to do is generate new understanding for ourselves within areas of scholarship and within areas of research. And when we do that, we have a burden of responsibility to share it with the wider academic community, which means writing papers, presenting at conferences, that sort of thing.
So to get a spot at a conference, for example, we might submit a little summary, a precis, of what we want to present on-- hence, an abstract. Or it might be a summary of a paper that we wish to put in for a publication. So that's our abstract.
But-- there's a "but" with this coming-- so some of the academic publishers are quite on the ball and have already said, if you're using Gen AI, we're unlikely to publish what you have put out because what they want to see is that the academic output from you is what's going on in your head, your ideas. Somebody alluded to ideas earlier, and you saw me scribbling at that point because, within some aspects of academia, the ideas are extraordinarily important. And we need to know, were the ideas from the person or from Gen AI?
So back to the abstract. It is, at the moment, potentially not a very good idea within the academic community to trust to upload your hard-crafted writing and abstracts to some of these tools. Now, Mike might disagree with me, so bring him in and see whether he disagrees or not. I'm a bit unnerved by it. And like I say, some of the publishers say no to it. I think some of them don't have their guidance in place, which is a bit unnerving as well. And others just haven't thought about it at all.
ROB MOORE: I think Mike is trying to answer in Morse code on his keyboard at the moment.
[LAUGHTER]
So we'll have a look at the votes first, and then we'll bring Mike in on question 6. And you can answer that in question 6. So the votes on that first question is, is it OK? And I think, generally, what we're saying is, yeah. Some of you think yes, some think no. But from what you said, then, Lorna, is it's OK. But one of the problems is it could be seen as not being your own work and it could disqualify you.
LORNA SIBBETT: Right.
ROB MOORE: And this is in a professional sense. This is not about your assignments or your TMAs. This is actually as an academic professional, publishing. So there are some bigger implications.
LORNA SIBBETT: So it's the same implications that a student has.
ROB MOORE: Thank you. Yeah.

LORNA SIBBETT: It's the same side of the coin as what the student has in terms of submitting a similar piece of work that we might ask of them. We would want to know, is it your work or is it the work of a tool? Whose thoughts are they?
ROB MOORE: Absolutely. Yeah. So it's making sure it's seen to be your own work and proving it looks like it's going to be a challenge. So when using generative AI, is it OK to upload a questionnaire and the data that you collect if it's got sensitive information? So if you ask AI to summarise it, is it OK? So what are your votes there? And we're seeing 50 of you say no, 2 say yes, which is exactly what I would expect you to say.
So, Mike, why is it not OK to upload sensitive data and asking ChatGPT?
MIKE RICHARDS: Wow, this is turning into a law course, isn't it?
[LAUGHTER]
But we've already done copyright law, where there's a second set of law you should be really concerned about, called the Data Protection Act, which you also may have heard as the GDPR and stuff like that. It's basically concerned with controlling personal data. By personal data, that can be everything from name, age, date of birth, sexual orientation, religion, addresses, national insurance number, you name it-- anything that could be used to identify a living person.
When we work in The Open University with personal data-- which includes your student records, your finance details, and things like that-- we have to obey this piece of legislation, called the Data Protection Act, which has terrifying penalties for anyone who breaches it. And you, if you are doing a research project or any form of assessment which you collect personal data, you are obliged by law-- with no exceptions-- to obey the Data Protection Act, which means you have to say what data you're collecting and how you protect it.
Now, by protecting it, that can mean you store it on an encrypted form on a computer, in a locked room, in a basement with armed guards, which is a reasonable way of protecting data.
When you say, I'm going to upload it to a large language model, which may be located in California, which has a completely different set of data protections and I don't know what's going to happen to it, is an immediate breach of the Data Protection Act, and all the lawyers in the world are going to fall upon you. And the university will not protect you. I can be quite honest with this one.
So whenever you're handling personal data-- which you may do as a researcher, particularly at postgraduate level-- you will be obliged as part of your assessment to conduct what's called a risk assessment. And one of the aspects of risk assessment is, are you handling personal data? If you are, you will then be passed on to various people inside the university, who will guide you through the risk assessment of handling personal data.
But even if you're just doing it as part of a TMA because you're interested in people's opinions, you have to take some consideration. Are you anonymizing them? So rather than having "John Smith," you would have "Subject A," and you would omit any data that could possibly be used to identify them. So you would omit their address, their email address, and things like that. You might be very concerned about things like age and stuff like that.
Again, if you're working on an assessment, the module team will have thought of this. They will give you guidance on what you should and should not do.
Can I quickly go back to a point about writing the abstracts and things like that? First of all, I hate writing abstracts. I don't know how Lorna feels about it. It is the worst part of writing any paper because you've

got to get 10 pages of content into one paragraph and make it reasonable. One thing people don't realise about large language models-- and it's completely counterintuitive-- they can't count. So don't ask them to write 500 words about something because they don't know what 500 words means. You can ask them to condense material. You can ask them to write three paragraphs and stuff, but you can't ask them to write 500 words and get 500 words out of it. It's a quirk of how they work.

