AI: Is it Good or Bad for Society?

https://mindmatters.ai/podcast/ep240

Announcer:

Robert J. Marks is a distinguished professor of electrical and computer engineering and director of the Discovery Institute's Bradley Center. His latest book is Non-Computable You: What You Do That Artificial Intelligence Never Will. On this episode of Mind Matters News, Dr. Marks separates fact from fiction about artificial intelligence.

Robert J Marks:

My name is Robert J. Marks. I'm a distinguished professor of electrical and computer engineering here at Baylor University. I'm the director of Discoveries Walter Bradley Center for Natural and Artificial Intelligence. I've been involved in artificial intelligence, specifically neural networks for over 30 years. I was one of the first editors of the IEEE transactions on neural networks. IEEE is the Institute of Electrical and Electronic Engineers. It's the largest professional society in the world, and I was also president of the IEEE Neural Networks Council for a number of years. This has been some time ago, so I've been involved in the area of artificial intelligence for again over 30 years.

Question:

What is artificial intelligence and where do we find it?

Robert J Marks:

Artificial intelligence depends on your dictionary. If you go into academia and go into the different disciplines, you'll find skirmishes as to what AI means. There's kind of synonyms which are used for AI, such as computational intelligence, machine intelligence, and to those in the field, this means a lot. They do subcategories and things of that sort. I think for purposes of our discussion, we should probably treat it just like the media treats it, which is any gee, whizz thing that can be done by a computer. If a computer does something and it's just miraculous and you go, "Gee, whizz," then that's probably AI and I like it for purposes of the discussion because what it does is it brings AI to the level of the computer, so when we talk about the capabilities of AI, we can talk about the capabilities of computers and if a computer can't do something, then AI is doesn't have the ability to do it either.

Question:

Where do we find AI?

Robert J Marks:

Artificial intelligence is indeed ubiquitous. In fact, it's around us and we're kind of numbed by familiarity. We see AI, for example, in things such as banking. It does automatic pattern recognition on checks. I can take a picture of my check and it does automatic optical character recognition to deposit my check without me going into the bank. It does facial recognition that allows me to log into my computer and my cell phone. It does voice recognition and Alexa and Alexa also does incredible things about finding exactly what you're trying to look for. It's not always successful, but it gives it a good try.

We see it also in places like Google Maps, in Uber, in accounting. AI made incredible inroads to accounting and it's changing the landscape there because it turns out that a lot of businesses that were using regular accounting before are now using artificial intelligence and you no longer need CFOs in accounting department, you can outsource it to a third party that uses artificial intelligence and does all of your accounting really, really quickly. We see medical applications and it's just applications in Bitcoin and blockchains and even Google search if you think about it as a type of artificial intelligence, so it is ubiquitous and is everywhere around us.

Now, there's been some landmark sort of occurrences of artificial intelligence. Probably the first was when Deep Blue beat Garry Kasparov in chess. Now, that was a long time ago. That's over two decades ago and this was astonishing because a computer could now outperform the best chess master in the world. Then we had IBM Watson take on the game of Jeopardy and take on the world's greatest players in Jeopardy and whip them in a head-to-head contest, so that was really remarkable. More recently, we had AlphaGo beat the world's champion at Go, which is the most difficult popular board game in the world, and it has been the holy grail of artificial intelligence for a long time to be able to beat somebody at Go and so that was an astonishing thing, too. We're hearing new astonishing things happening again and again as AI becomes more sophisticated and we see greater and greater number of applications.

Question:

Is AI good or bad?

Robert J Marks:

Artificial intelligence itself is like electricity or fire. It's neither good nor bad. It's how it's used. You can use fire in good and bad ways. You can use electricity to power your home. You can also use electricity to electrocute people, so it's neither good nor bad, it's simply a tool and it's a very powerful tool and it is going to remain a powerful tool, but it itself has no propensity for being good or bad.

Question:

What are the heady predictions of future AI?

Robert J Marks:

We have a number of often-made claims about artificial intelligence, for example, that artificial intelligence is going to write better artificial intelligence, is going to write better artificial intelligence, and pretty soon we're going to reach the so-called singularity, where artificial intelligence duplicates or exceeds the capability of human beings. This, however, assumes that artificial intelligence can be creative and artificial intelligence doesn't have the capability of being creative. We have to define exactly what we mean by be creative, but it doesn't have the capability, so that problem, that claim is totally vacuous, it has no substance at all.

Well, there's certain laws in the development of computer. One of them is Moore's Law, which says that there's a doubling in the capacity of computers every year, every 18 months or whatever, and it's been held for a long time. However, the exponential increase in the performance of computers is not sustainable. Anytime you have an exponential increase in anything, it has to level out at some point, and so that is going to level out at some point in the future. You also have claims that because of this increase in computational power that we're going to be able to do things that we couldn't do with AI currently. One of these is AI writing better AI that writes better AI that is ultimately going to exceed the capabilities of human beings and this is also not something which is possible

Question:

Is Ray Kurzweil's proposed singularity possible?

Robert J Marks:

Ray Kurzweil has the idea of superintelligence, that AI at some point will be able to write better software, better AI that in turn writes better AI that in turn writes better AI and pretty soon it will exceed the capabilities of human beings. Now, this assumes that AI is creative because if AI creates a better AI program, this is something that the AI program has created and artificial intelligence is not creative.

Question:

Some say we will be able to upload ourselves into computers one day. Is that possible?

Robert J Marks:

There's a claim that we might be able to upload ourselves into a computer and in some sense gain eternal life, I suppose. No, that's not possible. There are certain aspects about humans that are not computable and the aspects of humans that are not computable cannot be uploaded. You can only upload the computable parts of human beings and we have a lot of different qualities that are not computable.

Question:

What things can humans do that AI can't?

Robert J Marks:

There are some things that humans do that computers can't do, nor will they ever be able to do them. There's some obvious ones and those obvious ones include things such as love, compassion, and empathy. Some of the not-so-obvious ones that computers will never be able to generate, an artificial intelligence will never be able to generate, is creativity, understanding, and sentience. These are attributes of humans that will not and cannot be replicated by a computer.

Well, let's take first of all, the easiest of the three, which I believe is sentience, which includes qualia. Qualia is the experience that you have from your senses when you do something. Imagine taking a bite of a peach. It's a nice ripe peach and you bite into it. You feel the juice spurt in your mouth and you feel it on your tongue. You feel the aroma of the peach. You chew it and get all the wonderful juice and you swallow and anticipate the next bite.

Now, try to understand or try to comprehend communicating that to a person who has had no sense of taste and no sense of smell since birth. Now, you could explain to that person, you could explain the chemicals that were on the tongue. You could explain the biology of taste buds. You could explain the Newtonian mechanics of biting into the peach, but duplicating that experience in that person would not be possible. Now, if you can't explain it to a person void of the senses of smell and taste, how are you going to write a computer program to duplicate that experience? In artificial intelligence? It's simply not possible.

Question:

Can computers understand?

Robert J Marks:

Computers don't understand what they're doing. They can add the number six and 48, but they don't understand what the number six is. They don't understand what the number 48 is. In a great example of the inability of computers to understand, John Searle many decades ago gave the example of Searle's Chinese room. Searle's Chinese room is where Searle is isolated in a room with a bunch of file cabinets. In these file cabinets are filed a number of questions and answers written in Chinese.

Now, Searle did not write or read Chinese, but he was in this room with these file cabinets and somebody would come and they would slip a slip of paper in with a question written in Chinese. Searle would take the question, he would look at it, and then he would go to all the file cabinets until he finally found a match. Once he found a match, he would take out the card, he would write down the answer on the card, and we've refiled the card, go over to the door and slip the answer out the door. Now, externally, it looks like that Searle in this Chinese room understands Chinese. He doesn't. He's simply following an algorithm. When we had IBM Watson beat the world champions in Jeopardy, this was a big Chinese room. IBM Watson had access to all of Wikipedia, everything on the web, and was able to answer the queries made to it, but it had no understanding. It doesn't understand why it responded that way. Computers have no understanding.

Question:

Why can't computers be creative?

Robert J Marks:

Computers will never be creative. This was something which was noticed by a number of people including the genius Sir Roger Penrose who wrote a great book, which convinced me, which was called The Emperor's New Mind. Before we talk about whether or not a computer can be creative, we have to define exactly what we mean by creativity. Creativity is best defined by Selmer Bringsjord's Lovelace test. According to the Lovelace test, computers are going to be creative if the computer program does something which is beyond the intent or the explanation of the computer programmer. It's that simple. There is no instance of which I am aware that a computer has been a computer program has been creative.

Now, we can go to as an example when Lee Sedol was beat by AlphaGo in the game of Go. There was one point in the game where AlphaGo made this surprising move and everybody says, "Oh, my goodness." It made a move, which nobody anticipated. Was this creativity? No, AlphaGo was trained to play Go. That was exactly what it was doing. If somebody came to AlphaGo and said, "Can you explain me the rules of Go?" and it was able to explain it without being programmed to do so, then that would be creative. That would be beyond the intent of the original programmers. Or if you went to the AlphaGo and said, "Do you play poker?" and it would either answer yes or no, that would be something that would be a modicum of creativity. So no, we see no evidence of creativity in computers.

Question:

Why can't computers experience love, compassion, or empathy?

Robert J Marks:

I did make the claim that artificial intelligence cannot experience love, compassion, and empathy and other human attributes of that sort. It doesn't mean that they can't mimic it. There was a great Steven Spielberg movie called A.I. where this little boy robot was trained to emulate love, and there's this great scene where the little boy robot is activated to love and the little actor that played it, I think his name

was Osmond, he all of a sudden just, you just saw him a wave of compassion and love. This was a simulation of love. The robot itself did not experience love. It was something that was being mimicked, and I think as David Lerner said, there's probably very few things that AI can't mimic.

Question:

Why do many robots have human-like qualities?

Robert J Marks:

In order for AI to work seamlessly in society, it might have to be to the point where it can relate to human beings. This is done using what I refer to as seductive optics. You will see, for example, a robot that has facial expressions, and these facial expressions will be manipulated to convey human emotion, but that has nothing to do with the internal artificial intelligence. It's just dressing upon the artificial intelligence to give it the appearance of human-like abilities.

Question:

In movies, we see AI overtaking humanity. Is there any truth to this idea?

Robert J Marks:

There's a bunch of great movies, but of course, I think it's important to differentiate the computer science of artificial intelligence and the scientific fact with the scientific fiction, or science fiction, if you will. One of them, for example, was The Terminator with Arnold Schwarzenegger. In the original movie, he talked about his neural network and that his neural network needed to be reset and all of this stuff and then there was the concept that Skynet had become conscious and had taken over, after it became conscious it would take over the world. That sort of thing isn't going to happen. It's not algorithmic. It's something which is not going to be computable and it's something which is not going to happen in the future.

Question:

Can you explain the limitations of AI?

Robert J Marks:

A major limitation of artificial intelligence can be summarized with a single word: algorithms. Everybody's heard of algorithms, but usually algorithms are not well understood. An algorithm is nothing more than a recipe or a procedure. It's a step-by-step procedure to do something. A recipe for chocolate cake is a recipe and it is an algorithm and what it does is you have the input and what is the input? The input is all the ingredients. That's the input into the algorithm and then you add the procedure. It's how you cook the batter. It's how you preheat the oven. It's how long you cook the cake, it's how you put the icing on, et cetera, et cetera. So that's an algorithm, it's a step-by-step procedure. If I give you directions to come to my house, that's going to be an algorithm because I'm going to say, "Go down I-35. Turn off Route 6. Go so many miles. Turn left, turn right," et cetera, so that is another example of an algorithm.

Well, it turns out computers can only perform algorithms. This is all that software code is. It's performing an algorithm, a step-by-step procedure towards some end. And the interesting thing is that there are things which are non-algorithmic. Most undergraduate computer scientists know about something called the Turing halting problem, which has been proven to be non-algorithmic. There are

things proven to be non-algorithmic. If something is non-algorithmic, it means that it is non-computable. You cannot write a computer program to do it. The halting problem says that no computer program can be written to analyze another arbitrary computer program to say whether that program will stop or whether it will loop forever. It's just not possible, and it's not a conjecture, it's not wishful thinking, it's a proven fact, so there are things that are non-algorithmic. It turns out that this non-algorithmic aspect translates to human abilities, so as we mentioned, creativity, sentience, and understanding, we're talking about those as non-algorithmic properties, that there's something which cannot be reduced to a computer program. They are something that is above and beyond the capabilities of computation.

Question:

What about cases where AI is said to create music?

Robert J Marks:

AI is often purported to perform creative acts such as writing music and you'll hear, for example, somebody saying, "I wrote artificial intelligence that creates music." Well, let's peel that back and look at where that music came from. Here's a very typical scenario. One wants to write artificial intelligence that composes music, so one will go to a specific genre, such as Bach, and it will train the artificial intelligence using all of the volumous works of Bach, and then once it's trained, it generates music. Guess what that music sounds like? Bach. That trained artificial intelligence will never generate the work of a Stravinsky. It'll never generate jazz. It'll never generate one of my favorite composers, Charles Ives. It's just not capable of doing it. Artificial intelligence is only able to work within the silo of what it was trained to do. It cannot think outside of the box and thinking outside of the box is mandatory for creativity.

Question:

AI can generate amazing paintings. Doesn't that show creativity?

Robert J Marks:

One thing that happened not too long ago was there was a painting titled Edmond Bellamy. Edmond Bellamy was a painting that was done by artificial intelligence trained on the portraits of masters. Rembrandt was specifically mentioned as one of the places that they went to get paintings to train this artificial intelligence. So we have a case. Guess what the paintings that were generated by AI looked like. It looked like the paintings of the masters, right? I suspect, I don't know for sure, but I suspect they did not use Picasso's cubism in any of the training. I suspect they didn't use Jackson Pollock's paint-splattering paintings in the training. They simply looked at these great masters of the past that did portraits, and so again, artificial intelligence is unable to think outside of the box and thinking outside of the box is a characteristic of creativity. It could only take what is inside the box, what it was trained to do, and interpolate.

The idea of creativity, in fact, takes the status quo inside the box and it discards it. It discards the status quo for some new innovative ideas. You only have to go back to Einstein, for example. What did he do? He discarded the dogma of ether in space because they had no idea that electromagnetic waves could travel in a vacuum. He discarded the idea that light was a relative in speed and this was the accepted dogma at the time. This was inside the silo. He discarded that and he took these ideas motivated by a number of experiments, including the Michelson-Morley experiment, and from this generated the wonderfully creative idea of relativity, so that is an example of creativity. Creativity normally requires

some sort of avoidance of things within the silo that are to be done external to the silo in order to see creativity.

Question:

Can you comment on some of the quotes from thinkers who fear AI?

Robert J Marks:

There are many people who are really worried about the future of artificial intelligence and what will happen. Bill Gates said, and I have a quote here, "I don't understand why some people are not concerned about AI." Yuval Noah Harari wrote a book called Homo Deus which paints a very dystopian view of the future about the terrible things that artificial intelligence is going to do. He said, "Currently, humans risk becoming similar to domesticated animals because of AI." This is the idea that AI is going to write better AI, going to write better AI, and eventually, it's going to become master of the universe and it's going to become our masters and we're going to become like the pets of artificial intelligence. Even geniuses like Stephen Hawking says, "The development of full artificial intelligence could spell the end of the human race. It would take off on its own and redesign itself at an ever-increasing rate. Humans who are limited by slow biological evolution couldn't compete and would be superseded."

In fact, this is a theme which is echoed by Ray Kurzweil is that us writing AI that writes better AI is the next step in human evolution, that we have reached the pinnacle of how we can be evolved, so therefore, we have to go to a different venue in order to keep our evolution going, and that's going to be with artificial intelligence. Well, you'll notice that Stephen Hawking's claim assumes AI can write better AI and that's not possible because that assumes that AI can be creative and as we have defined creativity by the Lovelace test, artificial intelligence cannot be creative.

