Bingecast: Robert J. Marks on the Limitations of Artificial Intelligence (https://mindmatters.ai/podcast/ep97)

Austin Egbert:

Greetings. I'm Austin Egbert, Director of Mind Matters News. Welcome to another bingecast, where multiple episodes are fused into a single program. This bingecast features an interview by Larry L. Linenschmidt of The Hill Country Institute with our host, Robert J. Marks, about the astonishing limitations of artificial intelligence from a computer science perspective. This interview was originally aired by The Hill Country Institute at their website, hillcountryinstitute.org, and is included here with permission in its entirety. Enjoy.

Larry L. Linenschmidt:

Welcome to Hill Country Institute Live, the program that brings you together with Christian leaders, authors, and pastors to discuss major issues of our time. We seek to encourage and equip followers of Jesus Christ to know His heart and mind in all that we do. I'm Larry Linenschmidt, your host, and I thank you for joining us today. We have a special and unusual program today with a Christian leader who works in academia, who works with the artificial intelligence, neural networks, computers, and all of those things that those of us who aren't very technical wonder, "How do they work? How should a Christian think about artificial intelligence?" Stay with us. You'll want to hear the story of how Dr. Robert Marks, Professor at Baylor University, works with artificial intelligence and all of those computer capabilities that seem so complicated to us non-technical folks.

Larry L. Linenschmidt:

We invite you to visit our website, hillcountryinstitute.org, to hear podcasts of our past programs on topics including faith and work, science and art, fighting human trafficking, and many other topics. Our guests have included Os Guinness and Dan Davis and Andy Crouch, Mayor Ivy Taylor of San Antonio, and lots of other folks who are on the cutting edge of what Christians are doing today to interact with our world. The website also offers audio and video from our past conferences and seminars on faith and science and faith and art and C.S. Lewis and much more. Our radio programs are available on any podcast service as well as our website, hillcountryinstitute.org. On the podcast, they're Hill Country Institute Live.

Larry L. Linenschmidt:

The program's supported by donations. We ask you to consider donating or supporting the program in some way. For donations over $100, we have a copy of Jay Richards' book on the topic that we're talking about today, The Human Advantage: The Future of American Work in the Age of Smart Machines. Jay Richards is a scholar and previous guest on our program. To donate to Hill Country Institute, please visit our website, hillcountryinstitute.org, or call 512-680-7993, hillcountryinstitute.org, 512-680-7993. Now, let's welcome our special guest, Bob Marks.

Larry L. Linenschmidt:

Now, I'd like to introduce our special guest today, Dr. Robert J. Marks II. He's Distinguished Professor of Engineering in the Department of Electrical and Computer Engineering at Baylor University. He was

attracted to Baylor after 26 years at the University of Washington in Seattle by Baylor's vision of being the home of cutting-edge research while celebrating the Lordship of Christ. While he was at the University of Washington, for 16 years he was the Campus Crusade Faculty Advisor and he's still involved in working with student ministry through advising Ratio Christi and the American Scientific Affiliation Student Chapter at Baylor University. He was the first President of The Institute of Electrical and Electronic Engineers Neural Network Society. He's a Fellow of the IEEE, The Institute of Electric and Electronic Engineers, and he's Program Director of The Walter Bradley Center for Natural and Artificial Intelligence. He's authored, co-authored, or edited over 400 publications, and if I continue all of the things on my list to introduce him, we won't have a program, so let's just stop right there and say Bob, thank you for being with us today. We appreciate you're being introduced

Robert J. Marks:

Thank you very much. You mentioned 400 publications. That's kind of like saying somebody has pounded 400 nails in a piece of wood, so I should probably say that some of those publication are pretty good.

Larry L. Linenschmidt:

All right. I think so, and-

Robert J. Marks:

Okay.

Larry L. Linenschmidt:

If you keep invited back to peer-reviewed publications, then they have to be pretty good, don't they?

Robert J. Marks:

Well, not really. That's a topic for another day.

Larry L. Linenschmidt:

Okay, that's a good discussion. My friend Jim Garvin, who teaches at Baylor, and I talk about peer review and the process quite a bit, so I have some insight for it. I want you know we have great appreciation for Christians who are in science and technology and in your work in the area of information and technology and artificial intelligence is really on the cutting edge. We'd like for our audience today to get to know you a little bit personally, and then we want to get into the science of computers and robots and artificial intelligence or AI.

Larry L. Linenschmidt:

If you're listening and you hear AI, it's not just like Old MacDonald had a farm and I'm going back to the refrain, it's artificial intelligence. Of course, that's probably a silly comment because AI is on everybody's mind today. Bob, what attracted you to science and electrical engineering as a young man?

Robert J. Marks:

Well, I think that everybody, according to Scripture, has their gifts and God gifted me to be a nerd. I just like math, I like science. I love things. I can look at the latest gadget or the latest great invention and I can get excited and I can see God in that. My wife, on the other hand, she looks at it and yawns. She

sees beauty in a field of flowers that are beautiful. She says, "Oh, I see God there." I think everybody has their own predisposition as to what they're good at and what they're gifted at. Again, I was gifted to be a nerd and I appreciate it.

Larry L. Linenschmidt:

Well, and in this sense, you are what Tolkien refers to as a subcreator. You're looking at these things and how they can be put together and you're making something out of what God made to start with for us.

Robert J. Marks:

Yeah. In fact, I had a grandfather who had a third-grade education. He once told me, he says, "Nothing man has ever done that God hasn't done before in some sense," and I think he's probably right. We're just replicating what we see in nature. In fact, one of my societies here, IEEE Computational Society, used to have a logo which said that... not a logo but a saying that said that, "We look to nature to find out the new things that we can do." In fact, neural networks, if you talk about artificial neural networks, that's the attempt to replicate the human brain, so we're looking at a creation and then we're trying to borrow from that in an engineering sort of application sense, so yeah.

Larry L. Linenschmidt:

Sure. Well, let me ask you back to thinking as a young person, do you recommend science and technology as areas of study for young Christians?

Robert J. Marks:

Only if you're gifted in it. I think that there's this effort now for STEM... STEM stands for what, science, technology, engineering, and math? It's kind of like the idea that one-size-fits-all, but not everybody is supposed to go into the STEM fields. If you're not gifted in mathematics and a love for science, you shouldn't go in there. You should find out whatever area it is that you're interested in, go in there and flourish in the way that you've been gifted.

Larry L. Linenschmidt:

If you are gifted in those areas and you feel like God has particularly given you an interest like He did you, then is it a good calling? Is it a valid calling?

Robert J. Marks:

Oh, it's an incredible valid calling. I think that you don't have to go into ministry to actually be a good Christian or to follow Christ. I think you can do that no matter what area you go into. In fact, I feel it's a blessing that I'm doing what I'm doing now and can still serve The Lord. What a blessing to do what you love to do. I think Solomon said in Ecclesiastes that, "The only that is really good under the sun is the good feeling that you have after a good solid day's work." That's exactly the way that I feel and just feel blessed in life that I can do what I've been called to do.

Larry L. Linenschmidt:

Well, God wired us to work. We worked before the fall, and that's really important theologically because our work is fulfilling and good and honoring to Him, isn't it?

Robert J. Marks:

Absolutely, yes.

Larry L. Linenschmidt:

Well, you've been in an academic environment for pretty much all of your adult life. You were in secular universities as a student and as a professor, and now you've been at Baylor for many years. What difficulties, what opportunities, what's it like to be a Christian in a secular university setting today?

Robert J. Marks:

Well, I've been at Baylor about a decade and a half. I was with the University of Washington for 25 years. Looking at me, you would never think that, of course, I don't look that old, but nevertheless, sometimes it's difficult. I think that our travels as people who believe are not always easy. We hear of things today in many of the secular universities. We don't experience it here at Baylor because Baylor is, according to their publicity unapologetically Christian, so we don't hear it here, but at other places, you hear of parachurches, such as InterVarsity and Campus Crusade and such being kicked off of campus because the university prohibits them from requiring the leadership of the organization from being followers of Christ.

Robert J. Marks:

You can go in with any worldview or anything that you want to according to the university, and for that reason, because these parachurch organizations like Crusade, InterVarsity, and Ratio Christi don't do that. They are kicked off of campus and so that's very difficult. You might think that's new, but we had this problem back at the University of Washington where we were challenged for the Campus Crusade. Today, it's called Cru. We were challenged that we weren't allowing leadership positions for people that had opposite ideologies of what Cru stood for.

Larry L. Linenschmidt:

Yep. Well, that's going to be a constant struggle for us, isn't it? I don't know if we're a post-Christian culture or a pre-Christian culture at this point, but we need to continue to associate with each other, support each other, and share the Gospel as we can, don't we?

Robert J. Marks:

Yes, absolutely.

Larry L. Linenschmidt:

Well-

Robert J. Marks:

Yes, sir.

Larry L. Linenschmidt:

One of your interests is the Walter Bradley Center for Natural and Artificial Intelligence. Walter Bradley's a Fellow of The Hill Country Institute. He's been an important friend to our ministry and our faith and science work and supportive of just about every kind of work that we've done, from C.S. Lewis to how we relate as Christians to science. It's great that this center's named after him, and would you tell us, what is The Walter Bradley Center for Natural and Artificial Intelligence?

Robert J. Marks:

Well, The Walter Bradley Center for Natural and Artificial Intelligence has a website that I'll publicize it here. It's mindmatters.ai, and I heard politicians say if you mention a website you should mention it three times for the maximal effect, so let me say again, it's mindmatters.ai, mindmatters.ai. It was founded in July of last year and the purpose of this is to look at artificial intelligence from a perspective that you really don't see in the media. We try to be totally honest, we try to be scientific. We do embrace the Judeo-Christian ethic, but don't refer to Scripture. We're actually talking about the actual science and what goes on.

Robert J. Marks:

The Walter Bradley Center for Natural and Artificial Intelligence is looking at artificial intelligence, the hype that we see in the headlines today, which is pervasive, and trying to smooth some of that out and trying to bring a better perspective into exactly what AI is. We're also interested, and this is in the history of Walter Bradley, interested in establishing sustainable and appropriate technology in majority worlds. Now, many people haven't heard of majority worlds, and I didn't either until about a month ago, but, I guess, developing countries is derogatory and Third World is also derogatory, so majority countries is what comes to mind.

Robert J. Marks:

Walter Bradley went to different places and he took appropriate technology to the countries. Most countries don't need better computers, they need technology which is appropriate for them, so Walter Bradley did this especially with the chemistry of coconuts and actually has founded some sustainable technologies and sustainable businesses in different countries. In the Philippines now there's a company called Dignity Coconuts where they have taken the incredible number of coconuts which they have available and they have done stuff with them. And now you can buy Dignity Coconuts online at amazon.com where you can buy everything else in the world. What they have done is they have established now a business in the Philippines that is run by nationals that has helped with the poverty.

Robert J. Marks:

We have another example, which is we have just funded a gentleman named Brian Thomas who is going down to Haiti. Now, Brian is going around and started little micro businesses around. Everybody in the developing, I'm sorry, everybody in majority world countries has cellphones and they all need to charge them, so Brian went around and he set up people with solar panels that could charge people to come in and charge their cellphones.

Robert J. Marks:

This is a nice, sustainable technology that Brian is doing in Haiti and we're looking at expanding this in other places. The Walter Bradley Center is also into looking at how we can impact majority world countries with advanced technology, including AI. My area specifically is in the theoretical. I actually specialize in something, and it's a lot better than it sounds, it's a lot more fun than it sounds, but algorithmic information theory, which will actually tell us in the mathematics of computer languages what a computer can do and what a computer can't do. That's my area of interest and expertise, but the spectrum of The Walter Bradley Center is much broader than that as you can see.

Larry L. Linenschmidt:

Yes. Well, I know Walter's worked on wells, bringing clean water to so many people around the world, so there are many aspects to what that can do. You've kind of led us into computers, and let's think about, if we can, kind of lay the stage for what computers are, how they work. The cards that I had for programs back when I was at UT and the Business School, that's a long time ago and if they got out of order I was in big trouble. Today, we've...

Robert J. Marks:

I worked with those punch cards. It was... Hey, Larry, that really ages us, doesn't it?

Larry L. Linenschmidt:

It does, I'm afraid.

Robert J. Marks:

That was a long time ago. Kids today don't even know what those punch cards are.

Larry L. Linenschmidt:

Yeah. Well, let's just start there. When I read the term "classical computer," how does a computer function? Let's build on that to talk about supercomputers and kind of build into just a foundation of how these things work so we can then talk about the theory of AI and what it is and what it isn't.