ROB MOORE: Thank you. I like that a lot. ChatGPT can't count. We like that. I'm going to ask Heidi to run us through the last widget because I've noticed you've been very quiet, Heidi, because we haven't got to you. So what are the thoughts on, are there any consequences to cheating? And has anything been said in the chat about that yet?

HEIDI: Well, lots of things going on in the chat, as ever. Lots of different conversations going on. And I think we've got some AI experts in the chat ourselves, actually. The conversations are really, really interesting. I just wanted to pull out a few of the key points.

So Tanya has just said, the Data Protection Act is there to protect people, and uploading this data to AI is not right. So I completely agree with that. And having done my master's and then going on to do my PhD, with all of the data and the ethics and all of the approvals and things, that idea of uploading anything to a space where you've no longer got control of the data is absolutely terrifying because, as a postgrad student, they just drum it into of, like, you can't email the data. Everything has to be so super secure. So just that thought of uploading anything is, yeah, quite alarming.

Fiona says, what is the difference between getting AI to fix your work and getting another person to do it? It really feels like cheating to me now. So we get marks for how we structure things, surely. And then Danielle agreed with that point from Fiona. So really interesting that people are having those-- they're exploring that further.

And then I do just want to pick up on a brilliant point. It's completely off piste, but I just love this comment from Ruth, who said, for those of you who are a bit more geeky, something I'm trying to learn how to do as part of an experimental project with Friends of the Earth is to instal a local LLM to set up a generative AI tool that can embody a snail persona that can work on a Raspberry Pi for a snail robot. How good is that?

ROB MOORE: Excellent. We want you to come back and tell us how you get on. Come back. We want to see this Raspberry Pi snail crawling around. So once you get that up and running, please come and tell us. Come on Student Hub Live and demonstrate it to us. We'd love to see that.

In fact, we're talking to the STEM faculty later in the week, and we're looking at some of the large citizen experiments you can get involved with. So make sure you dial in on Thursday because we're going to be talking about that.

I am going to come back to Lorna, as our expert and academic conduct officer. So, Lorna, consequences, if you want to run us through them quickly. What is the worst-case scenario? Not that we want to frighten anyone, of course.

LORNA SIBBETT: Kicking you of the university. That's the worst case.

[LAUGHTER]

But there's quite a few steps along the way there. I mean, I think as I've said already this evening, the key thing for yourselves is pause and think about why we're asking you to do something within any task that you're being given. It should have a purpose. If you're not clear on why you're being asked to do something for a TMA, then ask your tutor. Check the forums within the module. And that hopefully will

give you clarification on it. I'm not even talking about, are you allowed to use Gen AI or not? Actually understanding what you're being asked to do and why you're being asked to do it is really, really important.

Also, when you're making notes from any module materials, sometimes what you see that people have done is they've copied and pasted little chunks of material from the module, popped it into a little Word document that sits on their desktop. Then they come to respond to a question, write an essay, a synopsis about it. Because they've done that about a month or so ago or two months ago, they forget that, no, that wasn't written in their own words. And instead, that's actually copied and pasted from module materials. And what we get presented to us as something that a student might authentically think is their own work is actually a little patchwork quilt of materials.

So guidance to you is be sure about what we've been asking you to do, and start at the point of your studying itself and taking really good notes that are in your own words so that you don't accidentally fall foul of it. What can happen to you-- with generative AI suspicions at the moment, those-- how do I put it-- we aren't using, as Mike has already said, the software tools like Turnitin that high match scores and so on because we don't trust those tools. So what we will look at is, does that submission that you have made seem to be written in the same style to your earlier TMAs? Can we see that trajectory? We might come back to you at some point and ask for your evidence to show us how you arrived at a particular answer. And I think, as I've already said this evening, what you might see in terms of TMA questions as they evolve over these next few years is that you will increasingly get questions that prompt you to make your construction of an answer very visible.