The irony in Hawking's statement is that he was a believer in something called Gödel's theorem, which was the genesis of Alan Turing's later work into things which are non-algorithmic and therefore non-computable. Creativity is non-com computable. Gödel said that there were things which were outside the scope of possibility given a finite set of axioms or assumptions, so Hawking said that, "Well, you know this theory of everything, which is a physical law, which encompasses the entire universe is not going to be possible because of Gödel's theorem," so you recognize that, but he doesn't take it the next step to Turing's work on algorithmic versus non-algorithmic and it showed that computers can only perform algorithmic properties and humans can do non-algorithmic or non-computable properties such as creativity, so Hawking was very close and he disagreed with his colleague Penrose on the result of that and I think Penrose in his book, in Emperor's New Mind, makes an incredible powerful case against computers ever being creative.

Question:

Can AI ever have common sense?

Robert J Marks:

There's a lot of research going on right now in trying to give artificial intelligence common sense. In fact, I think I heard it called common sense is the dark matter of artificial intelligence currently. There is a great story about common sense about Fred Flintstone when he got his fingers glued in the bowling ball and he couldn't get his fingers out, so he told Barney to go get a hammer, and so Barney came back with a hammer and Fred said, "Okay, Barney, when I nod my head, hit it." Now, that's funny because you think about how it could be interpreted, right, "When I nod my head hit it." There was a vague pronoun there that could have referred to the bowling ball or to Fred's head. Now, when we hear this, we know

exactly what Fred meant when he said this statement and we look at the incorrect observation as humorous.

Computers don't have that ability. They don't have that ability to assign meaning to vague pronouns. Some of my favorite examples are flubbed headlines where a headline can be interpreted in two ways and one of them is funny because we know it's wrong. But the other one is conveys a truth. I'll give you a few examples. Iraqi Head Seeks Arms. This was an example of a headline, Iraqi Head Seeks Arms, so the humorous part is we envision a big head that needs some arms to work around. Of course, that's not what it meant, and we know that that's not what it meant. The Iraqi head means the head of the Iraqi government needed military arsenal. Another one is Seven Foot Doctors Sue Hospital. Is it 7' doctors sue hospital or seven foot doctors sue hospital? Did a lot of tall doctors sue the hospital or did a bunch of podiatrists sue the hospital? We know immediately that it was the podiatrist, yet a computer would have problems in doing that.

There is a contest which is held yearly, which I learned from my colleague Gary Smith at Pomona called the Winograd schema. The Winograd schema is a contest to see if computers and AI can have common sense and the Winograd schema is something with typically a vague pronoun. I'll give you an example I got from Gary. It says, "I can't cut down that tree with this ax. It's too small. I can't cut down this tree with this ax because it's too small." Immediately in context, we know it's the ax that's too small. It isn't the tree that's too small, it's the ax that's too small, and so the question is, can a computer differentiate the meaning of the vague pronoun "it's" in "it's too small"? Currently, the last I heard is that the Winograd schema contest has come out with only roughly about 50% success, which is just like flipping a coin. It's a little bit above 50%, I believe, so currently computers don't have common sense. I'm not sure whether this is algorithmic or non-algorithmic. It might be something which is solved in the future, but right now it's a big problem in artificial intelligence to give the artificial intelligence common sense.

Question:

AI can't do all this stuff now, but what about computers in the future?

Robert J Marks:

Well, it turns out that there are certain things which computers today and computers in the future will never be able to do and this is constrained by something called the Church-Turing thesis. The basic idea of the Church-Turing thesis is that computers today, when they perform an operation, they execute a computer program. That same computer program could be performed by Turing's original clunky machine. Now, Turing's original clunky machine could take thousands, maybe millions of years to do it, but it still was operationally possible, and that's going to be true of computers in the future, except that they're going to be faster and more powerful, more memory, greater recall.

So the Church-Turing thesis says if there are fundamental algorithmic information theoretical limitations of computers today, they will also operate on computers of the future, so if a computer, if you have an operation which is non-algorithmic and therefore non-com computable, then it's non-computable. That's a fact. It can't be done by Turing's machine, it can't be done by the computers of today, and it won't be able to be done by the computers of tomorrow because the computers of tomorrow will still be limited to perform algorithmic or computable operations.

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 Egbert. The opinions expressed on this program are solely