Westminster Abbey
Learning

The Theology of Artificial Intelligence Q&A – Religious Education Masterclass 2025 teachers' notes

Thank you for downloading this resource. We hope that it will be a useful teaching tool in your classroom.

As we continue to grow our free catalogue of teaching resources, we'd really appreciate a few minutes of your time to let us know what you liked and what could be improved. Please complete this five-question survey.

Watch the Q&A here: https://www.youtube.com/watch?v=suypEMuysTM

Please note that all opinions are of the individual speaking and do not represent the views of Westminster Abbey. Questions in italics were asked by post-16 students.

Questions for discussion

Watch the recording of the question and answer session between Professor Andrew Davison and Laura Arends, Head of Learning at Westminster Abbey.

Question 1 – *beginning at [00:02:05]*

The media has recently speculated that there may be a chance to bring back our loved ones through AI. Does the resurrection of the dead using artificial intelligence go against the will of God?

In response, Professor Davison says that recreating a loved one through AI is not true resurrection, but a "good imitation". He emphasises that Christianity, Judaism and Islam teach resurrection of the body, not a digital likeness. To what extent does creating an AI version of someone who has died challenge or support religious ideas of the resurrection? Is Professor Davison right to say that "there would be no real permanence"?

Question 2 – *beginning at [00:06:10]*

Do you think that it is reasonable to believe that AI will be able to apply emotional intelligence, morals and belief systems to scenarios and actions in the future, and would this be more reliable than human responses to situations?

In reply, Professor Davison explains that AI is not truly intelligent or emotionally aware, but is instead a tool for data analysis. He says, "I'd rather have a clumsy real person with good will than a brilliant computer analysis". Can AI ever meaningfully replicate human emotional intelligence or moral understanding, or will it always lack the "person to person insight" Professor Davison refers to?

Question 3 – *beginning at [00:13:20]*

Do you think that there could be a time when religion succumbs to AI? Could AI become so powerful in terms of knowledge about the world and its creation that people will view religion as irrelevant or outdated?

Professor Davison responds that religion has never aimed to replace science and that true religious insight values kindness, creativity and personal relationship over raw knowledge. Do you agree? Does the growing knowledge and power of AI threaten religious belief, or might it actually deepen our sense of the mystery and value of faith?

Transcript

Speaker:

Professor Andrew Davison, Regius Professor of Divinity, University of Oxford

Chair:

Laura Arends, Head of Learning at Westminster Abbey

Introduction:

The Very Reverend Dr David Hoyle KCVO MBE, Dean of Westminster

Laura Arends:

Hello and a very, very warm welcome to you to Westminster Abbey's RE Masterclass this afternoon, on the Theology of Artificial Intelligence. I'm just going to invite the Dean of Westminster, Revd. David Hoyle and he's going to just say hello as well.

David Hoyle:

Hello everyone. Very good to have you at this Masterclass. Westminster Abbey is a working church with four services every day, and it's a lot of other things besides: a shrine with its own saint; a place of burial and memorial; and a gathering place in good times and bad. It's a place that's full of stories and not all of them are good. It's a serious place where we need to do some serious thinking sometimes. So, we are really delighted that you're here to think about some hard questions. You're in the best of company: Andrew Davison is the Regis Professor of Divinity in Oxford. That means he occupies one of the oldest and most famous professorial chairs in theology. He's a man of deep learning and yet he manages to carry all that very lightly, which is quite a skill. I've known him for many years and I always enjoy his company. I think you will. I hope what follows helps you with hard conversations about the things that matter. I hope it's serious and stimulating.

Laura Arends:

Thank you.

Andrew Davison:

Thank you very much for the introduction.

Laura Arends:

Good, well we're very glad to have you with us this afternoon, Professor Davison. Now, we've got some amazing questions that have come in from the students who have watched your fascinating lecture on the Theology of Artificial Intelligence. I personally learned a lot from it and so much so that I actually had to transcribe it as I watched it and make sure it would all go in. It was really brilliant. So, we've got lots of questions and we're going to start off with one from a student called Ruthann from the Harris Academy Sixth Form.

Question 1 *[00:02:05]*

The media has recently speculated that there may be a chance to bring back our loved ones through artificial intelligence. Does the resurrection of the dead using artificial intelligence go against the will of God?

Answer

Andrew Davison:

Well, what a wonderful question. I'm really glad to start with that one because I think it's very significant. Although I would say, probably more for pastoral reasons for human well-being than necessarily theology. And I'll come round to that at the end. So, we should always try to understand what we're talking about and this is indeed something that's been in the news quite a lot. The idea is of taking a general artificial intelligence chatbot, so something that's able to be spoken to in English, or any other language, and reply in pretty natural English. And then you overlay all the texts that you have from someone who's died. So, you might put books that they've written, or their diaries, or transcribe everything that they ever said on a video, and that gives you this collection of words that somehow captures their thoughts about things and that's as near as you can get to perpetuating their person. And you could ask it questions and it would answer back in the voice of the person who's died. So, it seems to me that that is full of pastoral questions and complications. So, for instance, is that really a good way of helping people with grieving and does it tend to help us not make progress with recognising that someone has gone? So, that will be worrying. And then there are all sorts of questions like, who owns this thing and would they be able to switch it off against the wishes of the loved ones? Or could they keep it going against the wishes of the loved ones? And how does it work in terms of money? I said in my lecture quite a lot of the questions here are really pretty old-fashioned ones about who makes money and who goes out of a job and that sort of thing, so you can imagine that someone has this amazing chatbot and then they say, well you can keep it going but you need to pay 10% more every month than the month before. So, all sorts of ways in which this could

be done pretty disreputably. But even if it was, I think, even if it was done reputably, I think there are questions about what is psychologically helpful.

And then that question about the theological aspect and the resurrection. I think it's important here to say what are we not talking about. We're not talking about the person still existing. We're talking about making a good imitation and when I think about my grandparents, who I love very much, I'm interested in the Resurrection of the Dead for them, not just for me. And there are these computer systems. There would be no real permanence. It wouldn't be the person. It wouldn't be Resurrection because Christianity, Judaism and Islam are interested in the Resurrection of the Body. They say we're bodily creatures; that's really important. And it wouldn't even be proper immortality because you'd only need a power cut or the Earth's not going to last forever. So, I think it's one of the most morally, pastorally significant things that's being proposed in AI at the moment. I don't think it really competes at all with theological ideas of Resurrection but it does raise all sorts of important questions about how this could be done responsibly and whether it could even be done at all responsibly.

Laura Arends:

That's very, very interesting. A lot of angles I hadn't really thought of there. Just to say to those of you who are viewing, if you want to come back to Professor Davison with any comments or thoughts, particularly if it was your question, please do put them in the chat, in the Q&A, and we will get that fed back to Professor Davison maybe a bit later on in the session. So, thank you, that was a really excellent question to start us off with. I've got another one here from Katherine this time, from the Sandy Secondary School.

Question 2 *[00:06:10]*

Do you think that it is reasonable to believe that AI will be able to apply emotional intelligence, morals and belief systems to scenarios and actions in the future and would this be more reliable than human responses to situations?

Answer

Andrew Davison:

Well, I have to say these are turning out to be terrific questions. I think that AI can be applied whenever there's data to analyse because, although we get very excited about it because it can do things this year it couldn't do last year, it really is an exercise in data analysis or statistics. It's just that we can do that more powerfully than we used to be able to. So, if there's ever a situation where you can get clues about emotional states that can be recorded in data, then a computer could analyse it. And there are some examples I know of giving people questionnaires and an AI system is better at picking up certain sorts of psychiatric problems or risks that someone could be a danger to themselves. It's actually better at doing that than even a trained psychotherapist. I think I read about that in an article in *The Atlantic*, so that in a way that's no different from interpreting an ECG when you have your heart waves recorded. AI is now better at analysing that than a human doctor is most of the time. And it turns out that there are ways in

which you can answer questions that give away information about your emotional state. The information is there in the data. Computers turn out to be quite good at finding the patterns. But what I think is important to say there is, “Is this intelligence?” The question was about emotional intelligence. Well, the thing about AI, despite the name, is it's not intelligent. There's no *thing* there that knows anything. And in the same way here there's no *thing*, there's no *person*, there's no *consciousness* that has a sense of people's emotional state. It's just data analysis and you could say that it's even further away from what we think of as being a person's response to a person. No. There are kind of maths problems, data analysis problems, that we might be good at or not good at, and a computer can usually beat us to it. If you turn emotional questions into data it can be trained to come up with a good response but we might think that that's even further away from actually having sympathy or recognising what's going on. In that way that's quite difficult to put into words.

There's also a question there about morals, about whether something could be trained to have a kind of moral sense. So, there's a way in which you could take textbooks of moral philosophy or theology and train a large language model on them. Or there used to be quite a tradition in the Catholic Church before the middle of the 20th century of enormous manuals of what to say to someone in a confessional. So, like every situation you can imagine and a sort of moral theological response to it that could be fed into a chatbot. I'm not sure whether it's ever been tried and it would be able to come up with responses within the parameters of what was in the books anyway. But what it wouldn't show is that kind of person-to-person insight and I don't want to preclude the possibility that there might be insights that could come from that kind of data analysis. But especially if you're in a morally charged situation, like you've done something wrong or your conscience is really afraid, and you want to get some comfort, I think that I'd rather have a clumsy real person with a good will than an absolutely brilliant computer analysis of moral textbooks. And on the moral question, I think I maybe said in the lecture, and if I didn't I should have done, or it's an interesting point, I think: these computer systems don't really do things. *We* do things with them. They don't *understand* but we can understand *with* them. They don't act of their own accord but we can act with them and they can be really useful tools, and I've written a bit about AI as a tool. Well in the same way, they don't have moral responsibility and the moral responsibility comes back always to the person who chooses to use it and I'm sure the law is pretty good at this. The law and the law courts and barristers and judges are pretty good at working out where responsibility comes back to. So, say you have an AI weapons system then and it kills the wrong person. I mean, if we think they can, it's legitimate for them to kill anyone at all but you use it in war and it kills a civilian, no one would ever be able to say, “Oh, it’s not my fault, the computer did it”. They would always come back to: did you, a person, make a mistake or was this system badly designed? And who's responsible for that? So, I think the moral case is really useful because in a situation of law you'd *always* have the responsibility coming back to a human being and I think that's because there's no ‘there-there’ when it comes to these AI systems, they're just tools and we can't shirk the moral responsibility. Ultimately, it comes back to a person.

Laura Arends:

Yes, the fact that you’d much rather talk to a person if you had an emotional problem rather than talking to a computer. I think it's very telling.

Andrew Davison:

Although I'll say on that, that I'm absolutely not dismissive of these things and I can imagine a situation where you might want to use a computer system to help get at the problem. So that article published in *The Atlantic*, an American newspaper, a few years ago suggested that there are situations where a really well-trained computer system can notice something and that that would be fine, but you already go to a doctor and fill in a questionnaire quite often it'd be just like doing that with even greater skilful computer systems. Or you might imagine like the church wanting to provide pastoral advice where it was a bit of a hybrid. You didn't have enough time to speak to everybody so you shortened the time but you also gave them access to... Well say someone was in financial difficulty, I think it would be responsible use of some of this technology where people could input some of their problems and it would begin to suggest some solutions and then you could talk it over with a person. So, I'm not against it completely used as a tool, but I'm not inclined to think it's going to replace the human for all the things that matter most to us as human beings.

Laura Arends:

Thank you very much. We've got another question from Sandy Secondary School, so thank you Sandy School, you sent us a lot of really lovely questions. This one's from Savannah.

Question 3 *[00:13:20]*

Do you think that there could be a time when religion succumbs to AI? Could AI become so powerful in terms of knowledge about the world and its creation that people will view religion as irrelevant or outdated?

Answer

Andrew Davison:

Well, there is a whole field of study which is about AI in religion, in religious practice and even sort of AI as religion and I'm not that kind of scholar, but I will point you towards the work of a British anthropologist of religion called Beth Singler. Beth Singler, she's in Zurich nowadays, and she's really at the forefront of thinking about AI in religion and as religion.

What I would say is, it's never been that healthy, I think, to imagine that religion takes the place of knowing about the world or that knowing about the world takes the place of religion. So if you think about science, you know, often these things are put in conflict: religion and science. But that's really a late 19th century myth. For much of history, the Church and other religious groups were great patrons of science. The people who *were* great scientists were very devout. These were just not seen as particularly in conflict. And at the beginning of his commentary on Genesis, John Calvin, the great reformer of Geneva, he says if you want to know about astronomy, don't go to the beginning of the Bible, don't go to the Genesis Creation story, go and talk to the astronomers. The Bible is not trying to be a textbook. It's about, he thinks, more important things but he loves science; he's a great humanist. Christian humanists are

interested in all sorts of forms of learning. So, I'd be inclined to say that healthy religious attitude has always been let's find out as much as we can about the world. Let's use all the tools we can to understand it. That's never going to take the place of religion and similarly religion is never going to try and take the place of that and it's not going to take the place of religion. So, I'd also say religion has tended to value the things that AI just can't bring to the table, like personal attention and personhood and we'd all say, I'm sure, religious traditions, that a single simple act of kindness is in a different league from all the knowledge in the world and no AI system is going to have an eye out of which can show an act of kindness.

I think also religions have been quite invested in creativity, especially in the Arts: poetry, music and so on. And there's a real question about whether AI can come up with anything genuinely new. Because what you do, you train it on the things that human beings have already done and then you say, do more of the same. I think there is genuine possibility of noticing some connections that haven't been noticed before and that might be a cause of novelty. It's possible maybe that you feed all the music of Mozart in and all the music of Beethoven and you get something that we didn't have before which was Mozart trying to be Beethoven or maybe we've had Beethoven trying to be Mozart because he comes after the time! We'll never know what it was like for Mozart to try to be Beethoven and maybe a computer could have a go at that, and it would be new. But generally speaking, AI just stirs the pot, it just reshuffles the cards and I'm not an expert on AI and creativity but I think there are some really interesting questions there. And in as much as religions are invested in humanity and creativity as being in the image of God, I think that's just something that AI isn't going to displace.

Laura Arends:

Excellent. This is a very short and pithy question here from Will from the William Farr school.

Question 4 *[00:17:19]*

Is AI making us lazy?

Answer

Andrew Davison:

Oh goodness! Well, I suppose the question for me to ask myself is, “Is it making me lazy?” And you can all ask yourselves, is it making you lazy? Well, obviously, superficially, it can do. It's possible to use it instead of writing an essay. It's possible to use AI to translate things rather than learning a foreign language, so in that sense it can make people lazy. But then the question is why would you want to do any of those things in the first place? So, if your purpose in going to university — and frankly paying a fair bit for it — is presumably because you want to learn something, then not only is it not very good at helping you learn, but if it takes the place of you actually doing the work on the essays, it's positively detrimental. It's the same with exams — even just pragmatically — if you're wanting to do well and, like here in Oxford, exams are in a hall for three hours without the aid of a computer (at most, you’re just typing on a keyboard), then it's again positively counterproductive to become reliant on AI, because you won't have it when you

need it, and it could just make you lazy. But you'd have a kind of odd view of what your aims were if you let it make you lazy.

But I always want to try and put in some positive words as well and I think it can be a genuinely useful tool, but in my experience, it helps those who already know something. It can be quite good at answering questions but you need to know what questions to ask in the first place and you also need to know enough to know whether it's wrong, because nine times out of ten if you're lucky, or maybe seven times out of ten, it'll give you quite an interesting answer, what a person might say. But one time out of ten, or three times out of ten, it just spouts a load of rubbish. And the thing is, of course, it doesn't know, it doesn't even have a sense of when it's on insecure territory. Maybe that could be improved but in that sense it doesn't. It helps those who can who already have a sense of the field because you can look at an answer and say no that's just barking up the wrong tree. And I mentioned languages: I found AI to be a really good language teacher when I'm trying to learn some French or I'm trying to translate some French. And previously I could have just looked up in a dictionary what a word means but if I ask chatbot what it means it will often give me pages and pages of background information. Or if I notice that a word can mean more than one thing, I might say, “What are the various meanings of this word in French?” and it'll give me usually a pretty brilliant description. So, I found that it can be like an infinitely available quite good personal language coach. So, the same thing that could make me lazy when it comes to languages, if I'm invested in learning, can actually be precisely the opposite.

Laura Arends:

An interesting perspective on them! Personally, yes, I think it can make me lazy if I'm honest about this. But I have to be careful not to succumb to it too often. Actually I don't know, Professor Davison, whether you were prepared to take a question from the floor as it were now? We've got one that's come in related to the question about life after death (which is resurrection after death) which says:

Question 5 – *[00:21:09]*

How is the idea of the bot with the memories of loved ones any different from Hick’s Replica Theory?

Answer

Laura Arends:

Now it might just be worth just revising with us what Hick's Replica Theory is first.

Andrew Davison:

John Hick was a member of the United Reformed Church, I think. A British theologian, representative of you might say the liberal tradition of the later 20th century. And he was interested in, at this point, continuity between life and death, probably in his book called *Death and Eternal Life*. {Gesturing to bookshelf} It may well be even behind me! Anyway, so he was in that sense relatively untraditional. He didn't particularly want to use the language of the soul persisting after death and so he thought about the

Resurrection of the Dead as more like the cessation of a person and then God recreates that person as a perfect replica. And in this way you didn't need to think about the soul surviving between life and death. In what way is it different? Well, one thing I'd say, is probably that most Christian traditions haven't taught that. So, Hick was going out on a limb by saying that and it was an answer to a problem that a lot of traditions wouldn't have because they think that the soul does the work there. But if you did go down Hick's line, then I had it in the back of my mind that this was actually quite a nice parallel, so absolutely full marks for whichever student has asked this question. The difference is that God, who created the person in the first place, is truly recreating that person. Whereas this isn't that person, because at least at the moment and maybe never, the computer, the chatbot, doesn't have consciousness. So, Hick's view is God is recreating that person with its own 'I', its own person, and that's not what we're talking about here. And also, there's just the difference of capacity or just how good a job it is. I mean in Hick's view it is the perfect re-creation of exactly that person. This is just scraping what data we can from the internet throwing it on the top of a large language model and, I suspect, this is going to be pretty disappointing because for one thing there's not much capacity for anything new ever happening. It's not like, in Hick's view, the person exists again and can make genuinely new, free, interesting creative decisions and thoughts. Whereas the chatbot is a not very good, I suspect; just re-animation of some past ideas. So, I think there's a big difference there but it's a lovely parallel.

Laura Arends:

That was from Jennifer, so thank you Jennifer for that. It was great.

Question 6 – *[00:24:15]*

We've got another question here picking up on something you were talking about briefly in your lecture and you mentioned the link between maths and theology. I wonder if you could expand upon that. That's a question from Abbey Grange Academy.

Answer

Andrew Davison:

Excellent. Well, in the Middle Ages before you got on to anything that was like theology you had to do lots of other things first. You had to do things like rhetoric, for instance, and all sorts of things to do with argument and speaking, public speaking. But I think even before that the first four things that you looked at were basically maths four different ways. There's music in there, but even in those days all music was about the mathematics of ratios and proportions so like to say, in the Middle Ages, you had to do a lot more maths and science, such as they were in those days, than you do today. Where now you can sail through a degree in theology or English without having to do any science. So, far from these being the Dark Ages they were really interested in everything that could be known. I had Plato in mind when I made that comment because Plato has this idea in the book called *The Republic* of the stages in the ascent to God. And he begins, there are probably several different stages in between, but I'll summarise it in three. There's the knowledge of the world around you, which he's rather down on: he thinks this is just shadows

of the eternal truths. Then, where you're ultimately trying to get to is the knowledge of the things themselves, so not copies of goodness, but goodness itself. Not, for instance, copies of the idea of a horse, but the very existence of *horseness* itself. And he has maths as an intermediate level because you get at it through thinking about the world, but it's always a little bit free from the world. So, when you understand what a circle is, you might only understand it in the first instance because you've seen circles in the world. But once you've got the idea of a circle, that floats free from any particular physical thing, and you can get further and further into maths. It's all about the relationship of ideas to one another, and he thought of that as a middle way between the copies of things that are down here and the eternal ideas, things that are up there, because you're slightly lifted away from the particularities of things towards universals.

Now, Christianity is going to have some trouble with that, although we've learned a lot from Plato, because we generally think that matter, and the body, and particularity are a good thing. God knows about every sparrow that falls from heaven. Particularity, finitude, materiality, these are not problems; these are part of God's good Creation. So, we want to criticise Plato a little bit there. But there's still probably something to be said for the role that contemplation of mathematics might play in human life or human knowledge and that might include just wonder. So, Aristotle says at the beginning of the *Metaphysics*, 'Philosophy begins in wonder'. And I, certainly in trying to get my maths up to... I mean, I just got a lot of maths because I was a scientist, but the last couple of years, I've been really trying to get into certain bits of maths because of its use for understanding AI. And there have been moments that have been a little bit like kind of mystical rapture — which is what Plato is after — when I thought, "This is just so amazing". I mean, the way in which it all hangs together, or the way in which maths can unveil, can reveal, the sinews of the world, the structure of the world, the connections of the world... I have found that, I mean, this is really where this point comes from — I was watching a video on linear algebra, of all things, which doesn't seem like a very interesting topic, on a thing called the Singular Value Decomposition, and I really had a kind of Platonic, mystical moment because it was just so extraordinary, and beautiful, and well-ordered, and intricate.

So, I think there's something to be said even just for wonder and appreciating these... I guess it does get you a little bit towards eternity. I mean, there are various different theories of the philosophy of mathematics, but you are contemplating things that are somewhat timelessly true, and that might be a stage on the way to contemplation. Although, of course, Christianity will always want to say the truest revelation of God was a Palestinian Jewish carpenter who trod around the hills and talked to people. So, I don't want to elevate the abstract too much, but it probably has a place.

Laura Arends:

That's fascinating, learning about that sort of link between maths and theology. We've got another question from the floor here. Actually, this is from Beth, and this relates to the second question when we were talking about computers being less sympathetic, perhaps, than a human being. It says:

Question 7 – *beginning at [00:29:18]*

If law is part of our moral imperative, why can't a data system be taught a moral response that would be less subjective?

Answer

Andrew Davison:

I think we may well see AI used a fair bit in law, and there is a way in which you could say that there's an objectivity to it, which is quite useful. A parallel example might be: say you had a very busy hospital and you got to the stage where you couldn't treat everybody immediately — it's called *triage*, where you work out priorities — and you might say that a computer is to be preferred there because it can be incredibly well trained, and you take sort of personal, subjective elements out of it. The truth is probably going to be a bit less rosy. For one thing, the way that AI has got really, really good is through systems that are so unbelievably complicated that we can't understand how they work. So, there's a whole area of AI research where people are trying to make *accountable* AI or *transparent* AI, so it doesn't just come up with reasons, but you can see what the reasons are — how it got there. That's still relatively tricky and in its early stages. At the moment you would have a decision, and there would be no kind of trail of how the computer got there, and that could be problematic. These computer systems are kind of unbelievably brilliant... until they're not. They can make a thousand great choices and then just a complete mistake. And maybe that's not so much of a problem if it's just — I don't know — a rich person's money; you know, maybe that's not quite as important as a person's life. But you've got to have incredibly high standards for how often it's going to make a mistake when you've got things like life and death, or justice, and who gets imprisoned.

Then there are also questions about bias. An AI is only as good as the data that you train it on, and there are plenty of examples of how bias has entered into AI systems. So, most of the people that were doing the work were men, and it seems like there may be... that women come off less well from decisions that AI makes. People are working really hard to try and put that right. And maybe it's not so much the programs that are at fault here — it's the training data. So, like facial recognition or voice recognition: most of the examples that have been used historically have been of white people, and the systems that have been produced as a consequence have inevitably have been better at recognising white faces and voices than other people. And it's just there's a bias there. So even if it's not intended, it's very, very difficult to get rid of. So, we talk about objectivity, but if there's bias in the data that we're using to train, then it's just going to be perpetuated by the system. So, my take-home message is: I'm not against AI, and I think we should think of all sorts of creative ways of using it, but we also don't need to be starry-eyed about it.

Laura Arends:

Thank you. Thank you. I hope that's answered your question, Beth — that's brilliant. We've got another question here from Jake this time:

Question 8 – *beginning at [00:32:55]*

Is it realistic to consider that AI could be all-knowing in the future, and to what extent may this affect human belief in the omnipotence of God?

Answer

Andrew Davison:

Well, I think these are great questions because they help us to think about the terms that we're using. So here, what's 'omnipotence'? Or maybe we should also be talking about omniscience here: to be all-knowing. And divine omniscience has not been thought — I think in the best writing anyway — to be simply a knowledge of all things, like an infinite sum of facts. God's knowledge, God's intellect, overlaps completely with his goodness, mercy, justice. We have to prize them apart because that's the way our minds think, but divine all-knowingness would be the same as all-powerfulness, which would be the same as being merciful. I mean, so that's one thing. But also, God's knowledge of things has been understood as like a maker's knowledge; not that God has to look at the world and learn from it, but that God's knowledge, rather, is productive. It's God's knowledge and thoughts that produce the thing in the first place. So, there are just all sorts of ways that we might say actually, God's knowledge doesn't just look like an infinite sum of facts.

Then, with omnipotence, obviously, knowledge is power. When you know things, you can do things with it. But I'd still think that was part of the human domain. AI is going to make us more powerful, but theology has a very long history of grappling with how human power and actions fit under God and we've generally not been considered to be doing a great job of that. The real challenge is to do what you do with kindness, and love, and mercy. And I'm not sure that if we're just more powerful, in some ways, it just makes things worse. So just the fact that we're more powerful doesn't really seem to change the parameters so much. It just asks us questions about what divine omnipotence means. God has been said to create all things out of nothing. It doesn't matter how big a computer is, even the size of the universe, it wouldn't be able to create all things out of nothing and hold all things in being. And I'm not sure it will necessarily have personhood or kindness. So, I think theology has usually said that we hardly know what we mean when we say anything of God, but that's because of God being so perfect. The words — human words — point in the right direction, but they don't capture what God is. And generally speaking, just putting 'infinite' in front of something doesn't really get us that far. It's like God's qualities are qualitatively different, partly because of this overlapping — the goodness is the mercy, the mercy is the wisdom, and so on. So, I think historical theology would not think that more of anything creaturely gets us very much nearer to being like God, especially if it's just a matter of quantity.

Laura Arends:

Thank you, thank you. We've got another question now, from the Abbey Grange Academy:

Question 9 – *beginning at [00:36:52]*

Do you believe that AI was God's plan for us to gain greater knowledge of his existence?

Answer

Andrew Davison:

Well, I'm almost inclined to say, "Let's wait and see," because if it turns out to be a complete horror show, then that might be good reason for not thinking that it was God's great plan. I think that Christianity, and most religions, would say it is God's purpose or intention, in general, for us to know things: to know the truth, and to know God. So, in that sense, I'm almost inclined to say yes to any question like that, because anything that leads us to better understanding, and especially of God, has got to be a good thing. But I think I need to add that knowledge only gets you so far. Certainly my Christian tradition, and I think other traditions as well, would say it's good to know about God, but it's even better to know God, or to trust God, or to love God. So, it's much more like: I know that my mother exists, but that captures very little about my relationship with her. Much more is that I have a sense of who she is, and that I love her, and that I trust her. And that would be pretty central. I mean, Luther was very big on this: just facts about God don't really get you very far at all. So, for that reason, existence — I'd rather go from existence and add: “But who?” and “What?” and “How?” and *the character* of God. And I'd rather go from knowledge to say: That's fine, that's part of it, but that the Christian tradition, and I'm sure other traditions as well, say what really matters is trust, and love, and a relationship. And I'm really not sure that AI is going to get us very much further on that front, compared to the things we've already got.

Laura Arends:

Brilliant, thank you, thank you. We've got a slightly more unusual question here, but I think it might give you a springboard to talk about quite a range of different things. Again, from Abbey Grange Academy:

Question 10 – *beginning at [00:39:05]*

Do you think that AI will be able to explain historical mysteries, such as how the pyramids were built?

Answer

Andrew Davison:

Well, as long as you've got data, and it can process the data, you might get somewhere. I'm not sure that it's going to be particularly good at one-off examples, because generally you need to train the AI on lots of things you already know. Like 100,000 pictures of cats labelled as ‘cats’, and 100,000 pictures of dogs labelled as ‘dogs’, that's called supervised learning, and that's the gold standard. So, if you had lots of data where we knew already how different things were built, you could train an AI system. Then you could give it an example that you didn't know the answer for yet, and it would be able to follow the patterns and give you a suggestion. This might be a case where we just don't have enough data where we already know the answer, to be able to train it so that it can give the answer when we don't know what the answer is. But in general, where we do have plenty of cases where we know the answer, you could train something to solve or think of examples. Potentially, I reckon it’d be quite good with paintings, for instance. If you trained it on every known example of lots of different painters, and we had very secure attributions, and you gave it a painting that it hadn't seen before, it’d probably be pretty good at working out who the artist

was. And that could turn up in the study of religion in something like — presumably, you could do it with texts, and it would be able to guess who the author was. Or you could give it lots of different scribes' handwriting, and then you give it something where we don't know which scribe it is, and it's able to work it out, like a well-trained human being could. I think that's all I'm going to say about that, unless you have any other angles on that you want me to discuss?

Laura Arends:

No, I think that covers quite a lot of different elements. Thank you. Now we've got a couple of questions coming on from the lecture, really. Thinking about towards the end of the lecture, you started talking about what a wonderful discipline theology at university is. I'm biased here; I read theology at Cambridge, so I definitely agree! But we're wondering, really:

Question 11 – *beginning at [00:41:48]*

Do you think that the study of AI will become an intrinsic element of theology — possibly even joining biblical studies, ethics, and philosophy as one of the foundations of study in the subject?

Answer

Andrew Davison:

Well, one of the things I love about theology is that it's about everything. My theological hero, Thomas Aquinas, says the subject matter is God and everything as it relates to God, which is... everything. So, I love that *omnivorousness*, and it served me well. It's delighted me, especially when I had the Science and Theology position in Cambridge. All of science was fair game for me to think about theologically, and that was marvellous. So, in that sense, if it's part of the world, theology should be thinking about it. And especially if it becomes an important part of the world, in terms of ethics and human thriving and well-being and jobs and money and all those sorts of things, like wages, then it would absolutely have to be part of the curriculum. I also think it's already got to the stage where it's going to be such an important tool that we ought to do that consciously. We should probably have 'how to use AI well and responsibly'. Just like we would also have classes on how to quote things properly and not fall into accidental plagiarism, or how to use a library catalogue. We don't tend to concentrate too much on skills, because what we're all passionate about is the content. But there always will be some training, even if it's just on the side, as you go along, in skills, and methods, and resources, and tools. And I think AI comes under that category. In certain realms it's going to be important to try to understand the thing itself, so to learn some maths. So, I think if you have a Science and Religion course, I always think you need to think about the nature of the thing you're studying. You need to put the hard work into understanding what this bit of biology is; you don't want to talk about it at arm's length or in general terms. So, if I had a Theology and AI class, then I would be wanting to encourage the students to really dig down into how it works — that kind of thing. If you're just teaching it generally as a tool, then maybe it's not quite so important to be interrogating what's going on under the bonnet.

I would say also — and this is a really good example for me — last week there was a terrific lecture on the Dead Sea Scrolls. Some people from the Netherlands were using AI to date the Dead Sea Scrolls and it was taking all the data that they had from carbon-14 dating and all the information that the computer could discern in the changing of handwriting styles. It meant that every bit of data that you had on every scroll reinforced itself, and you actually ended up with tighter, more defined predictions about how old any of the scrolls were than you would've got just looking at any of those bits of data by themselves. But what it threw up at the end was something that actually needs a fair bit of skill in statistics to understand. So, previously, you'd have just said, "We think this scroll is X hundred years old," and the most extra data you might have got would be something like an indication of how certain or uncertain we were about that. But that probably wouldn't even have been given in the books. What this data was throwing out was something called a probability density function, which is: across a range of dates, what's the likelihood that it's that date? When you sum it up, it comes to one. If you look at any particular region, it gives you the probability that it's in that region. And there will be examples where a scroll would have two humps, and you'd say, "Okay, so it's pretty likely to be either this age — approximately, give or take a bit — or that — give or take a bit. But it's unlikely to be the date in between". So the old average would be wrong, because the average, just one figure, would put it in the middle. Whereas now, we're saying it's probably around here or there. And there were even error bars on this probability density distribution. So not only what we think the distribution looks like over age, but even how certain we are about the estimation at that particular age. Now, I'm a scientist. I can take that in my stride. But nowadays, if you're going to be working on the Dead Sea Scrolls, or anything else that's been treated with AI in this kind of way, you just have to have quite a bit more statistics in order to be able to make full use of that data.

So, it's not even just the using and I'm not suggesting that anybody in the classroom would ever use AI to do what those amazing scientists had done, but in order to be able to even understand what the findings were, you might need new maths, or more maths than you had before. So, I think it is going to be an important part of training, and education, and research going forward. But it won't take away the hard and wonderful work of just sitting in libraries and reading books. I don't think.

Laura Arends:

Thank you, thank you. A lot more maths than my theology degree involved, certainly! One final question for you, Professor Davison. It's been absolutely fantastic talking to you this afternoon, and I hope the people viewing the Q&A have enjoyed hearing what you've got to say and answering the questions.

Question 12 – *beginning at [00:47:31]*

Final question then: What's next for AI and theology?

Answer

Andrew Davison:

Well, I think I began to give an answer to that with the previous question which is really taking it seriously: what it can do, whilst also not being led astray by unrealistic expectations. A lot of that is going to come

down to almost professional guidelines, really, or a sense as a student, as a researcher, or whatever, “How do I use it well? How do I know what its limitations are? How do I have to be transparent about the way in which it might have played a part?” So, I think not just *using* AI, but also taking a step back and thinking together about how we use it, and doing that work is important. It's obviously going to play more and more of a role in theological ethics as a topic. And like I said in the lecture, there's all sorts of stuff lurking behind the scenes here. For example: Who's being paid to label those pictures of cats and dogs? How much have they been paid? They're almost certainly in the developing world, and they're probably not paid very much. So, there are justice questions. Then, questions about copyright — that’s big in the UK at the moment. It's in the newspapers and perhaps coming before the government and Parliament to have laws about how or whether things are protected anymore. So, all sorts of ethical questions.

I’m interested, from my point of view, in the more philosophical angles. So, for my part, what I encourage people to think about is: What kind of thing is this? How does it work? How might theology's philosophical resources help? I plough a rather niche furrow here, which is: what does the great tradition of scholastic theology — often gathered around Thomas Aquinas, but not entirely — as a certain kind of retooled Aristotelian philosophy... what does it have to say about AI? I think there are lots of wonderful concepts within that tradition that can help us think about what AI is and how it works. That’s a niche interest, but I think there's more work to be done there.

There’s also lots of room for imaginative use of AI, where there's lots of data to be processed. Even that wonderful talk on the Dead Sea Scrolls — for almost every result they presented as better understood, it raised additional questions. There had been a thought that there were two different scribal handwriting styles in the scrolls — Herodian and Hasmonean, I think — and one was older than the other and they had used this amazing way of relating all the data to get perhaps better dates, and it seemed like the order and the idea that one succeeded the other wasn't quite as fixed as previously thought. But I asked in the questions: “How about there are various methods you can use where you just discount the old idea that there were two handwriting styles?” You could just present all the data and say: “Does it cluster into one, two, three, four, five?” You can cluster indefinitely, until you have as many styles as there are manuscripts, but you could also force it. You could say: “What number of handwriting styles captures 90% of the variation? 95% of the variation?” And it might be that three does that. Then four looks like it's just not a particularly useful subdivision. That's cluster analysis. So there immediately, out of that lecture, even after all this amazing work, further avenues were opening up.

The danger is that we might start to think the kind of questions AI is good at answering are the only questions and that everything should be treated that way. Or worse: *that if it’s not that kind of question, it’s not a valid question anymore*. And I'd be worried about that. I love to see the fact that there's money coming into thinking about my kind of questions from AI, but I'd be sorry if that then suggested that the only questions worth asking are the ones where AI can be applied. Because I think that’s just not going to be the case.

Laura Arends:

Brilliant. Well, thank you so much for giving your time to us this afternoon. I've learned loads and I'm sure our students who viewed this have learned a lot as well. So, thank you very, very much, Professor Davison.

Andrew Davison:

It's been a wonderful opportunity to be with you, and it reminds me of just how much work Westminster Abbey does in all sorts of ways. I know that you're great supporters, for instance, of homelessness work in London. There's a lot that you can see of Westminster Abbey and it's amazing and I recommend that you go and have a look at it, but it also does a lot of work that isn't immediately obvious, and it's very impressive.

Laura Arends:

Thank you very much, and a good afternoon to all of you.

Biography of the speaker

Andrew Davison is the Regius Professor of Divinity at the University of Oxford, where he is also a canon of the cathedral at Christ Church. Before that, he was Starbridge Professor of Theology and Natural Sciences at the University of Cambridge. He went to Oxford as undergraduate in 1992 to read Chemistry, staying on for a DPhil in Biochemistry. In preparation for ordination, he moved to Cambridge for an undergraduate degree in theology in 2000 and later worked in a parish in Southeast London. After that, he taught Christian theology first in Oxford, then in Cambridge, before the Starbridge appointment. He works on theology, philosophy, and science, including a recent book on life beyond Earth *Astrobiology and Christian Doctrine* and some papers on AI.

Context of the event

Westminster Abbey's Learning Department hosted the online event 'The Theology of Artificial Intelligence' on 5th March 2025. In this pre-recorded lecture, Andrew Davison set out some of the main achievements of AI and explored areas where Christian theology has a response: ethics and politics, philosophy and Christian belief, how we talk, and the place that AI ought to have in our lives. Students were able to watch and submit questions before the live event. Laura Arends, Head of Learning at Westminster Abbey, chaired a live question and answer session with Andrew, during which students' questions were answered.

To hear more about attending events like this with your students, please sign up to our schools mailing list.

You can also find related resources on our teaching resources page.