



OPEN
CENTRE FOR
LANGUAGES &
CULTURES



Distinguished Speaker Series

How LLMs such as ChatGPT and AI will Transform Adult Education

Prof. John Domingue,
Knowledge Media Institute, STEM, The Open University

27/09/2023

Edited transcript

Mirjam Hauck

The Open Centre provides access to high quality opportunities to learn languages and engage with cultures for learners from all walks of life and from around the globe. It increases awareness of the strategic importance of meaningful and effective communication between individuals and groups. We offer short courses that are designed to foster curiosity and understanding between individuals and people, both locally and globally.

We offer paid and free short courses on the Oh use virtual learning environment and we offer free tasters of the paid short courses on the OU's Open Learn platform. Our courses cater for different audiences and serve different purposes. We cater for leisure learners for continuous professional development learners, and all our courses are aligned with the Open University's social justice and inclusion agenda.

The distinguished speaker series is part of the Open Centre for Languages and Cultures and we usually invite colleagues who are experts in the topics that are covered in our short courses, and today, we are honoured to welcome Professor John Domingue from the Open University's Knowledge Media Institute - [KMI](#). He is a professor in computer science and was the director of KMI from 2015 until 2022. He's also the President of [STI](#), that is an organisation that specialises in semantic technologies to address challenges associated with communication. John has contributed over 240 refereed articles in fields such as semantics, AI, the web, distributed ledgers and e-learning. He has also led many prestigious projects over the years, most recently between 2022 and 2023, a project to develop a smart national educational content platform that incorporates cutting-edge AI to support further education educators, and since the beginning of 2023, has been at the forefront of examining the impact of generative AI on higher education. Today, he will talk to us about how large language models like ChatGPT, and AI will transform adult education.

John, over to you.

John Domingue

Thank you very much for that kind introduction. I'm going to give a quick overview in around twenty-five minutes of how large language models and AI will change adult education. Some background and motivation first.

There are lots of definitions of artificial intelligence out there. This is the one that you get from Google. So, it's a definition of AI that comes *from* an AI and basically, we define AI as getting computers to do things which humans would regard as intelligent, such as seeing, recognising speech, problem solving, learning, etc. There have been historically two main approaches to artificial intelligence, symbolic and non-symbolic.

So, by symbolic: these are graphs, linked nodes of concepts. The best example of this is a knowledge graph from Google, which drives their search. If you type 'London' into Google Search, you'll get a panel on the right with photographs and the weather, and that comes from a large conceptual symbolic graph. And then the other approach is more statistical, which underpins things like neural networks and machine learning.

Another big motivator for the work that we've been doing is [Bloom's 2 sigma](#) problem, which maybe many people know about. And that was a study in the 1980s comparing a standard class taught normally, to students taught via mastery learning, where students are not allowed to move on until they've mastered the previous concept, compared to students that are taught one-to-one, and students that are taught one-to-one performed two sigmas above the students who were taught conventionally. So one-to-one teaching can really improve student outcomes, and that's really the holy grail of approaches to using AI in education.

At the OU we've been looking at different AI approaches and we've come up with what we call an 'OU AI ecosystem', which is a goal for us. In the middle is - I mean, simplistically, you might think of as a brain - it's an enterprise data hub, where we have all of the content from the OU, everything we've ever written, and we have all the interactions with students, and then we make that machine understandable, and we have different mechanisms for understanding that. Around that we create a number of services: a digital assistant for every student (I'll give a little proof of concept of that in a moment); an assistant for a tutor (I'll show you something that we built towards this as well); something to help in the creation of content (I'll demo that later); something to support assessment, which, of course, is a big activity for anyone involved in teaching; and something we call the Dream Machine, which is an AI career coach.

I'll start with something that we built a long time ago: OU Analyse. We actually built this in 2012 in KMI. It's a learning analytics engine. So, we're basically applying machine learning to learn what the students are doing and predict outcomes based on the interactions with our virtual learning environment, and also their demographic data. A third of our four and a half thousand tutors, our AIs, now use this, and every tutor has access to this.

So, I'll just show this video. What we're doing is - I'm just moving the timeline back to 2016. We compare - this is comparing students in 2015, who started October, to students who studied the year before. So, you machine learn from last year's students, and that gives you your model for this year's students. I've zoomed in to a particular student, and now I'm comparing an individual student to their peers. This tool predicts with 95% accuracy in the best cases: is the student going to pass the next assignment? Are they going to pass the

course, overall? We've carried out lots of studies on OU Analyse and the main outcomes of our studies are that the tutor using OU Analyse is the second highest factor in students passing the course. The first is their previous performance. Students who use the tool get an increase in retention, the students do better, and later studies showed that students have a 7% greater chance of passing the course, which rises to 10% for BAME students if the tutor is using this tool. So, as I said, we built this tool 11 years ago and we regard it as a success story.

Let me move on to another service that we've built, actually a proof of concept, which is an AI career coach, which we call 'The Dream Machine'. If we think about students, especially students at the OU, they have dreams; they want to do a particular job, maybe they want to be a nurse or a doctor, or an astronaut, and they may think, what skills do I need to do this job, or they may want to change their job. They want to know what skills are in demand, what positions are in demand. Also, with the accumulation of qualifications and skills I already have, what are my opportunities? What's the best course for me to take? If you think about this, then the link in these areas between occupations, which may be linked to job postings and courses, which have learning outcomes, is a notion of a skill. We created a knowledge graph for this, and I'll show you a little clip of a video that we made about how the Dream Machine might work in practice.

Video footage

Create a huge graph of knowledge, about 2 million different link concepts. Now if you take my CV, you could see that in the past I was a project manager. So you can link a node which is Maxi to the node which is a project manager. Now automatically I'm also linked to the other nodes. You can know that I manage people, I am able to do planning and such things. So knowing what it knows about the job market, and knowing what it knows about you, it can link these nodes and find the best path to match you to the perfect job.

Okay, so that's just a quick glimpse of what we could do with symbolic AI, in linking students to jobs and supporting them.

Now, I want to delve into large language models, of which, of course, the best-known example is ChatGPT. I'll start with the basics. Generative AI systems are predictors, they predict, in this case, a next word after a phrase, or they can predict an image given an input. So here, we can see on this slide, we have a beginning of a sentence, 'The best thing about AI is its ability to...', and we have five predictions. 'Learn' is the most common word that comes next, then 'predict', then 'make', 'understand', and 'do'. So, when anybody uses a generative AI like ChatGPT, you put in some text, which is called a prompt and then behind the scenes, the system is looping and predicting the next word. Here's a piece of code: I've said, 'take that phrase, predict the next seven words', and it will predict the next word is 'learn' then 'from', 'experience', some notation, then 'it', 'it's', and 'it's not'. And that's what it's doing, it's predicting the next word, over and over again, in a loop.

Now, some interesting things about these models. The size of the model you're using is important and also how random it is. It turns out, if you always choose, the next word that is the most highly predicted to come next, the output tends to be very flat, and boring and repetitive. This is some text that comes from GPT2, the predecessor to GPT3, which is used

in ChatGPT, and if I just read it, you'll see: 'the best thing about AI is its ability to learn from experience. It's not just a matter of learning from experience, it's learning from the world around you. The AI is a very good example of this. It's a very good example of how to use AI to improve your life. It's a very good example of how to use AI to improve your life', and so on. So, you see, it's not so great. If I increase the size of the model, I go to GPT3, and instead of choosing the most common word, you choose words statistically at about the 80% mark, so about 0.8 value where 1 would be the one with the highest predictor, then you can see, I get text, which is a lot better. So, if I just read it, it says: 'The best thing about AI is its ability to learn and develop over time, allowing it to continually improve its performance and also be more efficient at tasks. AI can also be used to automate mundane tasks, allowing humans to focus on more important tasks. AI can also be used to support decisions and provide insights that would otherwise be impossible for humans to figure out'. So, you see this a lot better and that's from increasing the size and adding an element of randomness. When you use these tools, there's usually something called a 'temperature button' where you can balance between predictability and creativity.

Underneath these systems is a neural network, which is where you have nodes, it's inspired by the human brain, and we have things which emulate neurons, but very simply, to nodes which do calculations, and the nodes are linked to other nodes. The nodes are arranged in layers, and the strength of the connections is determined by a weight, and learning dominantly means changing the weight of the connections. Once the weights have been learned then the neural network propagates numbers, numbers come in and numbers come out and map to meanings, which are words for ChatGPT.

ChatGPT, built on GPT3 has 400 layers of nodes, it has millions of neurons, with 175 billion connections, 175 billion weights, which have to be learned. The human brain has a million billion in comparison. It's a very special architecture which came out from Google, called a transformer architecture, where the bulk of this is that you pay special attention to some parts of the input. It was trained, basically by reading the entire web, and also trained through human feedback where humans would interact with it and say this output was good, this output was not so good.

So, given all of this, what's the potential impact on education? I don't have time to say much in general, I can put some examples from the OU. But the latest things I've seen are, for example, the University of Michigan are rolling out generative AI to all of its students. They're providing access to ChatGPT at zero cost. They're allowing students to use large language models to query and question their own datasets, connecting to some other common services such as Google and Canvas, and, finally, they're providing a toolkit, where students who are inclined can construct, train, and host their own AI models. There are websites such as 'There's an AI for that', which have nearly 8 000 AI's around - over - 2000 tasks. Here, I've typed in 'education teacher', and you can see there are 153 tasks, and there are AI's which support this, and if I look here, I can see there are online courses, AI support school lesson plans, exam preparation, speech training. If I click on 'online courses', I get things like 'customising personalising study guides', 'online course creation', etc, so there's a lot of activity. The thing I should say about this tool is it's not just something you interact with. But also it becomes a platform to build applications, so what a lot of developers are doing is you string, a number of large language models together, so you may say, okay, I want to do the whole thing of planning a trip, and booking hotels, and booking sites, and I

will use a number of live language models, which do planning, do booking, do cultural insight, etc., and they're linked together.

Now I want to show you some of the things we've been doing at KMI, within Test and Learn, and also within the Innovation Foundry, a number of things that we've built and a number of activities that we've been leading on. The first is an AI module writing assistant, using generative AI to support the creation of modules, which is something that takes a lot of time at the Open University. We've run five workshops with module teams in computing, languages, psychology, the Faculty of Business and Law, and also with LDS [Learner & Discovery Services] and we came up with 170, plus challenges and nearly 150 ideas. A lot of the ideas are focused on improving the learning experience for the students. The second largest category were ideas that support content creation, then also providing advice, supporting the students on their journey. I'll show you a couple of examples of the ideas.

One was real-time assistant. So, whenever a student is reading any content from an educator, they could have a parallel chat window, where there's a personalization of that content, so, for example, you could say 'explain this concept to me like I'm 14 years old', or 're-explain it from a different point of view'. We could have personalised learning journeys, so, re-ordering how content is provided, changing assessments maybe based on the analysis tool like OU Analyse would produce. So, each student gets their own very personalised pathway through any learning materials.

And then finally, for the examples. Within some courses, students need to interact. For example, in psychology, students need to learn to be a counsellor and they can roleplay with an AI, so an AI could be a subject - maybe with a particular issue - and then the student plays a counsellor and gets AI feedback afterwards.

A couple of other things I want to explain before I give you a screencast of a demo. In KMI, for the last 10 years, we've been working on something called 'Core'. Core is the world's largest aggregation of openly available scholarly articles, we have nearly 300 million research articles which are semantically indexed. And now we've linked this to GPT. And I'll show you this in the demo. This was a PhD in it from Quan, and he did a study of what types of learning activities are supported in OU module content. If you look at this module content here, you can see that students are directed mainly to carry out tasks related to assessment, or assimilative. So, sitting back and absorbing material, and maybe with AI, we can maybe change the balance if you want to, or even make a change on the fly.

Okay, so here's the first proof of concept using generative AI. This is the AI generated AI module writing assistant. So, I'm going to select a module; I'm selecting one in Languages; I'm selecting the first unit, and then I'm going to say, 'I want the AI to write a conclusion for me'. The AI analyses the material, and then just starts writing the conclusion, you can see it, spitting out here, if I want to speed it up, I say, 'okay, auto-complete this, and then as an academic, I can read this content, see, if I'm happy or not. We did an analysis and found that 40% of the time, the academics are happy. Then I asked it to write a concluding sentence. and then it adjusts and writes a concluding sentence. So, you can imagine that we could move to writing materials, through interaction; we interact with an AI to support us. On the right, I can save conversations if I want to, I can load previous conversations. Now the AI has automatically generated some quizzes for this particular unit, I can have a look at these, some multiple-choice question were provided, some options and an answer, and I can look at another one, a short-answer question, I can look at the others and see if I want to include

them or not. This is core GPT, so imagine academics using this, if they want to understand what is the state of the art in their particular area. Here, I'm asking about papers related to science systems and semiotics. It's returned five, research papers from his corpus and gives me a short paragraph overview. Here's a similar response around 'accent and identity'. This last button, here the systems analyse the activities in the content so I can have a look at what activities the content is supporting and it gives me the option to change these. So, I can say, 'okay, this is a communicative type of activity, I want to make it an assessment' and then the system will automatically generate an assessment for me, and I can see if I'm happy or not with that. Then finally, this is something which is productive, the students have to produce something and I can say I want to change it to finding and handling information, and again, it's always the academic that checks the response.

Now, I'll move over to a student point of view, and show you a quick demo of something that we built on using generative AI to support students. It uses the same engine underneath. Here, I've selected a module and I select a unit, I select the same unit; I asked the system to load in material; then as a student, I can say, 'okay, I want to test myself, please generate some quizzes for me'. I asked her to generate six quizzes, and then I can type in my answer; I can answer each of the questions, 'what is an accent?', etc. I put in my answers and then I can ask the AI to give me a mark, and it will say whether it's correct, incorrect or partially correct and then, if it's not correct, it will give me a suggested answer that I can use as feedback.

This part of the system we have not validated yet, we're in the process of doing that. The students could say, 'okay, I have a test coming up, I want to have some flashcards to help me revise', and then the system provides some flashcards. I can save them to PDF if I want to. And then finally, if I want to, I can have a discussion with the system, so the system will ask me a question, I can give a response, and then the system will comment on my response, and we just enter into a chat dialogue around this. So again, this is always focused on a particular unit in a particular module to make sure that the system is giving them answers which are really tied to the course.

Okay, so this last thing is a demo where we've connected the engine that I showed you before, to a human avatar. There are a number of companies that have come up with human-like avatars which can understand speech input and can give a response. What you'll see is: I ask a question, which is not planned at all; the system does a speech understanding, sends the question to a generative AI; a generative AI comes up with a response, you see that come back as text; and then the human-like Avatar gives an audio response.

Video

What is the difference between structured and unstructured interviews?

In a structured interview, all interviewers ask the same predetermined questions in the same order. This makes it easy to compare answers between different candidates. In an unstructured interview, the interviewer asks questions as they occur to them and follows up on the answers given. This allows for a more natural conversation, but makes it harder to compare answers.

Okay, the key thing there is what we're adding is - there's a lot of AI into getting the facial gestures right and the pronunciation right, and it is improving all the time. In fact, yesterday, Open AI announced that ChatGPT is going to have audio input and speech output in the next

weeks. So, with this type of technology, we can basically turn anything that's written, any written content, into an interactive session for students. We did a survey a while ago, actually, on a far simpler version of an avatar to the one that I showed you, with 155 students, and the response was pretty positive that students are favourable to the concept. Most of them would use it, they think it would be useful to support their learning and they'd like to have it in the OU's virtual learning environment.

So, to wrap up, generative AI is based on neural networks. The ones that you're seeing have been trained over the entire web, they generate by repeated predictions. Lots of people predicted this technology is going to be as disruptive as the internet and the industrial revolution, and I and a lot of others think that adult and distant education will be highly disrupted by this. We are looking into technology seriously at the Open University.

Thank you. This was collaborative work with a lot of people, it was not my work, and funded by Test and Learn in the Innovation Foundry. Thank you very much.

Mirjam Hauck

Thank you, John. Big round of applause. Thank you. That was very, very insightful. It was like a whistle stop tour of all things AI and particularly with the OU lens on it. We have a few minutes for questions as...

Calvert McGibbon

I've unmuted myself. Yeah, really, really very interesting. John, thank you so much for putting that all together. I just wondered what we might think about plagiarism and where the boundary between supporting students' writing and actually just using AI to create data might fall on your view?

John Domingue

I've been working with Mychele Pride [Associate Dean, Teaching Excellence] on the OU policy, and you can imagine that many universities are panicking about this. So, I think the viewpoint the OU is using is a good one, which is that we treat this almost like a human agent. Students are not allowed to just copy from Wikipedia or the web, you're not allowed to ask another student or another person to write the essay for you. So, in that sense it's ... you're not allowed to use other sources or you can use other sources, but then you must cite them. That's the first part of response. The second and final part of my response is that the world is going to change. We know that this technology will come into Microsoft Office so then everybody gets their own personal generative AI. And then work becomes like - everybody has their own infinite army of interns. And we need to teach our students to decide ... I heard a great phrase the other day, that you shouldn't have analogue students coming out of your university, so we need to have students that are fully digital, and fully digital means understanding this. So, not a full answer, but I hope I've scratched the surface a little

Mirjam Hauck

Thanks John, Zsuzsanna, you've got your hand up.

