John Lennox: How Will Artificial Intelligence Impact the World by 2084? (https://mindmatters.ai/podcast/ep89/)

Robert J. Marks:

How will AI transform the world by the year 2084? That's the topic today on Mind Matters News.

Introduction:

Welcome to Mind Matters News, where artificial and natural intelligence meet head on. Here's your host, Robert J. Marks.

Robert J. Marks:

Greetings. In 1984, George Orwell forecasted dystopian future under communism. His classic book was titled 1984. Many of Orwell's predictions about communism were proven. What will be the effects of AI a century later in the year 2084? Replacing George Orwell is Dr. John Lennox who has written the book 2084: Artificial Intelligence and the Future of Humanity. How will AI, not communism, affect the future? Dr. Lennox is able to look at the AI phenomena from a number of different perspectives. From the technical side, he is an Emeritus Professor of Mathematics at Oxford University. On the theoretical side, he is a Pastoral Advisor of Green Templeton College at Oxford University. Dr. Lennox, welcome.

John Lennox:

Thank you very much. I'm delighted to be on your program.

Robert J. Marks:

That's wonderful. This is the second time we've talked and your first interview had a lot of traffic, so it's really great to talk to you again. You have written a book entitled 2084: Artificial Intelligence and the Future of Humanity. The first obvious question to ask is why did you write this book?

John Lennox:

Well, there's a sense in which I have been interested in futuristic scenarios for a long time. I'm not a sci-fi addict, but I was deeply impressed by CS Lewis' sci-fi trilogy. He raises the question in the third of those books, That Hideous Strength. He imagined scientists trying to increase their power by preserving a human brain. As I read that book and saw the issues it raised, it put into my mind the idea that it might be important to think through this stuff as it develops as it has, but the immediate reason for writing it was that I was asked to give a lecture on the topic in connection with the book of Genesis. I said, "Look, I think you come to the wrong person." They said, "No. Well, we think you're the person to do this." Well, I decided in the end to do it that because it initially had to do with artificial intelligence and the nature of humanity.

John Lennox:

As I started background reading of various people, I discovered that a lot had been said, a lot had been written, but there was a real need in my view of evaluating it. It has ended up with this book.

Robert J. Marks:

I see. Could you give us a quick overview, a thumbnail sketch, of what the book is about?

John Lennox:

Well, the book really has several purposes. I want to demystify the good side of AI so that people, particularly Christians, but not only, are not afraid of it. Secondly, I want to take some of the hype out of the science fiction side, artificial general intelligence, so we need to distinguish between those two things. That's very important and then discuss, for both kinds of artificial intelligence, the ethical, and philosophical, and indeed theological questions that are raised by them.

Robert J. Marks:

Very good. In talking about AI duplicating humans, duplicating us, it seems important to ask the question, what does it mean to be human? We're looking at AI possibly duplicating humans. I don't believe that will ever happen, but what do you think? Do you think that AI will ever have the capability of duplicating a human?

John Lennox:

Well, I doubt it like you because I think this is where we need to make the distinctions very clearly. Narrow artificial intelligence tends to do one thing superbly well that normally takes human intelligence to do, but the machinery and it consists of a computer with the capacity to dig into a large database and recognize patterns. There that's impressive and there are wonderful examples, particularly in medicine of it working very efficiently.

John Lennox:

The second kind of artificial intelligence, AGI, artificial general intelligence, is really the quest for super intelligence, so one of two kinds, either enhancing human beings as they exist and building a biological super intelligence or else discovering ways of uploading or downloading the content, say, of the human mind onto silicone so that we remove the dependence on an organic substrate. There, it seems to me, that the likelihood of building a super intelligence that exceeds human capacity in every direction is very slim because human intelligence is conscious. We don't know what consciousness is. No scientists knows what consciousness is. The most serious people recognize that that is a huge barrier. How can you build a conscious being when you don't even know what consciousness is?

Robert J. Marks:

Yes. I think there's different theories of consciousness. One is called panpsychism, I believe, which is the idea that since you can't explain it, it must be a property of the universe and that everything has a degree of consciousness. For example, a refrigerator, just because it has mass, it has energy, but it also has a consciousness associated with it. You have any thoughts on that and panpsychism?

John Lennox:

I have very few serious thoughts on that because it seems to me to be a misuse of the word consciousness. Once everything is conscious, the word loses its meaning. A refrigerator is not aware of itself. We are. There seems to be a vast difference between the two. I suspect there's very little mileage in eroding the concept so that everything could be said of consciousness and so that one can say, "Well,

there's no mystery." Of course some philosophers, I think Daniel Dennett, may be one, they answer the question by supposing it doesn't exist.

Robert J. Marks:

Yeah. I tell people I know I'm conscious. I'm not really sure about you. It's a hard thing to communicate. It's a hard thing to define, as you said. I don't believe artificial general intelligence will ever be developed, but if indeed I'm wrong, should AGI be given rights? You talk about this in the book a little bit. I thought about it and it turns out that there are non-human entities that are given rights, such as corporations. They're treated as a nonhuman entity. Do you believe that artificial general intelligence will ever mature to the point where it's given rights, like a freedom of speech or whatever?

John Lennox:

Well, again, I think consciousness is the big barrier. What do rights means given to an entity that does not appear to be in itself a moral entity? Now why I put it that way is the fact that even narrow AI, for example, we are involved these days in very rapidly developing self-driving vehicles controlled by artificial intelligence. The programming of such vehicles raises all kinds of ethical questions. For example, if an artificially intelligent, self-driving car is coming a road, and its sensors pick up an old man crossing the street, and on the other side, there's a line of children waiting for a bus, and it's going to hit one of the two, how does it decide between them? There are huge ethical questions.

