Westminster Abbey
Learning

The Theology of Artificial Intelligence lecture – 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 lecture here: https://www.youtube.com/watch?v=PfOrEVrZ_eM

Please note that all opinions are of the individual speaking and do not represent the views of Westminster Abbey.

Follow up questions for discussion

After watching this lecture, your students could debate one of these questions verbally or provide a written response for homework.

- Professor Davison emphasises Aquinas' view that tools are for the sake of the user. Should AI be seen as a tool, a partner, or something more?

- Professor Davison links human creativity in building AI to the theological doctrine of I*mago Dei* (being made in God's image). Is developing AI a reflection of divine creativity or a dangerous over-reach?
  Where do we draw the line between imitation of God and 'playing God'?

- The lecture explores how AI impacts work, inequality, and global justice. How should Christian ethics respond to the effects of AI?

- Professor Davison warns against over-focusing on speculative futures. Is it a distraction to focus on the idea of 'conscious' or 'sentient' AI or 'general' AI? Do you agree with his view that theology should focus more on how AI is used today rather than on distant possibilities**?**

Transcript

Welcome to this lecture on Theology and Artificial Intelligence which is part of the Westminster Abbey A-level Masterclass for 2025. I'm delighted to have been invited to speak on this important topic, not least because it shows something that's particularly important to me: that theologians can and should be interested in everything and that our religious traditions are a deep and constant source of ideas and insights as we think about life today. I'm Andrew Davison the Regius Professor of Divinity at Oxford.

Think how much work and the economy and human life changed during what we now call the Industrial Revolution of the18th and 19th centuries, driven by technological advances. Well, many people think that we are in the middle of another revolution like that, comparable to it, this time driven by progress in artificial intelligence and that topic will be the subject of this lecture.

I thought it would be good to begin with an example of artificial intelligence in action. This is Chat GPT answering some questions about Westminster Abbey.

….

I think that's pretty impressive. It's coming up with a discussion of the role of technology in the building of Westminster Abbey pretty much as fast as we can read. And I chose that question to make a point: that there's always technological innovation. What looks like a radical innovation in technology one day easily becomes just the common or garden background to life in a later age. And that means that much that strikes us as routine or just how things are done, like using pointed arches to build Cathedrals, was once cutting-edge technology.

As we get started thinking about this topic together it's worth asking a few questions about what we mean by some of the terms that are in play here. So, first of all, what do I mean by ‘theology’, which is my particular academic field? I'm going to be talking about God and religion and religion's great themes from the perspective of Christian theology which is my own tradition and of course that's also the tradition of Westminster Abbey which is a Christian place of worship. Other religious traditions, other religious intellectual traditions, will also have important things to say about AI. No doubt on some questions, perhaps on many questions, representatives of different religious traditions will have something quite similar to say, although I'm sure on other points they will differ. And I found that this sort of question, often quite a new and quite a practical one, can be a good topic for interreligious conversations, since it throws up new angles to be asked about old themes and topics.

If you study Christian theology and religion, what sorts of ways can you go about doing that? What sorts of skills can you draw upon? What sorts of academic disciplines find themselves in theological research, or for that matter, within a theology undergraduate degree?

Well, we can use the word theology to mean quite a wide range of things and you see some of them on the slide here. In my faculty, even from just the Christian perspective we will be thinking about the Bible and philosophy, church history, ethics, biblical languages, doctrine, and then things like theology and literature and theology and science. Sometimes, when we talk about theology, we mean one part of that when it refers to the study of Christian beliefs or teachings, which is what I've called ‘doctrine’ here. That just means ‘that which is taught’. So Christian theology understood as doctrine is talking about, thinking about, core Christian beliefs like God, creation, sin and salvation, Jesus, the church and the life of the world to come and so on. I'm going to be talking a bit about that but I'm also going to be bringing in discussion of philosophy and ethics and a bit of study of the Bible. Others of my colleagues in Oxford could talk about various other parts of theology. I suppose that this is an exercise I should say as well in theology and science, which is what I taught in Cambridge before my recent move to Oxford.

What do we mean by artificial intelligence? Well, it looks like we mean ‘an artificial or made thing that has intelligence’. Those artificial things would be computers or programs or apps. But I don't think that's quite right. I'd rather say that we're talking about using computers to do things that we manage because we are intelligent and which computers manage even though they are not. So it's a really important question and gets to the heart of much that we're talking about. Are these computers really intelligent or are they just good at doing in their own way what we do because we are intelligent? And I would endorse the second perspective there.

I've been working with some Swiss theologians, philosophers, and computer scientists recently and they pointed out that other names were suggested for this field early on. The name “Artificial Intelligence” stuck but there were other suggestions and some that may indeed have ended up being better. So, for instance, what if we didn't talk about *artificial* intelligence but instead we talked about ‘*extended* intelligence’. I like that name because it points out that the intelligent thing in the room is still the human being, not the computer. But we are now *extending* our intelligence and what we can do with our intelligence, by using computers as a sort of tool or instrument, (and I'll come back to the idea of AI as tool or instrument later on). So we might have preferred that what we're talking about would be understood as *extended* intelligence rather than artificial intelligence, but artificial intelligence won out and that's what we talk about today.

Let's think just a little bit more about this question of what artificial intelligence is, what it involves. I'm going to say one more thing about it which I think gets to the heart of what we're talking about and that is to say that artificial intelligence is an exercise in finding patterns. Let me say one more thing about artificial intelligence: what it is we're talking about. I think it's helpful; it brings it down to earth and perhaps punctures the bubble of the hype around it a little bit. And the point I want to make is, that artificial intelligence is really an enormous exercise in finding patterns in data.

You're probably familiar with the idea of curve fitting from maths. It is about getting some data and working out what's going on in that data by drawing a line through it. Sometimes we draw a curved line but the simplest example is drawing a straight line. And here we've got some data: it plots people's weight against their height and as you can see there's definitely a pattern in that data and we could probably approximate most of the information there with a single straight line. We've done that here and once you've got a line drawn like that, or it could be a curve, but we can do it here with a line, then you can set the data aside and use the pattern you found to answer some questions. So, this is data about height and weight of children in a school class. Once we've drawn the line we can say, ‘What's the best estimate of how much someone would weigh who's 148 cm in height?’ And the answer, as you see, is about 52 kg. Well, AI is basically doing that. We feed in enormous amounts of data and we use all sorts of clever computer techniques, really clever mathematical techniques, to find the patterns in the data. Here, with this example, we're capturing what we want from the data in two numbers or parameters, so we could change the angle that the line makes rotating it clockwise or anticlockwise and we could move the line up or down. And that uses up all the ways that a line could be different so we're approximating the pattern in the data in two numbers.

What makes contemporary artificial intelligence so clever, frankly so amazing, is that it can extract many, many, many aspects of the pattern or the meaning from the data that it's given to work on. So, the version of chat GPT you saw a few minutes ago seems to have extracted the pattern or meaning of

what it's been shown from lots and lots of examples of human writing. It's sort of extracted the pattern or the meaning from all of that and encoded it in something like 200 billion numbers or parameters. So, I gave you some data, we had a line we're able to work out what's going on and summarize it with two numbers of parameters. Chat GPT is probably working with something like 200 billion numbers. And that's a good way to think about what's going on here. It's like a computer with 200 billion dials and they've each been twisted to one value or another, amazingly, in a way that can extract or represent a great deal of the meaning in all the English texts that it's ever encountered. We have data on the screen plotted in two dimensions on a flat page. The best contemporary AI is working because it's done something like drawing a line through data in 200 billion dimensions, or maybe something like a line with 200 ways of wiggling through data in one billion dimensions.

Anyway, it's an absolutely extraordinary achievement and I find that way of thinking helpful. AI is a glorified, and glorious, exercise in finding patterns in data. Typically, you first train it and then you use it. So, perhaps you show it pictures of fruit; you ask it to guess what the fruit is. It's probably hopeless at it. You tweak some of those dials and have another go. Maybe it's a tiny bit better or a tiny bit worse and you use that information to nudge the dials bit by bit, one go after another. And if you do that billions of times you'll have your dials set pretty well for being able to give the system a picture and it being able to label it correctly as one kind of fruit or another.

I mentioned church history as one of the subdisciplines within theology, part of what people in my faculty are experts in and part of what you can study if you come to do a degree in Theology and Religion. Well, you might think there's not going to be very much artificial intelligence in the stories of church history and by and large you would be right, but I can think of at least one story and I thought I would pass that on. So, there's a legend about St Albert the Great, the most impressive of medieval scientists, about him making an intelligent robot in the middle of the13th century. And the story goes on, that one of his most brilliant students, in fact his most brilliant student, St. Thomas Aquinas, took offence at this thinking robot and destroyed it. Now, I'm not saying that happened. I don't believe there was artificial intelligence and thinking robots in the middle of the 13th century but that doesn't make the story any less interesting. We can ask what was it about Albert and his work, what was it about thinking machines and the nature of intelligence, that captured the imagination of the people who told the story and were interested enough in it to pass it on and write it down. And which side are we supposed to be on in this exchange? Are we supposed to with Albert and be sorry that the machine was destroyed, that he'd so ingeniously made? Or are we supposed to be on the side of Thomas and think that this was a monstrosity that should never have been put together in the first place? Well, there's one little bit of church history or a church legend that got passed on which I think is of interest and relevant to our topic.

More broadly thinking about history, I want to say that it's always useful to have a historical view on things. And if you're coming to study theology and religion at university, if you're thinking about that, if you're doing really anything at university at all, or life in general, it's always good to have a sense of the history of things and with AI that is certainly true, not least because it really helps us to understand where we are now; to have a sense of the history of how we got here. We'll see that the story of artificial intelligence in the past few decades advances, both in maths and software on the one hand, but also in advances of computing power or hardware on the other. Progress has come in fits and starts. There has been enthusiasm and then disappointment; enthusiasm again. So, there was a great deal of

excitement in the 1960s and into the 70s which failed to produce what was promised and there was sort of deflation there. Then in the 80s the idea of ‘expert systems’ took off when people tried to isolate human expertise and code it into computers. That didn't prove anything like as powerful or useful as people hoped, so funding dropped and many seemingly promising companies went out of business. These periods of disillusionment, of lack of investment, are often called ‘AI Winters’.

Well, the whole idea of artificial intelligence revived in the 1990s and although there have been some important insights and advances in the maths and the computer coding, a good part of the sudden and ever-growing success of AI has come from the availability of more computing power.

You might have heard of Moore's law: the idea that the number of transistors on a chip will double approximately every two years, which is a good indication of the power of this computing unit. Well, it turns out to have been a really good approximation or prediction. If you have a look down there at the graph on the slide, you'll see that this prediction of doubling every two years, which is the line, is followed pretty closely by what's actually happened. Notice that this isn't what we call a linear scale on the y-axis, the one that runs up and down, but a logarithmic one, which means that every time we go up by certain amount it's going to double, in fact. Here you see that the labelling on the side sees it increased by 10 at each of those units which is to say that the power grows exponentially; it takes off really, really steeply.

So, I'll give you some examples of just how extraordinarily the computing power that we have available has grown. It's not easy to tell exactly but it's said that ChatGPT4, which is the chatbot that I used right at the beginning, cost about $70 million to train. And 10 years ago, that amount of computing power would have cost a great deal, more many hundreds of millions of dollars and even as recently as 2010, so only 15 years ago the same amount of computing power which cost, we think, about $70 million would have cost something like $125 billion, if even that much computing power was available in the whole world. Another example: what would have cost about a million pounds in computing power in 2010, so only 15 years ago, costs only a few hundred today. So, over the course of 15 years, what would have cost a million pounds is down to just a few hundred. And a lovely part of this is the story of what's been played in this by the PlayStation, especially the Sony Playstation 3 of 2006. It turns out that, even though scientists and lots of other people wanted to use artificial intelligence, that wasn't a big enough constituency, it wasn't a big enough purchasing block, to really drive the computer technology and make it worthwhile for these chip manufacturers to make big innovations. But that's not the same, that's not true, of the enormous number of people who want to play computer games. And so, it turned out that wanting to do that better led to some really significant advances in chip making and how we build computers. Especially the graphics processing units where you work at what's been displayed on the screen. And in particular, there was this really important advance around the time of the Sony PlayStation 3. So, all the amazing artificial intelligence advances of our own day, which include some really spectacular scientific ones like the ability to predict the structure of proteins, which I'm sure will revolutionize medicine, all of that really stands on the shoulders of lots of gamers wanting to have better and better graphics. And that was a large enough group of people who were paying real money to drive the development of these more and more advanced chips.

I want to encourage you to think, then, about history and the story of artificial intelligence over the past decade or two. I also think of course it's worth thinking about artificial intelligence now and how it might

be in the near future and, in fact, I think that's probably where our focus should be. But I also want to point out that a lot of thinking about AI, not least a lot of theological thinking about AI, is about the speculative far future. People might not think it's the far future but they're talking about things that we have no real reason to know whether they're even possible, and certainly the sense that they just lie around the corner may be really quite wishful thinking.

Thinking in particular of artificial *general* intelligence; so, the creation of a system as intelligent in the same sort of way that we are, in as much as it can do many different things in particular, this idea that when we get there these entities, these computer systems, may be conscious, free, moral or immoral and I just don't know that we're ever going to get there. Possibly we will, but it's quite an extrapolation to say that because we've got computers that are good at recognizing cats or fruit or whatever, or even that we've got these computers that can extract the sort of sense of patterns and meaning in human texts and come up with new human sentences or new English sentences or whatever language it is in response to prompts, it really is just a complete jump from anything like that to say that we're going to have computers that are, for instance, conscious and I think it's a bit of a red herring. By all means let some people think about those sorts of questions, 'Can computers be conscious and so on', but I think there will be a real problem if we put so much emphasis on that, that we weren't thinking about all the ways in which computers are busy transforming the world here and now and that will be particularly important when we come to think in a moment about ethics.

So, in what remains of this talk I'm going to focus in a little bit more on some theological detail and I've got some examples of the sorts of things that we could think about.

So, there's the whole area of ethics, which is an important part of Christian life and Christian thought and the sort of thing that goes on in theology faculties at universities. And ethics is going to throw questions like, 'Is AI destructive or can it be constructive?', 'What would it mean for us to use it well?'. If we're thinking about using AI then, maybe something that doesn't immediately come to mind, but if we're thinking about AI in theology, one of the questions we can ask is, 'How can AI itself be helpful in theological work and the study of beliefs about God, about religion and so on?'. Then we've got what I call Doctrine, the core teachings of the Christian faith and we can ask, 'How does AI relate to aspects of Christian belief?' And then I also want to point out that Christianity, like many other religions, has been a custodian of wonderful philosophical ideas. It's gleaned them from many different sources. In the case of Christianity, both Aristotle and Plato have been particularly important and then it's developed those ideas in its own ways, in distinctive ways. It's developed its philosophy in conversation with other religions, especially Christianity and Judaism and Islam have had a very rich and important, historically important, conversation about philosophical matters. So, there are some of the things that we can talk about in the rest of this talk.

I suspect that it seems obvious that if we're going to talk about AI and theology or religion then we're going to get to ethics and you'll be right. It's a very important area and there's work going on in my own faculty among the ethicists who think about AI questions. All I want to say at the beginning of thinking about AI and ethics is that ethics is *important* for religion but it's not *all* of religion, and I think sometimes especially if people are looking at it a little bit from the outside, they think that religion is mainly about following certain rules, right and wrong. Of course, that's important but there are all sorts of other aspects as well: there's worship and liturgy; there's community. There are all sorts of ways in

which religion will want to shape people and move them to live good lives, but it won't be just about lists of rights and wrongs and laws. It's also going to be about forming us, shaping us into be into being certain sorts of people out of which certain sorts of actions will follow quite naturally. So, especially if we're mainly thinking here about rules that wouldn't be all of religion but it's certainly an important part.

One important set of ethical questions around AI are its impact on the environment. People talk about extraction and especially the idea of extracting energy from the earth and all the oil and gas and coal and so on that we might dig up to power these quite energy-hungry computer systems. I'm not sure that it's, I mean it's important, but I'm not sure it's quite as important as people necessarily say it is. It might come to one two per cent of human energy use. There are all sorts of other things that we do as well and I think that if, for instance, we ever flew we would be much better to think about that in terms of our impact on the planet than we would thinking about AI. It could be that getting too het up about the environmental implications of AI could be a bit of a fig leaf if there are all sorts of other things that are maybe more impactful that we could change and making airplane flights would be pretty much at the top of that.

There are other extraction concerns, though, when it comes to computing. You need all sorts of rare earth metals, some are rare elements, that are often found only in a few parts of the world and some of those can be quite war-torn or unstable parts of the world and the big inflows of money and all sorts of conflicts about who has influence in those parts of the world; the fact that the money doesn't always go to the people who live there or the people who are poorest, I think they are, you know, equally important extraction questions about what goes into the making of the computers that we're using.

Often with ethics the danger is that we're going to think about the sort of flashy and new and exciting questions, of apocalyptic questions in this case. I'd want to say that Christianity is invested in all the old-fashioned questions, as well about ‘Who has a job? and ‘Who doesn't have a job?’ and ‘What kind of job it is?’ and ‘Is it worthwhile?’, ‘Is it unbelievably tedious?’, ‘Is it properly paid?’, ‘Where do the profits go?’, all those sorts of things and there are all sorts of questions like that around AI. So, quite a lot of the work to make the data usable, because these computer systems have to be trained on enormous amounts of data, gets done by people in the developing world not paid very much. That's something for us to keep an eye on, I think, to be concerned about. I think these questions of work are absolutely central so if it's true that what we're going through, what we may be beginning to go through, is something like a new Industrial Revolution then all sorts of people are going to be put out of work as happened back then and we need to think together as a whole world and as countries about all sorts of questions of taxation and law and training and all sorts of questions about work and about money and the accumulation of money, so which jobs will be around in the future. These are pretty important questions, ‘How does it touch upon my own life?’ I'm getting emails from a publisher at the moment saying you need to sign a form which will give us the rights to use the books that you've written to train these AI large language models and interestingly, there's no I do not want to consent. But it's just an email that comes about every other day for each of the books that I've written for them saying, ‘Oh just sign here and you'll get a pretty small proportion of the of the proceeds,’ and I think there are therefore really significant questions about intellectual property.

It's good for us to think about how AI shapes human interaction and attention. That's one of the most pressing of all ethical questions, I think. There's a lovely picture here by Eric Pickersgill, an American photographer, which I'm using here with permission and he did a series of photographs, 'The Removed' series I think they're called, where he showed people in all sorts of American settings as if it was the present day with them on their phones, except he removed their phones from their hands and I think it's a really powerful and pretty scary invocation of how odd it is that we would be, for instance, together like this. It's on the Blue Ridge Parkway which is one of the most beautiful sets of views you can possibly have in the United States, of amazing forests And what are they doing? They're not looking at one another; they're not talking to one another; they're staring at their phones and that's made all the more obvious because the phones have been taken away.

We don't necessarily think of social media as being about AI but it absolutely is. AI is used to work out what posts to show you, to put together your reels or your series of posts that you scroll through and it's a really very important clever and economically significant use of AI. People talk about an attention economy that what people are bidding for, what they're making money out of, is your attention and it is worth thinking about. And it's a good religious question, attention, that these enormously powerful computer systems are being used to grab and monetize our attention.

I mentioned apocalyptic outcomes. There are certainly people who are thinking about this and I'm sure it's right for some people to be thinking about it. For instance, about whether AI could become sentient, could become conscious, then could far out strip us and improve itself and that would take off and take over the world. It's the subject matter of many a science fiction film. I think that those sorts of future questions are worth addressing but, like I've said, I think they can take our eye off the ball and it would be possible that we'd be so fixated by those sorts of apocalyptic questions that we didn't recognize what was happening all around us. And there is plenty of potential for far less spectacular catastrophes than the idea of the production of some sort of self-conscious AI system. All sorts of catastrophes. For instance, stock market trading is increasingly being done by AI models and you need all sorts of clever breaks and controls in there because AI tends to be really clever until it's not and of course it doesn't know whether it's being clever or not, so it'll often do something really spectacular many, many times and then suddenly it turns out there's a gap in what it can do and it gets it spectacularly wrong! And that could easily cause a catastrophe, for instance of a stock market crash. Or there are all sorts of other questions like that.

There are also humble questions that it's easy to just pass over but which are actually increasingly and incredibly important. Questions around bias for instance; so not necessarily any one huge bias but if the whole system is biased all the way through then whole groups of people are going to be disadvantaged and there's all sorts of interesting work on this: that the people who work on these AI systems tend to come from certain demographics. Other demographics are less well represented and so when it came to things like the production of faces if you're making art and you're doing that of a typed in prompt or the recognition of faces or the recognition of voices it turned out that certain racial groups were treated much more sympathetically or accurately than others. Similarly, certain sorts of voices, perhaps men more than women, for instance, the systems were good with one sort of user, less good with another. So, there are all sorts of questions about bias and that is becoming an important topic and it turns out to be more difficult to deal with than perhaps people thought at first.

