Open Centre for Languages and Cultures:

Distinguished Speaker Series

Monday 21 November 2022

Edited transcript

Presentation 1: Dr Heriberto Cuayáhuitl, School of Computer Science, The University of Lincoln: *Robots Interviewing Patients*

Andrew Gargett

For this session, we are delighted to welcome Dr. Heriberto Cuayáhuitl, a senior lecturer in computer science in Computer Science at the University of Lincoln. He received a PhD from the University of Edinburgh in 2009. He has an international research profile in academia and industry in the discipline of machine intelligence, carrying out work in dialogue systems, machine learning robotics, and has published over 80 research papers in these areas. He has served in a variety of research leadership roles, including as lead organiser of an international workshop on machine learning for interactive systems and as guest editor for the journals, *ACM Transactions on Intelligent, Interactive intelligent Systems* and Elsevier's *Computers, Speech and Language*. In addition, he has an extensive background in industry with placements at *Speechworks International* - now, *Nuance Communications* - as well as the *German Research Centre for Artificial Intelligence*, and also *Samsung Research*.

We are very honoured to host Heriberto's talk today on the future of robotics in healthcare. Please, Heriberto, take it away!

Heriberto Cuayáhuitl

Thank you, Andrew. It's my pleasure to talk to you today about this particular topic: robots interviewing patients. Basically, the topic is, as we can see in the in the background of this image, a human robot talking to a human patient about health topics. One of the main reasons for me to try to attempt to create a robot that can talk to humans about health is the shortage of healthcare staff that we're having here in the UK. It is predicted that in the years to come we will have about a quarter of a million less staff than expected We can see the trend in this plot:

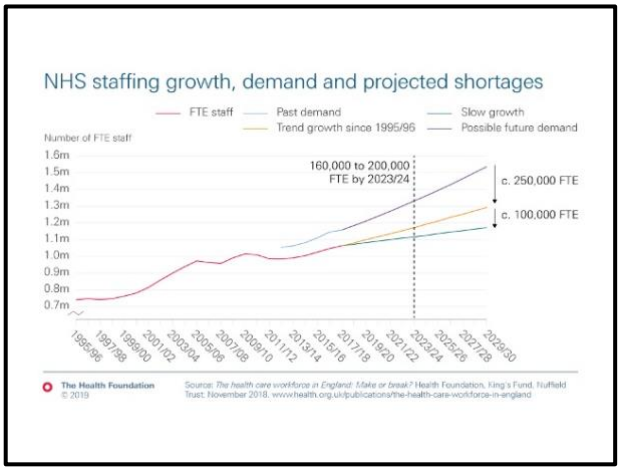

So here we can see how staffing has been developing over the years since 1995. The yellow line is the expected growth, the purple line is the potential growth that is expected, and the gap in between these two lines is about a quarter of a million by 2029. So, it looks like there is a substantial shortage of healthcare staff. And that is just here in the UK.

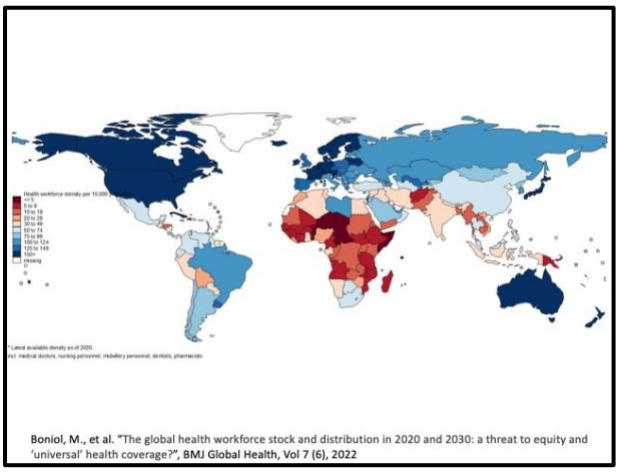

Let's have a look at this other plot: This plot shows a healthcare staff per 10,000. people.

The darker blue shows about 150 plus staff per 10,000 inhabitants, and those countries in blue or light blue, they are in a good situation. Those countries in red, light red, dark red, they are in a less good situation where there is a substantial reduction in healthcare staff per 10,000 inhabitants. So if we look at the worst case scenario, it gets up to less than five healthcare staff per 10,000 inhabitants so it can get real really, really, really serious when it's when it's time to think about the health workforce to take care of people. There is certainly something that needs to be done. It looks like all around the world, not just in developed countries, but also in under-developed ones, and in the less developed countries, the situation gets worse. So it looks like the message is, is this: 'Help Wanted', there's help wanted in all countries, whether they are developed or not. So, this is one of the motivations.

If I may tell you one more motivation driving my initiative on this topic: when I was a teenager, I had an accident, a car hit my leg and part of my knee got broken. When I was taken to the hospital, the doctor that was looking after me, his diagnosis was, 'we need to amputate his, his lower leg', from my knee, the lower part of my leg, I would lose that part. That was the diagnosis of this doctor. My parents took me to a different hospital, and they had a different diagnosis, and they didn't amputate my leg, I'm so grateful for that, that I didn't lose part of my leg. One of the things that I really enjoy is exercise, and when I go out for a run it always reminds me of that; we humans can make mistakes that can have devastating consequences. And unfortunately, sometimes our medical staff, obviously not everyone, but sometimes they get things wrong, and when they get things wrong, it can go seriously, seriously bad. And it looks like our healthcare staff could get some help in other ways, not just by having more healthcare staff, for example, but AI could also make them stronger, for example, in the diagnoses that they make.

So, there are the two reasons for bringing machines and AI to healthcare: one, to address the issue of shortage of healthcare staff, and two, to hopefully make diagnosis more accurate.

Okay, so, if I think about the idea of creating a robot or a machine that can take some tasks that a GP, a general practitioner, performs, what kind of tasks do they have? Well, it looks like the top skill is communication skills. Another one is working with people and another is working within a team. When I think about this, it looks like a machine that can communicate is crucial. It would be useful to have a machine that can communicate because patients, most of the time, they have to communicate, to express their situation to receive help. And that's why a communicating machine - or a communicated robot - could be a useful thing. So that part was the motivation behind today's talk. Now I'm going to start the talk - and by the way, please feel free to interrupt me whenever you want.

Okay, so I will tell you about four things: a proof of concept that I have created; data collection that I have done; some results that I have got; and some things that could be done as part of future work.

Okay, so for a prototype. So here, it is a little bit interesting because we have a situation that we call in computer science, a chicken and egg problem: I need data to create a robot system and by robot system, I mean the software of the robot, the software running on his brain - the robot itself is just the hardware. So, we need these two things: the hardware and the software, so we can have a thing that can communicate with patients, right? I need data to create my robot system, but I need I need a robot system to generate the data. So that's why we call it the chicken and egg problem. So how can I solve this problem? To address that, I have created a prototype that uses teleoperation. In reality, I still have a medical doctor talking to a patient, but in the middle, there is a robot, and these three entities are in separate rooms. The medical doctor is in one room independently talking to some computers, and the robot and human they are in a different room. The patients don't know that they are talking with a human being, they only know that they are talking to a robot.

We call this teleoperated robots because the robot's task is basically just to mimic what the human is doing. What that means basically, is it has to say what the human is saying. And, and it does that by speaking, by sending images and speech to the medical doctor and it's also moving its body. The medical doctor can see the robot – the virtual robot - and he can see also the patient that he's talking to. He knows when the robot is talking and when the patient is talking. He also can hear everything that the patient is saying and he can also hear everything that robot is saying. So, basically, the medical doctor gets into the body of the robot. That's basically how this works. And it's quite interesting that once one wears headphones, one really feels like being inside the robot because we can hear and we can sense what the robot is seeing and saying. This is the scenario that I used for collecting data. Because in the end, I'm interested in creating my robot system.

So today I'm talking to you about the data that I've collected, but I've already used our working prototype because my recruited patients are already talking to a real physical robot. So, the robot streams audio and video to the to the patient and to the medical doctor. And we use speech recognition to recognise everything that the medical doctor says, and that is sent to the robot so that the robot uses its own voice to say sentences to the patient.

Okay, so this is the procedure that I followed to collect data: I recruited some participants from the medical school and some others from the College of Science. Many of them come from Computer Science but also from a few other schools. So, there were 26 native speakers of English and 27 non-native speakers of English because I did this data collection with a couple of medical students, and we wanted to see if there were some differences between and native and non-native speakers.

They were provided with consent forms so they could approve the data being used for research purposes. They were provided with a patient profile; they didn't use their own data. They were given a name and date of birth, for example, the symptoms that they had experienced so they could talk about that during the interview. Participants were not instructed at all on how to interact with robots. Something that I realised later is that some of them were very, very inquisitive about that, asking, 'well, what am I going to do? How's it going to go?'. But it turns out that later they discovered it' just a chat as if they were

meeting with a GP – well, almost. So it wasn't anything very sophisticated, it was just a chat, as you will see in a moment. After the interview the patients filled in the questionnaire, and they provided some quantitative and qualitative data that I will try to summarise in a moment.

Now I'm going to play you a video. Just a couple of comments before I play the video: I don't have captions but don't worry about that because I just want to give you a bit of a sense, an idea of how the medical doctor was actually in the head, in the body of the robot. The robot is moving its body, is leaning sometimes, is moving his arm and moving its head. You will see that during the during the video.

[The video is played at this point and the soundtrack is unintelligible]

Don't worry if you cannot hear much. The idea is that what you're seeing is what the robot was saying, and when you see the camera moving it's because the robot has a couple of cameras in its head and it's using one of them to look at the patient. And yeah, the robot says something in its own words. Sometimes the robot doesn't say exactly what the medical doctor said, but most of the time it does.

There are only some minor errors in the words that the medical doctor says and what the robot is actually saying. If anyone is very interested in this type of material, I could show that later on, or you could get in touch if you wish to see these videos.

Okay, so now let me go back to my slides. I collected about seven and a half hours of interviews. It doesn't really sound like much but this about 26,700 seconds, 5.5 gigabytes, 854 megabytes, over a quarter of a million of images or video frames and the dialogues have on average about 590 words; they are rather lengthy.

The robot speaks a little bit more than the patients, about 355 words versus 232 for the patients. If I look at the total set of sentences, there were 2,137 unique sentences that were spoken in the data. So that number shows that patients were not really reading the patient profile, they were explaining in their own words, what they thought their situation was, their health situation.

Now, the questionnaire.

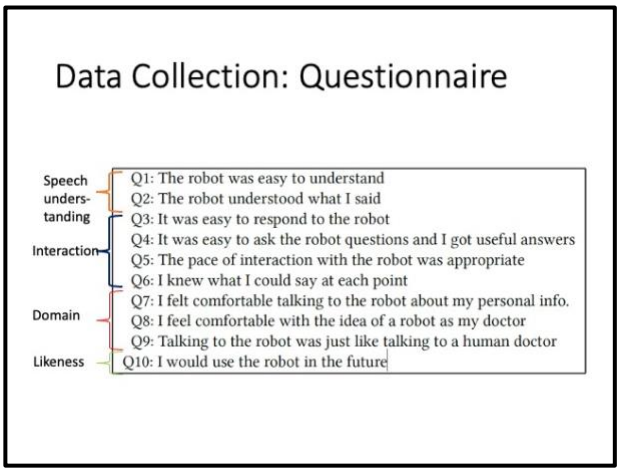

Once the interview was over, the recruits - the patients - received this questionnaire and they had to answer it using Likert scales with five points. First question - well, there were two questions regarding speech understanding: 'The robot was easy to understand', and they could agree or disagree with that, with five numbers. 'The robot understood what I

said’. Regarding interaction I had four questions: ‘it was easy to respond to the robot’; ‘it was easy to ask the robot question; ‘the pace of the interaction with the robot was appropriate’; and ‘I knew what I could say’. Now regarding the health domain, I have three questions for that: ‘I felt comfortable talking to the robot about my personal information’; ‘I feel comfortable with the idea of a robot as my doctor’; they feel like they are talking to a human doctor and to what extent they don't; and the last one is about likeness: ‘I would use the robot in the future’.

Every participant we recruited answered these questions, and here are the results.

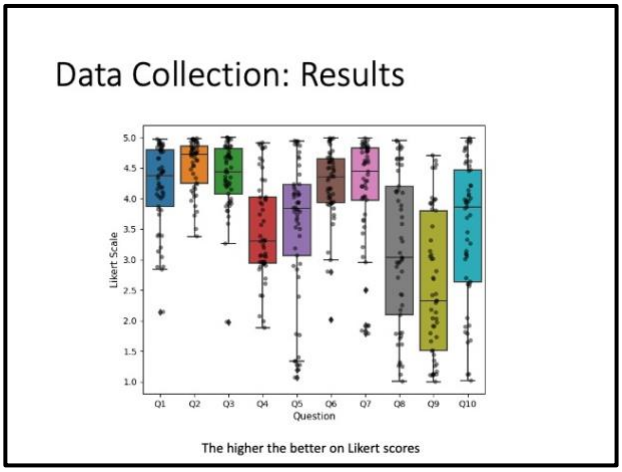

We have these two metrics that assess speech understanding, and they are both very high. That's not surprising, because ultimately, participants are talking to a human, but it's also good to know that they are able to recognise that as they are able to understand pretty much what the robot is saying. So that's very good. Regarding the interaction metrics, it was easy to respond to a robot. Let's see, matrix q’s four and five, it was easy to answer the robot’s questions, the pace of the interaction with the robot was appropriate. So, these two metrics, well, definitely they are not as good as the previous ones. The pace definitely can improve. Ultimately, I think it was easy to ask questions, but this metric – q 4 – has to be seen in the light of the interactions, because some of the recruited patients answered this question neutrally, because they didn't ask any questions. So, they say, ‘well, I didn't ask any questions, so yeah, I am not in favour or against’. But it looks like q five is definitely one where we can improve. And six and seven: ‘I knew what I could say’; ‘I feel comfortable talking to robots’ are good, are relatively high. Q eight and q nine are the lowest and that is ‘feeling comfortable with the idea of a robot’ and ‘talking to the robot was like talking to human doctor’. So certainly – yeah - we knew that. We have a machine that's clear, it is different talking to a machine than talking to a human. But this was quite interesting because some people are already biased before the interview, they look like they are already biased: ‘I just don't think that I want to talk to a robot’. But it looks like there is a division because look at this, there’s about half in the upper scale and about half in the lower scale, so it's a bit divided. When we finally look at whether people would use the robot in the future, it looks like most of the opinions are in the ‘yes’. Most say, ‘yes, I could use these robots in the future’. Some people are leaning against that, but there are some advantages that were mentioned by people. People think that talking to a robot about personal information is much simpler, they feel more comfortable interacting with a machine than with a human. When we look at our results, 70% ‘agree’ or ‘strongly agree’. So, these types of interactions are relatively successful. I have lots of numbers here, but I'm

not going to comment on the details. I want to skip these details and tell you a little bit about the qualitative analysis.

People provided some feedback in free text form and here you can see some of the questions:

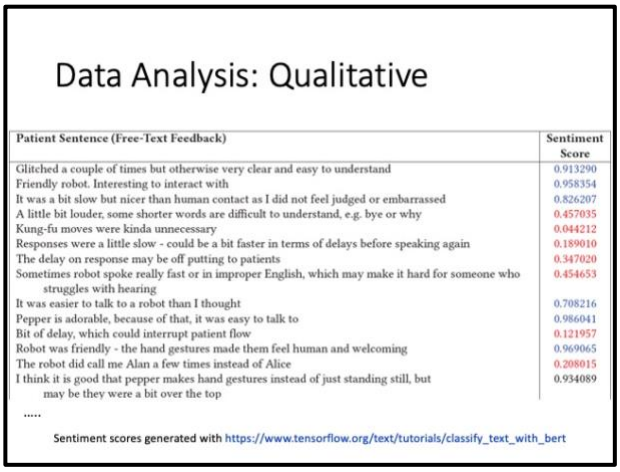

Data Analysis: Qualitative

Patient Sentence (Free-Text Feedback)	Sentiment Score
Glitched a couple of times but otherwise very clear and easy to understand	0.913290
Friendly robot. Interesting to interact with	0.958354
It was a bit slow but nicer than human contact as I did not feel judged or embarrassed	0.826207
A little bit louder, some shorter words are difficult to understand, e.g. bye or why	0.457035
Kung-fu moves were kinda unnecessary	0.044212
Responses were a little slow - could be a bit faster in terms of delays before speaking again	0.189010
The delay on response may be off putting to patients	0.347020
Sometimes robot spoke really fast or in improper English, which may make it hard for someone who struggles with hearing	0.454653
It was easier to talk to a robot than I thought	0.708216
Pepper is adorable, because of that, it was easy to talk to	0.986041
Bit of delay, which could interrupt patient flow	0.121957
Robot was friendly - the hand gestures made them feel human and welcoming	0.969065
The robot did call me Alan a few times instead of Alice	0.208015
I think it is good that pepper makes hand gestures instead of just standing still, but may be they were a bit over the top	0.934089
.....	

Sentiment scores generated with https://www.tensorflow.org/text/tutorials/classify_text_with_bert

I used an automatic sentiment analyser to score each of those sentences, the higher the number the better. On some of these comments in red, people were definitely commenting on things that can be improved, for example: ‘a little bit louder could be useful’; ‘some words were difficult to understand’; ‘sometimes the moves were unnecessary’; ‘responses were a little slow’. Yes, we know that there is a little gap between what the robot says and what the patient says and what the robot responds to. There's, there's a little gap in between, because the human is speaking in between and their spatial condition as well – the speech was being streamed to the robot. Overall, we got .62 or .63 - we can round it to a 63% sentiment score, and it looks to be more on the positive side, although there were opinions that gave us room for improvement, where future prototypes can get better.

There were lots of positive comments, and that's very good, that was really good to see. In the future, I would like to use a fully autonomous robot, using a technique in AI called ‘reinforcement learning’. I'm not going to explain how to do that, because that is a separate talk. I would be happy talking about that in the future if you wish.

Can we do better than the results that we have got? The answer is ‘yes’ and ‘no’. Yes, in things like the slowness in the interactions - we can get that quicker. We cannot do much about the fact that the machine is a machine. We can’t really make it more human. We could make it a little bit more natural, possibly that could be improved a little bit. But in terms of recognition, that's one of the challenges. We had a human during the recognition, the language, understanding. If a machine has to do that, then it wouldn't be as good as a human. Arguably, at least we would need more data to match that performance. Matching this performance is challenging, improving others that are low is certainly possible. But I'm looking forward to seeing whether we can get more people on board with whom I could use this type of machine or systems in the future. There are a number of things that we can do in the future, such as extending the interviews, taking vital measurements, providing medical diagnosis. In the future, I could see a robot not just interviewing a patient, not just a static robot, but a robot that can move around the hospital, monitoring patients, bringing in food to the patients, and so on and so forth. And for that, we need more data.

I'm going to stop here.

Andrew Gargett

So firstly, thank you very much, Heriberto. That was a wonderful talk, I very much enjoyed that.

Okay, so we have a good five minutes or so for questions. Maybe I could start on the question side of things. There are so many questions I'd like to ask, and one in particular. So, the situation with the robot asking more, or having, it seems, more contributions or more terms, at least using more words than the patient? There's a couple of things there: one of them is a question of whether maybe that was a feature of the particular questions that the doctor was prompting the robot with? Or on the other hand, maybe polar type questions, 'yes', 'no' type questions would tend to just elicit shorter answers, maybe elaborations if the patient felt like it, and the other thing is, would that be then something characteristic of doctor-patient type interactions, and/or is that just simply something designed into the robot?

Heriberto Cuayáhuitl

Here you can see an example transcript of the actual video that I tried to show. So sometimes the medical doctor could be a little bit verbose. And actually, I noticed because I used two medical doctors, and one of them was more verbose than the other. Actually, in the slides that I didn't show, when I looked at the number of words, robot words, for example, this doctor spoke about 50-60 words more - well about that 50 plus more words, - than the other doctor: one doctor 384, on average, and the other 328. So, the reason here is that some doctors could be a little bit more verbose than others, but they are trying to do two things. On the one hand, they are trying to make some acknowledgment of what the patient is saying, and then asking questions. Putting these two together, they seem overall, to require more words than the actual patient because the patient only focuses on answering the question. That's why the patients speak less.

Just to take the advantage of this slide of this page that we can see here, sometimes my robot misrecognised what a medical doctor said. So, for example, the medical doctor said 'ward', but the speech recogniser recognised it as 'world' and the speech recogniser picked 'about' when the doctor said 'brought', or the doctor said 'A&E' and the speech recogniser said, 'a knee', which sound very similar, somewhat similar, when all these words are being streamed together. Sometimes it was a little bit strange, but I think that overall patients were able to understand what the robot was saying, as judged from these results that we can see here: 'they're always easy to understand' – q 1. The results are very high, substantially high. Again, there's room for improvement. Maybe it was due to those initial conditions. But, yeah, it looks like the doctor is trying to do two things to answer your question and the patient is doing one.

Andrew Gargett

Okay, thank you very much for that. There's actually a question in the chat. So Jess asked, 'in how many languages will the robot be able to communicate?'

Heriberto Cuayáhuitl

That's a really interesting question. Thank you for that. Well, these days, it could be in any language. And something that I didn't mention, and it might be interesting, is this robot can be a virtual robot, it can be on mobile devices, given that now so many people have mobile

devices. These robots, if we wanted, it could run on your mobile device. The virtual robot could run on a TV, I could be in front of my TV and having a consultation. This is something that, that I can imagine happening very soon in the future. I'm going to go to my GP, well guess what you don't need to, you can go to your living room, and then you are certainly, talking to GP. But on the question about languages, what is very interesting is that these can be imported to so many languages around the world, and those countries in need of more healthcare staff, they could receive quite a lot of help by having some automated systems that could sometimes take the role of a GP. Not all the time, because sometimes, there is the need for a human touch, or there are more complex situations the machine wouldn't be able to cope with. Nonetheless, for many situations, it looks like machines could be able to interact with humans and it could be in any language. Once we have that in English, we can import it to many different languages.

Lucy Moss

Andrew, I think we've got time for just a couple more questions. So we've got hands up from Frank Monaghan and Rosina. So Frank, do you want to ask your question?

Frank Monaghan

Thank you very much for the talk, Heriberto, it's really interesting.

I've got two questions: would you imagine working with linguists on this to look at the difference between the human interaction and the robotic one to see whether there are things there? Because we know from other research that metaphor, for example, is an incredibly important part of how patients and doctors communicate and, obviously, robots might have more difficult with that. But the other thing that was playing on my mind was, we seem to be assuming that the robot needs to get more human. I was just wondering the extent to which we need to get humans more used to the qualities of the robots so that people might have more confidence that the robot knows more, it may have more diagnostic capabilities than a human can keep in its head. So, whether we can convince people about the advantages of the robotic. You mentioned, using tests, so they could take blood pressure, they could take other measurements and so that would enhance confidence in the machine that might push up some of those lower scores that you were getting on q eight, q nine.

Heriberto Cuayáhuitl

Thank you for your questions, Frank. So, these days AI is relying a lot on data and data may come from text data, or maybe numerical data, that we can use to predict diseases for example, but it looks like the current fashion of AI is to gather as much data as possible so that we can create systems that can behave as the decisions that are being made in the data.

With regard to linguists, I think that there is an opportunity for bringing in different experts, because it looks like although the conversations may be [*unclear*] from the behaviour of servicing the data, we may want to add some constraints to the interactions. And that's where the linguists or domain experts may come in to sometimes specify the behaviour of the machine to give it a more human touch, as you were commenting. So it looks like these type of AI systems could have some sort of hybrid system, one that is driven by the data and another that is determined by the experts, whether they are linguistic experts or medical experts.

Andrew Gargett

Rosina, did you want to ask you a question now?

Rosina Marquez-Reiter

Thank you, Heriberto - a very interesting talk indeed.

I wanted to talk to you a little bit about the transcript, some observations on the little dialogue that you showed. It seems that the patient seems to be orienting to the interaction with the robot, if you like, as a test-to-learn situation, as if there was almost a written questionnaire, if you like, with the exception of what he said that he smokes, and he shouldn't be smoking, which are cultural 'no no's', right? So, in terms of the number of words that you are saying the doctor says, versus the patient, your difference might be there, amongst other things, because normally, doctor-patient communication has a lot more than 'yes, no' responses that you get from the human, apart from the hesitation device that you wrote down, which is, 'eh', right, which is quite interesting. So, they are orienting to the fact that they're talking to a robot already and learning how to talk to the robot how patients talk to a doctor.

Heriberto Cuayáhuitl

Yeah, yeah, certainly. These types of interactions, they are on the one hand … we know in advance that medical doctors are trained to interview patients in some particular ways, but in the end, it looks like humans choose their own way of verbalising things. That's not surprising. With regard to patients, well, it certainly would be useful to have a wide variety of patients in order to see the differences between different groups of patients, because for example, our 'patients', unfortunately, were recruited from the university and they are mostly university students. They wouldn't be representative of the entire population of patients and certainly they may be a little bit biased by the patient profile that they were given. Ideally, we would have wanted to have real patients, but for the moment, we want, or personally, I would like to create a prototype where I can say, you know, this thing really works, can we trial it in a in a hospital or a clinic? But yeah, these times are still ahead and hopefully, even if it's not me, if it's somebody else, hopefully this will happen in the future for the benefit of, of society. I hope that I'm trying to address your question a little bit. Thank you very much.

Andrew Gargett

Sorry, to come in, but thank you for the question to us and to everybody else. So, can you all just please join me in thanking Heriberto again for his talk.

Heriberto Cuayáhuitl

Just one last comment. If there is anyone that does have an additional question, please feel free to send me an email, it would be nice to hear if there are any more questions if you want to, to know more, or to discuss more. Thank you very much.

Presentation 2: **Dr Andrew Gargett, School of Languages and Applied Linguistics, The Open University: *AI Matters***

Andrew Gargett

So, this brings us now to the final part of the proceedings. Welcome to the final presentation on *AI Matters*, which is a course in a series of courses, *AI in the Workplace,* for the Open Centre for languages and Cultures at the Open University.

Overall, the aim of this course is to help you the learner to acquire a critical appreciation of AI technology, to develop awareness of the risks of such technology, and ways to reduce these risks. The course has several specific aims. If you were to join us a learner, you'd be able to formulate your own definition of AI by the end; become aware of the major historical, social, political, and economic issues in AI; be able to critically evaluate primary source information on these issue and formulate your own responses to these issues; and identify key features or risks in AI projects and practices using these features to flag risks in actual AI projects.

The course presents ideas and views of prominent AI researchers and industry leaders keeping the material up to date in this way and highly relevant to the current challenges of AI technology.

There are eight units, each estimated to take four or five hours of study, so that's approximately forty hours total for the entire course. Each unit includes interactive tasks, as well as unit quizzes at the end to test understanding, and there's also the option to engage in forum discussion with others in the course - this is moderated by a learning advisor.

The course is open for between six to eighteen months once you register, depending when you register that is, and to allow learners to complete at their own pace. Once you've joined the course, you'll have read-only access for up to three years, and the total cost for the course is £195 p. On this slide here, you can see the list of available units.

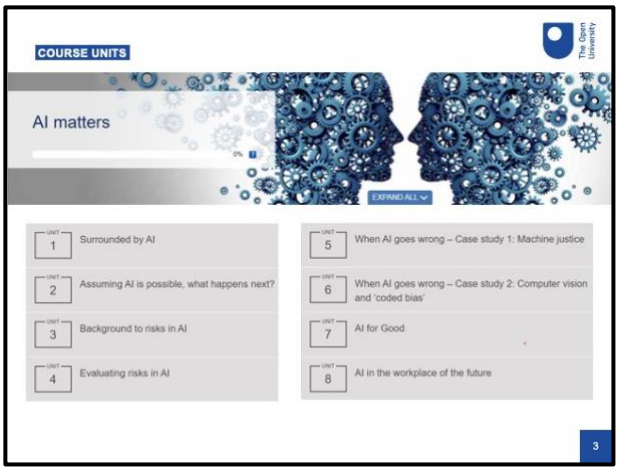

The overall flow of the course as follows. After providing suitable background to AI technology as 'risks' - units 1,2, and 3 - the course moves into considering responses to these risks from unit 4 onwards, introducing learners to an ethics evaluation toolkit from the UK Information Commissioner's Office, and there are guides on how to use this for AI projects. A series of case studies follows in units 5 to 7, from the use of AI in the criminal

justice system in North America to bias in computer vision systems. And finally, examples of AI being used to benefit society. A final unit explores the impact of AI on the workplace now and into the future for a general course wrap-up.

So, the interactive activities. Throughout they explore complex and very contemporary challenges to AI technology, using mixed media such as online talks, in-page activities, and finally quizzes to check progress.

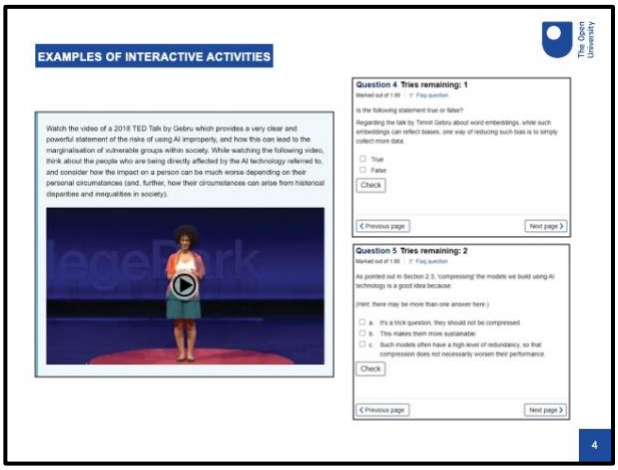

Here, we can see an example of one of these activities from unit two. And this examines the work of MIT graduate and former Google employee, Timnit Gebru, who's work is on exposing bias in computer vision technology. On the right, you'll see also some examples of where the learner's understanding of these activities are later tested through these quizzes.

Here's another example of interactivities from Unit 7:

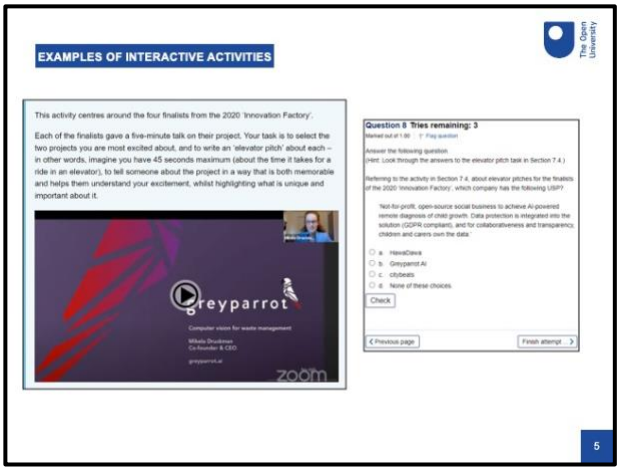

In this example, learners are guided through the creation of an elevator pitch inspired by finalists of a UN supported competition that focused on solutions to global challenges facing contemporary society. You can see here the presentation from the winner of this competition, Greyparrot, for their project on using computer vision for waste management. And on the right is listed an example quiz question testing the understanding learners gained from this activity.

In addition, as already mentioned, the course has both in-page forums and course forums that are moderated by a learning advisor. The forums enable learners to interact with other learners who have different experiences and perspectives to offer. And you can see here an

example of learners interacting with the learning advisor on the theme of myths of AI, which is from Unit 1 of the course.

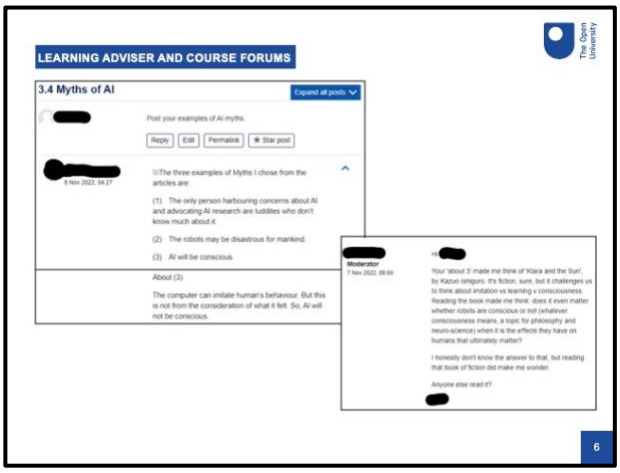

In addition, there are in-page links to forums also moderated by the learning advisor and this enables learners to react immediately to course content. This example here shows discussion between learners and the learning advisor on the theme from Unit 2 about mechanical people being presented throughout history as artificially intelligent.

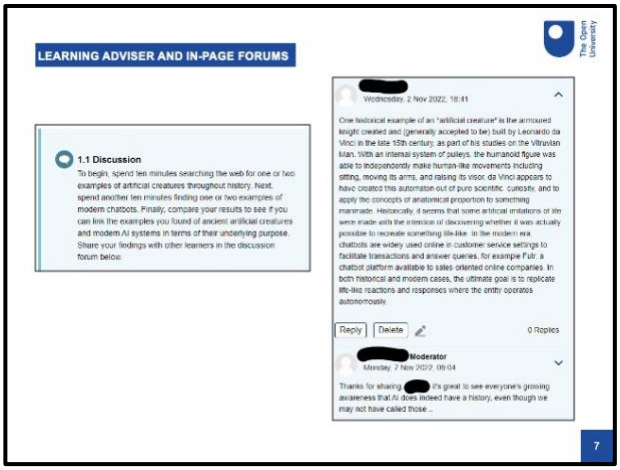

This theme is of course, highly relevant to our modern understanding of AI technology, particularly expectations around technologies such as chatbots. And dare I say there may well be some connections to some of the material that Heriberto was just talking about, too, which should be very interesting to follow up.

Also, learners are awarded digital badges at the end of the course for getting 70% or more on each of the unit quizzes. These badges can be shared on social media or added to email signatures, and can be used for continuing professional development, job applications, promotion, and the like.

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

Many thanks for your attention and if you have any further questions about the course, or indeed anything that arose during today's talk, today's session, then here's an email address: oclc@open.ac.uk . Also, here's a link to *AI in the Workplace: AI matters* as well:

https://www.open.ac.uk/courses/short-courses/lg003 . And finally, if you would like to find out more about the range of courses in the OCLC, then please visit the link: https://www.open.ac.uk/courses/language-short-courses.

So that brings us to the end of today's session. If anybody does have some questions they would like to ask about the course, please feel free to use this email to get in touch with us. And I'd like to, to thank Heriberto once again for his wonderful talk. And also to thank everybody here for their attendance.

And finally, we do have a final announcement regarding upcoming talks as well. Lucy, did you want to take over at this point?

Lucy Moss

Yeah, just to mention that we've got the **next talk in the series on the 12th of December**, which was rescheduled. It was due to take place earlier in the year, but it coincided with the Queen's funeral. So this is our **UNESCO chair, Professor Alison Phipps**, our colleague in the University of Glasgow doing a talk on ***The Languages of Crises*****,** and that's now open for registration: https://www.eventbrite.co.uk/e/distinguished-speaker-series-alison-phipps-and-the-languages-of-crises-tickets-470951798507

Thanks for coming, everybody!

Andrew Gargett

Many thanks again for everybody coming along; it was a wonderful session. Thank you