John Lennox:

The problem is that the machine itself is amoral so that the ethics that's built into it have got to be the ethics of the programmers. Therefore we have to inquire what right have they and on what principles will they function? That's a concern to many people who are working at the higher levels of implementing AI in great corporations and so on. What ethical principles are we going to go by, because one thing is very clear to everybody, the technology is outpacing the ethics by a very large factor.

Robert J. Marks:

Interesting. I agree with you, by the way, that ethics is never the responsibility of the artificial intelligence. It belongs to whoever did the programming and who was ever the end user of it. You also mentioned self-driving cars. I think that there we have a case of the failure of a type of general artificial intelligence that has failed to deliver what it was promised to deliver. If you can't do self-driving cars, people and corporations are abandoning the research into self-driving cars, how can we ever have an artificial general intelligence? It's going to be an interesting thing to watch. As all these things happen, John, technology is changing everything. I had withdrawal symptoms a couple of weeks ago when I was separated from my cell phone. You just rely on these things. In what sense do you think AI technology will change what it means to be human?

John Lennox:

Well, it is going to do certain things. It's already doing them. One very important example of that is the change in the job structure in society because, as more and more processes get automated, and they reckon in the United States, I think within 20 or so years, at least 30% of current jobs will fall under artificial intelligence and therefore will no longer exist for human beings.

John Lennox:

It's going to create a huge problem of re-employing such people because employment will increasingly depend on advanced knowledge of technology and not everybody's got that. Now I'm speaking of a highly developed nation like the USA. I was recently in South Africa and they were saying that AI is creating a huge problem because when people lose their jobs there, there isn't the educational infrastructure to educate them to take over these new technological positions. That's creating a problem.

John Lennox:

The thinking members of the community are asking questions about what does it mean to have a large proportion of your culture unemployed because from the biblical point of view, and I think this is very important because psychology confirms this, that jobs mean a great deal to people. They supply part of the meaning of their lives. If you deprive them of jobs, the recent research has shown it can lead to all kinds of pretty serious psychological problems. We've got to think that through. You can repeat this in all kinds of other spheres, but that's one example.

Robert J. Marks:

Very interesting. It's been said that AI is the new electricity. It's neither good nor bad. You have addressed some of the potential negative uses of artificial intelligence or the negative impacts of artificial intelligence, but expanding on that, what are some of the big threats that you see in the use of AI technology in the near future?

John Lennox:

Well, the threats are best explained by comparing them with the advantages. Let's take a very simple and practical example, which is extremely useful. That is in the field of x-ray technology. Let's imagine that we construct a database of a million x-rays of people's lungs and we get the best medical experts in the world to label the diseases on those images. That becomes a big database for a computer AI system. Then you or I, we get problems with our breathing, an x-ray is taken. In a few seconds it's compared with a million pictures in the database and out comes a diagnosis. Now, we're already at the stage where such diagnoses tend, on average, to be quite a bit better than you would get in your local hospital. There's a positive thing.

John Lennox:

Now let's go down the scale and think of the AI associated with our smartphones. We buy a book and a few days later up pops a little message that says, "People that bought that book also are interested in this book," and your attention is drawn to buying the second book. Well, that can be very useful or it can be very irritating. What many people do not realize is that that system is actually harvesting a great deal of information about us, about where we go, who we meet, what our buying preferences are. It's being sold on to third parties without our permission. This is what is in a way called surveillance capitalism. Shoshana Zuboff who's an emeritus professor at MIT has written quite a chilling book about this as being very dangerous for society, this theft of our property, our data property, without our permission. The problem is that we willingly wear these trackers and we're sacrificing, in a sense, quite a lot of our privacy in order to gain these alleged advantages. Whether they are advantages or not as something that we need to seriously think about before we're engulfed by it.

Robert J. Marks:

One of the big negatives about artificial intelligence you also mentioned in your book was the use of facial recognition, especially currently about China. What I found really interesting is that there's an intersection here between your book and Orwell's book. With the face recognition and all of these things happening with our loss of privacy, we are living in an age of big brother. I think that they have a long way to go. I share my amazon.com account with my daughter who also orders off of my account. They haven't figured out I don't need ads for baby diapers yet. They have yet to differentiate between my use and other people's use. It's far from perfect. Dr. Lennox, thank you. This has been a lot of fun. We've been talking to Dr. Lennox. His new book is entitled 2084: Artificial Intelligence in the Future of Humanity. We're going to provide, on the podcast notes, a link to his website. Dr. Lennox, if I remember, the site is really simple. What is the site for your book?

John Lennox:

Yes, it's 2084book.com

Robert J. Marks:

2084book.com. Please visit there. There's a very nice video that tells us a little bit about the book and you can also order it there. I'm sure like everything else in the world it's available on amazon.com.

John Lennox:

Yes. It is.

Robert J. Marks:

Excellent. Excellent. Do you know if it's out on audio yet? I'm a big audio fan.

John Lennox:

I think it may well be. Of course there's lots more stuff on my own website, Johnlennox.org, but I'm pretty sure it's on audible and it's on Kindle. I know that.

Robert J. Marks:

Okay. Wonderful. Wonderful. Yeah, if you do go to amazon.com, I encourage you to look at some of Dr. Lennox's other books, which are also excellent. I recommend the book. I've read it. It's a very easy, but it's a very deep read. I think you'll enjoy it. Until next time, be of good cheer.

Conclusion: