Oxford Mathematician John Lennox on AI

https://mindmatters.ai/podcast/ep236

Announcer:

John Lennox is a professor of mathematics at Oxford University, internationally renowned speaker and author of several books on the interface of science, philosophy, and religion. On this episode of Mind Matters News, Dr. Lennox answers questions about artificial intelligence and the transhumanist claim that AI will turn humans into god's. Let's listen in.

Question:

What is artificial intelligence and where do we find it?

John Lennox:

Artificial intelligence is a very rapidly developing field of computation today. It splits into two types. There's narrow AI and there's AGI, artificial general intelligence, and it's very important to keep them separately. So let's look at narrow AI. That really is the AI with which we are all familiar. It's the stuff that actually works. And the typical AI system consists of a very powerful computer with a large database and an algorithm that recognizes particular patterns within that database.

Take a simple example where in the midst of a Covid pandemic at the moment, and it affects people's lungs. So let's imagine that our database consists of 1 million x-rays of diseased lungs. They are labeled with their respective diseases by the top medical people in the world. That's our database. Now, a photograph is taken of my lungs if I catch Covid 19, and the algorithm compares the photograph that's taken of my lungs with the million in the database and comes up very rapidly with a diagnosis.

Now these days that diagnosis is likely to be much more accurate than I would get from my local hospital. And systems like this are being employed to find vaccines for Covid 19 and do many other wonderful things. So there's a very positive side to this kind of AI. The second thing to mention about it is that characteristically such a system does one single thing that normally requires human intelligence. That's why it's called artificial intelligence. It's not real intelligence. The machinery, the computer doesn't think. It's not conscious. There's intelligent input from the doctors and the computer scientists and the people that developed the algorithms. What comes out looks as if the diagnosis looks as if it's been made by human intelligence, but it hasn't. It's been made by this sophisticated piece of equipment.

Now, it's very important to grasp that. Now, another example is if we stick with pattern recognition, which is enormously popular, facial recognition technologies have developed a very advanced state, and you can see that they're very useful for police forces to pick out a terrorist or a criminal in a crowd and be used therefore to capture people who are a danger to society.

But there's a downside because facial recognition technology is being used at the moment in certain parts of the world to invade the privacy, not only of individuals, but of whole people groups and actually control them and suppress them. Now, I mentioned that example to say that very rapidly AI, narrow AI raises huge ethical questions.

Now remember, this is the stuff that's actually working, self-driving cars, autonomous vehicles, AI system built in there, but you have to build into it some kind of ethical decision making. If the car sensors pick up an old lady crossing the road and it can see, so to speak, that if it misses her, it'll hit a

cue of children waiting in line for a bus. What does it decide to do? So you've got to program that ethically. And of course the ethical norms that will be built in are the ethics of the people doing the programming. That raises huge questions. Even more difficult is the matter of autonomous weapons and so on and so forth.

Now, the final example I would give you is the fact that we're all involved in AI. That is any of us who own a smartphone, it's tracking us all the time. What many of us don't realize is that, for example, we make a purchase at Amazon. A few days later, we'll get a pop-up saying, people that bought this book were interested in that book and they're trying to encourage you to buy something else. Okay, we put up with that. But what we don't often realize is we are feeding a vast amount of information into that system, and large companies are selling that on to third parties without our permission. And this has become a real problem, and it is really what is called by Shoshana Zuboff who's an expert on it and has written a book talking about it called Surveillance Capitalism.

So we have surveillance in that sense, which is economically vastly profitable. We are contributing without our permission and we're doing this voluntarily. Now, this is eventually going to hit us really hard and we begin to see that there are problems with it. What I'm saying here is AI is here to stay. It is going to rapidly be, if not already, a trillion dollar business per annum, and it raises vast problems in every direction as well as doing a lot of good.

And as a Christian, I encourage young, scientifically and computer literate people to get into AI, firstly, to do the good stuff and develop, say, medical research that will help humanity, but also to be there to make a contribution to the ethics, which is very necessary. That's a huge topic on its own. So that's narrow AI.

AGI on the other hand, is basically represented by the quest for super intelligence that has two parts. Firstly, the idea of enhancing current human intelligence by bioengineering, by cyborg technologies, by drugs and all kinds of things to create a super intelligence based starting with human beings. That's the agenda of one of the best known people on this topic who's not a scientist. Yuval Noah Harari. In his book, Homo Deus, The Man who is God. And what he is trying to do is achieve the transhumanist agenda, which really is too pronged.

First, he feels that the problem of human death will be by technology fairly soon, and therefore, the main effort of artificial intelligence and related technology will be to enhance human happiness so that we turn human beings into God's small G, Greek style gods having super intelligence. Now, the other research is on trying to replicate human intelligence on silicon to make it independent of organic material. And it's in that kind of context that you get notions of people hoping that their brains will one day be capable of being uploaded onto silicon. And so they live forever.And we are all aware of people who go to great expense to have their brains frozen on death so that they hope that one day they will be woken up and their brain contents will be put onto silicon.