So, for example, rather than just submitting an essay, we ask you to submit the plan of an essay and get the final essay afterwards. So we might even get you to use Gen AI in that. So worst-case scenario initially is you might get reported to be investigated by an academic conduct officer like myself. You'll get a notification that you're being investigated. You'll get a little letter-- email-- from an academic conduct officer, saying, we've noticed this. Would you like to tell us something about how you arrived at that answer? And you've got 10 working days in which to respond to that. If you don't, then the academic conduct officer is in the position where they have to make an answer devoid of any of your input. But if you do respond to it, we will listen to what you have got to say. We'll listen to it to understand your process.

It may well be that your response could really help us within that module to clarify guidance on a TMA or otherwise within the module. So you could do something that's actually really supportive. So the academic conduct process isn't purely about being nasty at you and slapping you. It's about us recognising what is working and not working within module assessments on the structure of the university and informing into that.

So a first-time academic conduct case, we're unlikely to do anything too severe. You could end up with a warning, guidance, study skills, that sort of thing. It's repeat offences that are going to get you in trouble. But en route, there'll be a lot of support offered to make sure that you understand better why we're asking certain things of you so hopefully you're not going to repeat it.

ROB MOORE: Absolutely. We're not trying to catch people out. We're trying to maintain the standards. And you're all working hard to achieve a degree. And we want to make sure that the standards are there and also improve your academic abilities and the way that you write your assignments.

And I'm glad to say that everybody who answered said, yes, there are consequences if you're caught cheating. And I think we accept that. Today isn't about trying to frighten you or get you to be petrified of going anywhere near AI. Today is about saying, yes, there's a tool out there. For some things, it's going to be good. But make sure you understand what it's doing. Make sure you've got the right to upload things that you upload because you lose control once you've uploaded it. And ask questions.

And I know Mike has got some top tips. So the last thing we want to do today is for Mike to run through his list of do's and don'ts. And this is how to make sure you stay on the right side of generative AI. And with that, over to you, Mike.

MIKE RICHARDS: OK. Yeah. I mean, absolutely everything Lorna has just said and everything Rob has just said. We're not trying to catch you out. If you somehow get flagged for whatever reason that your tutor thinks you may have used generative AI and you haven't indicated it, there will be a conversation. It's not going to be an immediate thrown in the dungeon, we'll never talk to you again.

The first thing I would like to say is, if you haven't already, please read the module guide. We put these together. We try to cover everything in there, and it becomes this huge document, which looks incredibly intimidating. It will have a section about assessment, and, increasingly, the section on assessment will mention generative artificial intelligence.

To go back to an earlier point that one of the viewers said, is it going to be confusing that some modules are going to allow it and some won't? We are in transition, as modules get updated, to deal with generative AI. But also, please remember some of our programmes are accredited with other institutions and other bodies, and some of them will have rules that differ from our own about whether you can use generative AI or not. So look in the module guide. It will tell you how you can use it.

Then, have a look at the Open University general student information about using generative AI. It will give you the rules about how to say you've used it, what is acceptable, what is not acceptable. If you're in any doubt, have a conversation with your tutor. You will not be the first person to raise the points with them that generative AI exists. They will give you some guidance. They might need to refer up to the module team, and the module team might need them to refer further up. So you may not get an immediate answer. Be patient. We will let you know on how you can use it and how to use it safely.

We've said it so many times-- please don't upload any OU material-- and that includes any module materials and any assessment materials-- to a large language model. It could be used in a training set. It is a breach of copyright.

As we've said, the OU is researching the use of large language models. There may be soon-- and by soon I mean not immediately-- an Open University large language model that you will be able to have a proper conversation with about module materials. But this is an experiment we're having at the moment to make sure that it's safe, reliable, and it gives the right answers. Because one thing we haven't quite mentioned as much as we probably should of-- large language models sometimes get things very wrong. And for anyone who doubts it, Air Canada lost an awful lot of money recently fighting a legal case, where their large language model chatbot gave a passenger incorrect advice about fares. And Air Canada said, well, it's just a computer. And the court said, actually, it's your computer. You own it.

So we don't want to do that. We don't want you to use one of our large language models and get wrong facts.

And the final point is, if you use generative AI in your assessment, and it's allowed, don't copy the material straight out of the large language model and paste it into your assessment unless that's part of

the question. So if the question said, here's a prompt, and copy the response into your solution, that's fine. But if you've asked it about-- I don't know; let's go back to the original example-- the pyramids, and you just copy and paste its answer into your TMA, all you're doing really is quoting another source. And we don't tend to give marks for quoting sources. Quotations are there to support an argument, not be the argument itself.

So by all means, within limits, use a generative AI to develop an answer, but expand it yourself. Find other supporting materials. And the other thing-- going back to this point that the AIs make stuff up; it's a phenomenon which gets a bad name, hallucinations-- you have to check it. Just because it says there are pyramids in Egypt. You should go away and check there are actually pyramids in Egypt. Hint-- there actually are. But use another source.

It's very much like the guidance we give you for using online resources such as Wikipedia. We say, use Wikipedia to develop your answer, but we usually then have said, and several other sources. So spend some time supplementing anything that comes out of generative AI by going to credible external sources and referencing those-- because I'm going to mention the word "referencing"-- to make sure your argument is coherent.

I've used generative AI. And occasionally, I'll go, really? And I then go down a rabbit hole of trying to prove that. And, yes, occasionally, it does get stuff wrong. And the best bit is you don't know when it's going to get stuff wrong. Sometimes it trips over really common things. Like just now I told you they can't count. And it's quite true. Large language models can't count. And they get very detailed, very sophisticated things, completely correct.

So you do need some time analysing its answer, supplementing it, and making sure that the information you're providing in your assessment answer is correct and in your own words.

ROB MOORE: Thank you for that, Mike. And that's some really good tips there to make sure that you stay on the right side of the AI wave.

As always, we're really running out of time, as we always do in these sessions. We could easily do another half-hour, 40 minutes. But, Heidi, any final points that you want to bring out from the chat?

HEIDI: Yes. So a question there for Mike. So we've got some concerns from people in the chat, asking, does the OU really have a dungeon? You said there that we're not going to put people in the dungeon for using AI. So I think we need to expand on that.

[LAUGHTER]

ROB MOORE: We had a cellar bar. I'm not quite sure I'd count it as a dungeon. If there's a dungeon, I haven't found it yet. But I'm not ruling it out.

[LAUGHTER]

HEIDI: That's great. Thank you, Mike. And before we wrap up, Alex has been asking me throughout this session whether you can give us the name of your LEGO parrot that you introduced at the beginning. And I haven't got round to asking you that question yet. So what's the name of the parrot that sits behind you, please, Mike?

MIKE RICHARDS: We haven't-- we should throw it open to a poll-- we haven't actually given it a name yet.

HEIDI: Good idea.

ROB MOORE: OK. Yes. So if you've got--

MIKE RICHARDS: The only reason we got a parrot is because there's a very good paper about large language models, which says, basically, they're parrots. They babble words to you. They don't know what the words mean. And that's completely true. A generative doesn't know what any words mean. So the phrase was-- they called them stochastic, which is a fancy term for statistical-- parrots. So all the work we're currently doing, we're putting little parrots-- we're hiding parrots everywhere in our new materials we're writing.

So for those of you who are going to be around for a while, look out for the parrots. But no, he hasn't got a name. Give us a name.

ROB MOORE: Excellent. Well, my suggestion is Polly Carbonate.

[LAUGHTER]

We are right at the end now. So I'd just like to say thank you to Lorna and to Mike. I'm sure we'll keep coming back to this topic because it's going to grow. The advice that we give is going to change because, as things develop, we're going to learn new things. There are going to be new ways that we can incorporate this. So keep checking on this.

I know in the links on Student Hub Live, on our web page and in the chat, there were some links around where to find the latest guidance.

So thank you, Lorna. Thank you, Mike. And thank you, Heidi.

And just a final reminder to everyone, we have a feedback form that links into tonight's session. If you like this topic, if you like the way that we've done it, let us know. If you didn't like the way that we did it, let us know as well because we want to make sure these sessions reflect what you want to get out of them. And the best way to do that is fill in the feedback form, and then we know what you want more of, what you want less of.

A couple of things just to remind you of for this week-- tomorrow, at 7:00 PM, we have the end-of-year party with Heidi and Isabella. And if you've been to those before, you know it's going to be manic. And I believe the Silent Knight is making an appearance.

Then we've got-- Margaret is going to be doing her third workshop on critical thinking. That's on the 19th, at 7:00 PM. And then I'm back again with Heidi, on the 20th, at 11:00 AM, looking at sustainability and how you can build it into your studies and make yourself more employable.

So it's been really good to see you. Also, again, thank you, Lorna. Thank you, Mike. Thank you, Heidi, and Jay and Mary on the chat. I've had a great time, and I'm really looking forward to seeing you all again. So in the meantime, have fun and see you soon. Cheers.

[MUSIC PLAYING]