In all of this it seems to me it's important for us to think together about what kind of values we want to hold and what sorts of goals and ends we want to pursue and I'm pleased that in Oxford there's a new initiative, the “Oxford Collaboration on Theology and Artificial Intelligence” which I'm involved with and one of the things that the ethicists in that group are really driving is to be talking with people who make and use AI to come up with professional standards, discussion of professional values. The sort of commitments that people might make just like doctors, for instance, to make professional commitments about doing no harm and that sort of thing.

If we move on from ethics to think about how theologians might use AI, it might be that what I've just said is so sobering that you might think that really our best advice would be just to leave well alone. But I've mentioned already that what was startling and incredibly technological yesterday can very easily become the simple background to life today or tomorrow and it seems to me that Christ's injunction that his followers should seek to be as ‘wise as serpents and innocent as doves’ is a good one. It's there; it's not going to go away; other people are going to use it. I think the imperative is for theologians, people of religious convictions and those who don't, to think about how to use this in a wise way and innocent way. So, a saying attributed to one of the Booths, the couple who founded the Salvation Army, they put together a hymn book and they brought in quite a lot of modern tunes and they asked, ‘Why should the devil have all the best tunes?’, meaning ‘Why couldn't popular music be used in the worship of God and drawing people into the church?’ and it seems to me that as a religious person myself I would want to be asking, ‘Why should the devil have all the best AI?’

So, what might some of the examples be of how AI could be used in theological study? Well, a lovely one that's come up recently has been to decipher the text in various scrolls, ancient repositories of learning which have got burnt and can't be read any other way. So particularly in Herculaneum, which is just near Mount Vesuvius there's at least one Library, there may be others, and when the lava and the ash descended and killed so many people, these scrolls containing all sorts of things, some of which have been lost to the world were turned into basically just carbonized lumps. And when these were discovered, and for the last 100 years or so, people have been trying to unroll them but for the most part they've just been destroyed any of these scrolls people have tried to pick apart. Well, it turns out that using scanning techniques like computerized tomography which you use in medicine, you can work out the structure of these scrolls but they're still very irregular, they're tightly bound, they've got completely deformed by the burning process and the challenge is to be able to trace within this data that you've got where the individual pieces of parchment might be or how it rolls around in a in spiral. And it turns out that AI has been used to do this. Just in the last year there have been some wonderful stories about this so if you give your AI model all the data that's come from the computerized tomography scan it can work out where the roles of the parchment are and virtually unroll it for you and we're able to read things that haven't been read for a very long time and no doubt amongst that Library there will be things that are of interest in the history of religion. And who knows, there could be all sorts of other scrolls which contain things that we're only going to be able to read by using this amazing artificial intelligence technology. That's a bit of a compressed description of that but if you put Herculaneum scrolls and AI into your search engine you'll find some wonderful stories.

Also, of course, all sorts of important things that one could do with AI around translation and learning languages. I'm reasonably optimistic about this. I think that it's good for us to have access to text that we wouldn't be able to read otherwise and I can see that it might seem like using AI translation widely

will take the initiative away from learning other languages but there's nothing like knowing a language for being able to really appreciate the text that you've got in front of you and I would say that AI could be a really useful tool in learning languages. So, when I'm reading something in a language that I know somewhat then, nowadays, I'll ask ChatGPT about some word that I don't quite understand and it will tell me what part of speech it is and perhaps produce a whole chart of that particular verb and how it works. And I have found it to be, just even the last month or two, really quite a useful thing to interact with in trying to improve my languages alongside having a human French teacher who I speak to each week, for instance. So, I like the idea of us being able to deploy AI to make individual language tuition because it's something where you just need input all the time, do a little bit every day if you can. I don't want to put human language coaches and teachers out of work, not at all, but if what we could maybe get once a week with them could be supplemented by AI systems, I think that would be really impressive.

And another example of AI that I've seen used in theology is working out the story of how a collection of manuscripts were copied. So, we're often very interested in finding out of all the manuscripts that have survived of something around the world in various libraries what's the original text or at least how close can we get to the original text because people copied these things out. That's wonderful; they disseminated the texts so because they copied them, they've survived but every now and again they would make a mistake and it's a really important part of scholarship to have a sense of all of these different variants and trying to work out what the most plausible story is, we can tell about where we started from and then which manuscript was copied from which other one. And this is exactly the sort of thing that artificial intelligence is good at: finding patterns in data and it is already being used to trace the story of how manuscripts are passed on and help us to get as close as we can with confidence to original texts.

I've mentioned already the importance of ethics and thinking about this ethically we are in really interesting, but also contested, times about the use of AI. Certainly, anyone who works in a university at the moment is interested in how it might be used in assessed work and exams and that kind of thing and the important thing is that we know what our professional or institutional standards are, what's expected of us, and we can play our part in developing those standards and shared ideas; about what makes for good and disreputable use of AI. But I speak to myself and I speak to everyone at every stage of their studies. The important thing is not to be afraid of using this technology but also to know what the expectations are: morally, ethically, and, of course there's an extra set of ethical questions for us to attend to and I would hope that the church would be interested in this because it's historically been interested in in poverty and improving the lives of people around the world: there are questions about access to this technology. So, I have wonderful access to it and the doctors that might treat me in this country, in the United Kingdom, probably have access to the latest technology, for instance, for reading ECGs. (AI turns out to be really good at looking at ECGs which are about the operation of the heart and working out what's going on.) I suspect that people in many of the world's poorest countries *don't* have access to this, either professionally like I do or in terms of health care or very many other ways at all. So, one of the things that we might be interested in thinking about in ethics is a kind of just access to this technology. In a way you might be expecting that I would talk most of all about AI and Christian doctrine or central Christian beliefs. That's what I work on above all else but here I want to say that I think, as I've said before, there's a bit of a danger of our thoughts being captured by ideas of artificial *general* intelligence. So, systems that, like us, are intelligent in lots of different ways, and the questions about

could that go on towards we call ‘sentience’ or ‘consciousness’ or something like that. So if you have that sort of computer system in mind you can ask questions like “Could I have a soul?’, ‘Could it be made in the image of God?’, ‘Could it sin?’, ‘Could it be redeemed?’, ‘How might it relate to Christ?’, ‘Could it have an eternal destiny?’ and so on. And I suppose those are all good questions. But for my part I want to say we should probably pay the most attention to AI as it really is *now* or as we have every reason to suppose it will soon be, and I'm not sure that we will *ever* give rise to a computer system that is conscious in that way. It's an interesting question, whether we think it's possible. I'm not averse to it being possible; maybe it is. I just think that it would be easy for us to ask those sorts of questions and ignore the ways in which AI as it is *now* is both changing the world and is itself raising all sorts of fascinating questions.

So, I think about the doctrine of Creation in particular, I think it's really significant that AI is able to look at the data about the world, data about human culture, and is able to, as I said, really fit, be able to extract the meaningful order of everything that's pointed to it. I think it's significant from a theological perspective. We theologians tend to think that the world and human culture is deeply meaningful and AI is able to latch on to that and see the patterns that are out there. I think that's really important. And in terms of creativity, the very fact that human beings can do this, can make these sorts of computer systems, seems to me an important expression of human beings being in the image of God. God the great Maker. To be in God's image, is amongst other things, to be a ‘maker’ so we can explore this idea of human beings made in the image of God and the sort of creativity that we show in making artificial intelligence systems but of course whatever you make bears the effects of the of the kind of thing you are. People tend to pass on what they are into what they make and so there are all sorts of questions about Christian ideas of sin and whether the things that we make are likely to bear the trace of our sinful impulses.

I'm also interested in the way in which thinking about AI brings mathematics back into theology because AI is at its roots an exercise in theology and one of the things that's been most fruitful and interesting for me in getting into this topic has been to reconnect with some of the maths that I learned as a scientist. And this is, in fact, something that's really quite medieval. In the Middle Ages if you were on your way to being a theologian you ould have spent a few years, in fact, thinking about mathematical topics of various sorts and nowadays we tend to let that drop out, at least for the study of the Arts and Humanities. But in the medieval mind there was a way of ascending towards God as beautiful and ordered and good and true which came through mathematics and I quite like the idea that perhaps the study of mathematics might come back into theology at least for those of us who are interested in thinking about AI.

I'm going to end by pointing to two possible angles for thinking about theology and AI, thinking about AI in theological terms. One of them is work that's only just come my way and my colleague in my new faculty, Hindy Najman, has done some fascinating work on ideas of authenticity and the way in which things can be copied or derivative in biblical thinking and in later texts that have been working with Biblical ideas. We're very interested today in copyright and rightly so. I've mentioned that already, but in the ancient world there was no such thing and quite often people would write things and pass it off or associate it with some other tradition. So there's an important work by Dionysus the Areopagite which is actually a good few centuries into the Christian tradition but was described as being by a figure in the first century and had quite a lot of authority. Because of that there are various books in the Bible which

go by the name of authors, where perhaps it's more work of something in that community rather than by the author itself. And later on there are all sorts of, like I say, with Dionysus, all sorts of books that are described as being by writers that have a biblical authority. What I find really interesting about this work in biblical studies is to say that we should get out of our contemporary mindset and try and understand what these themes of derivation or what we would call plagiarism or authenticity/inauthenticity, what they look like in the ancient world. So, I've got a quotation here, "The connection between creative continuation, impersonation and interpretation is well evident across many different literary and religious traditions. Attribution, (so, saying that a work was by somebody) was increasingly understood to be an interpretive gesture rather than an historical claim to authenticity. As these writers point out, "Much ancient writing involved imitation." In fact, much writing down the centuries has involved imitation. So, I think there are really important ethical questions about the way in which AI works in a derivative way. It chews up what lots of people have written and does new things and there are all sorts of questions there that are not to be treated lightly. But I found this discussion of themes of derivation and inauthenticity, and working within existing traditions, and taking things and rehashing it as being really quite useful in thinking about these topics that come up in AI. Of course, it might be that we think that the way in which computer systems use existing material and rehash it to make new text, as we saw in that chatbot earlier on, is actually not very much like human examples at all but if so, that itself is worth thinking about. What would be different and there may well be important differences between how one human being can live within the tradition of another and reuse ideas and in the same way in that process contribute something new. How would that be different from the way in which a computer perhaps simply cannibalizes what other people have produced without adding anything new of its own?

And finally, I want to talk about my most recent paper on this topic of AI which has been about thinking of AI as a tool or instrument. So, I work quite a bit on Scholastic theology in philosophy which grew up in the Middle Ages but continues right to the present day. And one of its figures, Thomas Aquinas, thought quite a lot about how we use tools, so he had both a sense of tools that put things in their place. So, it's a sculptor that carves a sculpture and not primarily the chisel. It would be strange to say that the chisel produced some work of Michelangelo and forget that Michelangelo was involved in it. And yet for his time particularly, against the backdrop of his time, he had quite an elevated view of tools. He recognized that tools allow us to do things that we couldn't do without them. He recognized that they do play a real part in what we produce. He recognized that there's something about the nature of the tool which leaves a trace on the kind of thing that's produced. So, I found this idea of thinking about AI as a tool quite important. I about talked earlier on about the idea of thinking of AI, artificial intelligence, as actually *extended* intelligence; it's *we* that remain the intelligent users but we use these systems as tools or instruments to be able to allow us to be able to do things that we couldn't do otherwise. And one of the really useful things I think about Aquinas on tools, which is definitely relevant for thinking about AI as a tool, is that he says tools are for the worker. We get it completely the wrong way around if we said that the worker was for the tools.

So, I'm going to end with some Thomas Aquinas and what he has to say about the value and place of tools or instruments. I think it's a good example of how our religious traditions will sometimes take a step back, think quite carefully and philosophically about something in their own time and here he didn't have particularly advanced tools or instruments in mind and yet with his wonderful, analytic, careful, well-informed brain he was able to say some really useful things about the nature of tools. In

fact, a great deal more than I've got time to put on this slide but there are some points here about what it means to use something as a tool which have really important ethical consequences, I think, real ethical insights. He says, “A tool is not valued for its own sake but because it's useful for the person who uses it. So it should be that the careful work that is devoted to tools must actually be done for the sake of the human agent who uses them, whose use of the tool is the goal of what we're doing and not the tool itself.” That's quite a profound insight, I think. And in all this work that's been done to develop AI, we should recognize that it's only a tool and that tools are for the worker, for the sake of the one who uses it and we should develop them in such a way as to benefit us and our best goals.

There's another quotation from Aquinas about tools there. I'll leave that for you to read and think about.

It's been a pleasure to give this lecture on some aspects of how artificial intelligence might feature in theological thinking from all sorts of different perspectives: ethics, philosophy, Christian doctrine, some insights from biblical studies. And I hope that, as well as this particular topic, I've been able to give you some sense of theology as a vibrant and really consequential discipline in the 21st century and I wish you well for your studies at A level and also encourage you to think about the possibility of studying this fascinating subject or actually very large cluster of subjects at university level. Thank 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.