Now, all of that is highly speculative. There's much more science fiction than science. However, there are scientists, leading scientists like our own astronomer, Royal Lord Rees, who think that we are going to change the nature of human beings so that within a few hundred years the beings that exist will have no knowledge of us really at all and no idea what we are like. And that has created not only speculation, but great concern. Elon Musk is worried about it. Somebody else said that in developing this kind of stuff, we are summoning the demon.

So AGI is largely the stuff of science fiction. They tell us that there's a singularity coming. That's Ray Kurtz file who's a guru in this field. This is now not the singularity of the original big bang. The singularity is when human beings become merged with technology and the technology essentially takes over. Now, I see huge problems with this, but you ask me what it is, and that is what it is.

Question:

Does the rapid growth in computer processing contribute to these ideas of technology taking over?

John Lennox:

Computer power is increasing rapidly, and there's Moore's law that says it roughly doubles every 18 months or so. There will be an eventual limit to that, but the idea is we'll be able to go so far that we'll be able to solve all these problems. But I have very big questions about that. One of the deep questions in this whole field is the connection between intelligence and consciousness. You see, I said earlier that artificial intelligence is artificial. There's a wonderful paper written by one of the pioneers in the topic, Dr. Professor Millichamp who said, this is the title of his paper. I think it's marvelous. The artificial in artificial intelligence is real.

In other words, it really is artificial. And therefore some argue. Look, we don't need to develop consciousness, say on a base of silicon. All we need to do is replicate human intelligence. In other words, get something to pass the touring test that it appears to other people interrogating the system that it is human, but it needn't be identical. It needn't actually be conscious at all. And John Searle in Berkeley some years ago pointed out that a system could appear to be intelligent without being intelligent at all. It is famous Chinese room experiment.

Question:

What is Johns Searle's Chinese room experiment?

John Lennox:

Well, the idea is that you've a closed room and a man inside it, and he has got various instructions as to how to deal with Chinese symbols. He doesn't speak any Chinese, and people come up to a hatch in the door and they feed in a question in Chinese. The man in the room doesn't understand it, but he compares it with all the instructions he's been given. And the instructions say, if you see this, then hand out that and he hands out something he doesn't understand either, which is the answer to the question that was fed in. So the people outside think that they've had an intelligent being responding to their question when they have a no such thing.

Question:

Will a computer ever be conscious?

John Lennox:

How could you answer something like that when we have no idea what consciousness is? Consciousness is the huge barrier. No one that I've ever spoken to or read has any idea what it is. So speculation is irrelevant. I doubt it very much because computers appear to have certain limits. Roger Penrose is very interesting on this topic. He argues that you will never be able, even in principle, to simulate the activity of the human mind on a computer because the human mind can do things that a computer cannot do. They're not computable. And one of them is, for example, in mathematics, the proof of girdle's famous theorem that arithmetic is not complete.

Now these are very complex things, but it's very interesting to see that someone as bright as Penrose with his recent Nobel Prize argues that no, it isn't computable, so computers will never be able to do that.

Question:

Can you explain the AI concept of transhumanism?

John Lennox:

The origin of the word transhumanism, interestingly enough, is not secular at all. It wasn't first used by a scientist, but it was used in a translation of one of the books of Dante, his Paradiso. It occurs in a passage where Dante is trying to imagine the resurrection of his own body. And he says, I quote "words may not tell of that Transhuman change". And the political scientist and author of The End of History, Francis Fukiyama, he regards transhumanism as the world's most dangerous idea because it runs the risk of infringing human rights.

Now, I would go further than that. I would say this is where the great danger of AGI and its transhumanist agenda resides. We are almost in the first generation that is capable not only of implanting chips in human brains and so on, but actually of altering the fundamental specification of human beings by reconfiguring their germline and therefore determining the kind of beings that will exist in the future.

Now, one of the reasons I wrote my book, 2084: Artificial Intelligence and the Future of Humanity was precisely this, that AGI whose agenda is largely driven by atheists so far as I can see, is raising perhaps the deepest question of all that we face today. What is a human being? What dignity, what value, what rights does a human being have? And many years ago, C.S. Lewis wrote two books that I think are seminal for understanding this. The first one is called The Abolition of Man, and the second is a science fiction novel called That Hideous Strength. And he talked about the danger of a group of scientists being able one day to alter the very specification of human beings. And he points out that what they will make are not human beings, but artifacts made under their specification.

So he says their final triumph will be the abolition of man. And that is what concerns me because the law of unintended consequences is likely to have huge effects, and there is danger in this world that rogue scientists will start taking human embryos or human beings who exist and experimenting on them and producing a race of freaks that is irreversible. So all of that tells me that we need to think very seriously about human nature, and that's where my Christianity comes in. I want to bring the biblical worldview where we find that humans model one, mark one humans like we are, are given an infinite dignity and value because we're made in the image of God. And I believe we need to bring that insight to the table and bring I think a bit of reflection on the biblical worldview into the situation.

Question:

Are there things humans can do that AI will never be capable of?

John Lennox:

It's always very difficult facing questions what do you think AI can do and cannot do? But I do think that Penrose has got an argument that artificial intelligence by definition will not be able to do all the things that conscious intelligence can do because in the way stands the problem of consciousness. And until somebody tells me what consciousness is, I will remain a very deep skeptic about these possibilities.

Now, that is not to say that AI will not solve a lot of problems. It is already even narrow AI solving a lot of problems. But speculation about the advent of a super intelligence is a different matter except for one thing. Now, the interesting thing is that physicists like Max Tegmark of Princeton, he has written about AI and he has set up, interestingly, a whole set of scenarios. What happens if we create a super intelligence? Is that going to be the last invention of humanity? Because the super intelligence will

simply take over, will it be benevolent or will it keep human beings, a few of them as pets, or will it simply regard them as a nuisance and destroy the lot? And there's a whole range of speculation of what happens when man is turned into a God like that.

Now, I find that interesting, why? Because the idea of human beings becoming gods is extremely ancient and deeply rooted in almost every culture in the world. And it begins on the opening pages of the Bible where the fundamental attack on humanity comes from the suggestion that if the humans will only disobey God and eat the forbidden fruit of the tree, of the knowledge... of good and evil, please note not the tree of knowledge. God is not against knowledge, but if they eat that forbidden fruit, they shall be as Gods knowing good and evil.

They did eat it, and they discovered that knowing good and evil is an awful thing to know, but that infected the whole world and it misrepresents God so that many people today think that God is their enemy if he exists at all, and that he's trying to suppress them and not trying to rise to their full potential.

Now, the interesting thing is this, the transhumanists, they think that we are going to use AI and other technologies to raise humanity to the full potential of being Gods. We're going to solve the problem of human death in the sense that low humans can die. They won't have to die. Of course, that's going to be only for the very wealthy, but we'll leave that aside at the moment and then the hope that will upload our brains and live forever. But these are utterly fascinating things. Now, they're not going to solve the problem of human death by technology.

I'll tell you why. Because no technology, and it's been tried before. The Soviets tried to create a new man and they tried to make utopia and other utopias have been attempted. They've all failed. Why? Because they bypassed the problem of sin and evil that's endemic to humanity. Now, the biblical, and here's the irony of the whole thing. I read Harari and his books sell in millions. He says, we're going to solve the technological problem of death. And I say, you're too late. The problem of death has already been solved 20 centuries ago. Jesus Christ rose from the dead. And not only that, he promises that if we face the endemic problem of human evil and sin and repent of the mess we've made of our lives and the lives of others, then he will give us new life that will transcend death. It'll even transcend Covid 19.

So the solution to death is already there for the believer in Christ. And then when it comes to the uploading, and here it's very ironical, they're hoping that by technology we're going to upload brains onto silicon. Well, you don't need to do that because Christ promised those that trust him that he will return for them and upload them into heaven.

Now, when I compare the two kinds of uploading the transhumanist one and the biblical one, I will choose the biblical one because there are far more evidence for its truth. Now, that's a subject of course for another time. But the point is that this whole program is a parody of what the Bible teaches, and that's why it's very interesting. And the biblical worldview within it is contained the fact that all through history, human beings will make the attempt to build their way to heaven.

Barbell the Tower of Babel was a famous example. Many of the ancient emperors claimed divine honors. The Caesars claimed divine honors. And even in recent times, we had people that seemed to claim divine honors. The Soviets tried to build a new man. Who knows what we're going to try to do. And perhaps the chilling thing is the Bible indicates that in the future there will be a man who will claim to be God and will do all kinds of apparently supernatural things to get the whole world into as power economically.

Now, that sounds exactly like one of Tegmarks scenarios, and my plea is very simple. If people today are going to take Tegmark scenarios and the scenarios of other people, 1984 and so on, seriously, then I would just invite them to have a look at the biblical scenario and compare it with the others and see

how much sense it makes because I think it makes tremendous sense. And because of that, there has been enormous interest in the book that I've written that risks going into these areas of Christianity that aren't often talked about in the public domain.

Announcer:

This has been Mind Matters News. Explore more at mindmatters.ai. That's Mindmatters.ai. Mind Matters news is directed and edited by Austin Eggbert. The opinions expressed on this program are solely those of the speakers. Mind Matters news is produced and copyrighted by the Walter Bradley Center for Natural and Artificial Intelligence at Discovery Institute.