Zsuzsanna Bárkányi

Yeah, so I just wanted to ask how widely available these tools are, or will be for us, you know, at the University for our students - for all of them and for us when we are creating materials, and also linked to that, how it works with other languages?

John Domingue

There are a number of things one needs to do with this technology, to get scale. One is - there are three - you need to somehow tune it on your own material; there are different ways you can do that. So, instead of speaking with an average web voice, it speaks with an OU voice, they are you. You need to carry out experimentation to make sure we get the fit right between what the students or the academics need and the tool we build. And we need to train staff because there's a sort of different mindset needed for interacting with these machines. Each of those requires investment. I have a small amount of investment from Kevin Shakesheff [Pro-Vice-Chancellor, Research & Innovation] and we're working to get bigger investment at the OU, and we know that getting investment anytime, and at this time, is not easy, but that's what we're working on quite hard. But the goal of the team is to make this available to every academic, all the relevant professional services people like LDS [Learner & Discovery Services], and also all the students. So that's the goal. For me, it's just a matter of getting the required investment, doing the testing. We'll always welcome people to interact with us on this and willing to use early beta versions.

Mirjam Hauck

Azumah, I think you are next.

Azumah Dennis

I'm just wondering how we're managing the equitable or inequitable dimensions of AI. Perhaps the issue is resolved if we give it the database to work on, and partly in that, I'm responding to what might well be urban myths or just things which are popularly said, from visuals that don't recognise black faces, or when you ask AI for 10 philosophers, it returns 10 DW male philosophers, nine of whom are dead. And how we can make sure that the knowledge we are being presented with is global, is equitable, and can meaningfully be something that I as a human being, feel speaks for knowledge that I would reproduce? Perhaps I'm making the assumption that I'm as good as or even better than the AI, but that's where my questioning is: how do we guard against it reproducing inequalities?

John Domingue

Yeah, great question. So, I'll answer the question in a sec, but I just want to sort of turn it on its head. One of the things I say when I'm making an investment case is that it is important

that the OU enters this space so we don't just leave it to the big tech companies. There should be voices in this space that have social justice in their DNA like we do. Now, to answer your question, there are a number of techniques that we're applying. One is you seek to train the model on curator content that you say, 'this content is the right type'. The other thing you do is, whenever you ask it a question, when you send a prompt, you say, 'okay, here's a question, I want you to give an answer. When you give your answer, only use this material in your answer'. So, in fact, what we do first is we do a search of our material, say these are the relevant chapters from all of our huge library content, use this. The other thing we can do is when a response comes - actually there's two more things - one is, when we ask the question, we actually augment the prompt, so, we inject things into the prompt, such as, 'please be equitable, think about gender equality, think about racial equality'. So that's added in the prompt. And then when the response comes back, you check that as well, you say, 'okay, just look at the terms and entities and look at the spread of those terms'. So, these are called 'guardrails'. In terms of, 'is it as good as what we do, or an OU academic?', the only way to do that is through extensive testing, and so we've done some, but you need to continue to do that.

Mirjam Hauck

Thank you, John, we need to move on to the second part of this session. And I will thank you again, John. And I'd like now, Andrew to share his screen and introduce our latest short course in the Open Centre for Languages and Cultures in AI. Over to you, Andrew.

Andrew Gargett

Gosh, that was a wonderful talk and now I'll attempt to follow it with some details about the new short course that's been recently released within the OCLC and this is the short course 'Language Technology in the Workplace', and my name is Andrew target, and I authored the course.

The course provides a solid understanding of applications of language technology, across all aspects of the modern workplace and there is particular focus on responsibility, transparency and fairness in such technology.

Overall, the aim is to help you to acquire a critical appreciation of such technology to develop awareness of the risks of using language technology and ways to reduce these risks. The course has several specific aims, which are as follows - you can see these on the screen. You will be able to formulate your own definition of language technology. You will have described major historical, social, political and economic issues and language technology. Upon completion, you will also have critically evaluated primary source information on these issues, be able to formulate your own responses to these issues, and these responses will be driven by fact-based critical research that you will undertake during the course. You will have learned about key developments in the history of language technology, and also about some of the major recent developments in this field, and to be able to identify key features of risks in language technology projects.

The course presents ideas and views of prominent language technology researchers and industry leaders. And the aim here is to keep the material up-to-date, and as highly relevant

as possible to current challenges of language technology, many of which John has already presented in his talk.

Now, given such areas of language technology have been historically misrepresentative of society more generally, as came up in the questions after John's talk, the course attempts to address such issues in two main ways. Firstly, in the details about how such technology can contribute to these kinds of issues in representation, but secondly, the course was created utilising a co-design methodology to enable collecting contributions from OU students representing marginalised groups, who came along to workshops and collaborated with us to provide us with their own perspectives, voices and lived experiences in order to help us shape the course content and design. The project, which I'm referring to here was one that was jointly run with Mirjam and also Mychelle Pride, through a series of workshops at the end of last year.

The course consists of five units, each estimated to take four to five hours of study, approximately 25 hours therefore in total. Each includes interactive tasks throughout and there's a quiz at the end of each unit. There also is the option to engage in forum discussion with others on the course moderated by a learning advisor. The course will be open for between six to eighteen months depending on when you register, and this is in order to allow learners to complete at their own pace. They'll have read only access for a further three years. And overall, the cost of the course is £195.

Here on this slide, you can see a list of the available units. The overall flow of the course is as follows. After providing suitable background to AI technology and its risks in units one and two: 'Language technology all around us', and 'Responsibility transparency and fairness in language technology', the course moves into detailed consideration of technical details of how language technology works in unit three, and this is aimed to a general audience. Unit four introduces a specific use-case of language technology in education, including discussion of the impact of technology such as large language models. Many of the issues that John raised in his talk are covered throughout the course on this specific technology. Finally, in unit five, we complete the use-case on education by looking at how language technology can be evaluated for potential risks, harms, as well as less than optimal performance of specific language-related tasks, and the central role of evaluation in effectively and safely designing such technology now and into the future. Unit five also includes a general course wrap-up.

Interactivities throughout the course explore complex and very contemporary challenges for language technology and they use a range of media in order to do this such as online talks, in-page activities and quiz items in order to check your progress. On this slide, we can see an example activity from unit five that examines a presentation by academic and sign language researcher Mark Wheatley, on the challenges of co-designing technology for users of sign language. On the right of this slide are two examples of where understanding of this activity are later tested through interactive quizzes.

Here's another example of interactive activities and this time from unit one. In this example, learners are provided with a wholly new digital resource created specifically for the unit and for the course. It's an interactive timeline collecting major areas in the history of language technology development. The timeline includes natural language processing, or NLP, the area focusing on developing methods for getting computers to communicate with humans, also machine translation, whereby NLP is used in order to carry out automatic translation of

one human language into another. A good example of this being Google Translate, and also neural natural language processing, which is the use of artificial neural networks, inspired in part by natural neural networks in humans and other animals, which are systems that can learn to how to do something rather than having to be programmed to do so. Various activities throughout the course are framed around these timelines. And on the right of this slide, you can see an example of a quiz question which draws on the information presented in the timelines.

The course forums are moderated by a learning advisor, who is there to support learners chiefly by interacting through forums. The forums enable learners to interact with other learners who have different experiences and perspectives to offer. Also, the learners who successfully complete the course (up to 70% on the units' quizzes) are awarded a digital badge at the end of the course. These badges can be shared on social media or added to email signatures and can be used for continued professional development, job application, promotion, and such. Here you can see an example of a digital badge for the first course in the series, the AI series of courses on OCLC, which is the 'AI Matters' course. The digital badge for this particular course is currently in production and will be ready very, very soon.

Finally, learners can also interact with all the courses they're enrolled in at the OU using the OU Study app. This includes the OCLC courses, and this provides access to learning materials on the go, the option to download these materials for later use, as well as reading and responding to foreign posts.

Thank you very much for your attention. If there are any questions at all, about any of the material from this talk, there's an email address [oclc@open.ac.uk] you're more than welcome to try to contact to get some more details from us, or of course, contact anybody who was presenting, so Mirjam or myself, about any of these details that you would like to follow up. And there's more so more information here about the short courses as well [<https://www.open.ac.uk/courses/language-short-courses>]. Thanks very much for your attention.

Mirjam Hauck

Thank you, Andrew, another round of applause. Thank you. That was again a whistlestop tour, but this time around a new short course. We have time left for questions for both John and Andrew.

Julia Molinari

Hello, everyone and thank you both, John and Andrew, for a really, really clear and helpful explanation. My question is a really horrible question. It's about AI policy, and I'm asking it, because I'm being asked left, right and centre about what the OU AI policy is. I get this from students at the Graduate School with whom I work, I'm getting it from colleagues. I don't expect you to have an answer. I sort of know what the answer is, but I just thought given that we're all here together, that you might have an update on when you will be able to communicate its approach. Other universities have kind of got theirs up there - not all of them, but others have, and at the moment, my strategy is to collate what I think is

reasonable from what I'm seeing in other universities to pass on as good practice. But I just would like to be able to say, 'this is the OU stance on it'.

John Domingue

The first thing I would say is it's not a horrible question at all; I was expecting something far worse. So, again, I don't want to be known for this, but I think that the colleague at the OU that is the expert on OU policy is Mychelle Pride. She is changing roles, but she led the Task and Finish group on OU policy. The OU has published some statements, but maybe not the one that's most useful for you, and there will be others in the pipeline. So Mychelle will give you a definitive answer.

Mirjam Hauck

Yeah, time for one more question. Ian, please.

Ian

Yeah, I had a quick one. It's about Grammarly, because there's some aspects of Grammarly that provide AI support and therefore, are we discouraging students from using Grammarly? Because it's got an AI content or is there a point where some AI is good, some AIs aren't so good?

John Domingue

A fabulous question. So, with others in the group supporting Mychelle and writing the policy, we looked at other university policies, and some university - I can't quite remember, and I wouldn't name if I could - had a really stupid policy if I may say that technical term. They said students aren't allowed to use any AI, even in the background research to their assessments. Now, if you watch online programming, if you bank online, if you listen to music online, if you use Google search, you're using AI. So, AI is embedded everywhere so you cannot ... it would be ludicrous to ban AI. I think it's - for my own purposes, in my personal view, this is my personal view - you have to think about what role did the AI play in generating the output and making that role transparent? It's almost the same discussion you have when you're writing a paper, and you know those discussions you have about the author or the order of the authors: what role did you play? So, for me, it's like that.

Ian

To just extend it slightly, then, I mean, it's to support a dyslexia student, I need the same sort of support. Word has the editor, is that the same sort of thing?

John Domingue

Good question, I guess it quickly becomes very grey, but again, my personal view is just encourage students to be transparent, and also if they can be transparent about their

particular context. I know for example - it's slightly different ,but students whose language, first language, isn't English, have been getting caught out with these AI spotters, because sometimes the English as written by someone who isn't a native English speaker somehow looks like an AI. Today, in fact, using an AI will not be caught out by IT. So, I think in a sense, the problem hasn't really changed, but it's always about being transparent, about who did what. But I'm sorry, I'm not more helpful than that.

Ian

No, that's good. Food for thought. Thank you.

Mirjam Hauck

Okay. I think now we have two minutes left and I want to hand over to Qian. Are you still with us?

KAN, Qian

Yes I am! I, just want to say our next joint event with our OCLC and OU Online Confucius Institute is a talk by Dr. Jinying Zhao, her Chinese name is ZHAO Jingyi - from Cambridge University's Needham Research Institute, and she's also Needham Research Fellow at Clare Hall, Cambridge University. So that's the topic [The Good Life: what ancient philosophers can tell us about living well]. I'm going to put the link to the event there.

[\[https://www.eventbrite.co.uk/e/the-good-lifewhat-ancient-philosophers-can-tell-us-about-living-well-today-tickets-700803400837\]](https://www.eventbrite.co.uk/e/the-good-lifewhat-ancient-philosophers-can-tell-us-about-living-well-today-tickets-700803400837)

It's going to be on 17th of October at one o'clock. So if you go to the link all the details about the talk, that blurb, how to register is on that page. And hope to see some of you there. I know it's a totally different topic.

Mirjam Hauck

It's a joint Open Centre and Online Confucius Institute event.

Thank you, everybody, for coming. Thank you, John and Andrew.

This has been recorded and the recording and the transcript will be available, you will get the information, and we look forward to welcoming you on October the 17th to our next event.

Bye for now.

End of transcript