Robert J. Marks:

Well, computers as we know them today were actually popularized by Alan Turing. You might recognize his name. He had a biography in the movie The Imitation Game, where Alan Turing was played by Benedict Cumberbatch, which is a fun word to say. Alan Turing was really a genius and he was the one that actually invented the idea of computers. He used a type of computer to break the Enigma code for the Nazis back in World War II. That was the subject of the movie The Imitation Game. Basically, Turing showed the idea that computers were limited by something called algorithms, and we hear algorithms a lot. "Such and such is doing an algorithm and Facebook has initiated an algorithm to do something." The question is, what is an algorithm?

Robert J. Marks:

The algorithm is simply a step-by-step procedure to accomplish something. If you go to your shampoo bottle and you look at the back and it says, "Wet hair, apply shampoo, rinse, and then repeat." That's an algorithm because it tells you the step-by-step procedures that you need to wash your hair.

Larry L. Linenschmidt:

Well, that's a pretty short algorithm for me since I don't have much hair, but go right ahead.

Robert J. Marks:

Isn't that right? Well, the interesting thing about that algorithm is if you gave that to a computer, that computer would wash its hair forever because it doesn't say repeat once, it just says repeat-

Larry L. Linenschmidt:

Oh.

Robert J. Marks:

So you would go back and you would do it again and again and again. A proper shampoo, make sure your shampoo says repeat once so you don't spend all of your time in the shower shampooing your hair. An algorithm I like to think of as a recipe. If you look at the recipe for baking a vanilla coconut cake, for example, it will tell you the ingredients that you need and then it will give you a step-by-step procedure for doing it. That is what an algorithm is and, in fact, it is what computers are limited to do. Computers are only able to perform algorithms.

Robert J. Marks:

One of the things that we can identify that humans can do that computers can't do are things which are non-algorithmic. If it's non-algorithmic, it means it's non-computable. Actually, Alan Turing showed back in his initial work that there were things which were not algorithmic. It's very difficult, for example, to write a computer program to analyze another computer program. Turing showed that specific instantiations of that were non-algorithmic. This is something which is taught to a freshman computer science students, so they know what algorithmic and non-algorithmic/non-computable is. Again, non-computable is a synonym for non-algorithmic.

Robert J. Marks:

We have a number of aspects that we exhibit that are non-algorithmic. I would say qualia, creativity, sentience, consciousness are probably things that you cannot write a computer program to simulate. That's kind of the background and I think I went down a rabbit trail here a little bit, but that's-

Larry L. Linenschmidt:

No, no.

Robert J. Marks:

That's-

Larry L. Linenschmidt:

That's-

Robert J. Marks:

Kind of-

Larry L. Linenschmidt:

We're going to [crosstalk 00:20:06] tie into that some more. I want to come back-

Robert J. Marks:

Okay.

Larry L. Linenschmidt:

To what computers can and can't do and the uniqueness of the human, but in building computers, I have a cellphone that I understand has more power than a room full of computers 50 years ago that Army intelligence used. It's a massive increase in computing capability, isn't there?

Robert J. Marks:

Yes there is, but what we've done is we've been able to by increasing the speed and using parallel sort of computers, we have just increased the speed of the computers. There is a principle taught to computer scientists called the Church-Turing Thesis, which basically says that the computers of today can do... they're remarkable, but Alan Turing's original machine could also do what the computers today do. The only thing that computers could do today is do things a lot faster. Therefore, that is really good, that is very useful. We can do some things that we couldn't do before, but in terms of the ability of the computer, they are still restricted to algorithms. I'm not sure if you've ever heard of the quantum computer-

Larry L. Linenschmidt:

Yes.

Robert J. Marks:

Which is kind of the new rage where you use this strange, weird world of quantum physics in order to get computational results. Even quantum computing is algorithmic and is constrained by the Church-Turing Thesis, so we're doing things a lot faster now. With quantum computers, we're going to be doing them like lightning, but still, all of the stuff we could do we could do with Turing's original machine. Now, with Turing's original machine, it might take us a trillion years in order to do it compared to today, but nevertheless, the capability is with Turing's original machine. We're just getting faster and faster and we can do more interesting things because of that speed.

Larry L. Linenschmidt:

Well, when I visited places like The National Renewable Energy Lab and they have literally a large room full of the computing capability they have, it's a large structure and it's up to date in terms of speed, and what you're saying is all of that has really done from Turing and World War II and his great work on the Enigma code is increased the speed and increased maybe the complexity, but I'm not even sure that's the right word. It's just bigger and faster.

Robert J. Marks:

Well, no. If you... Yeah. If you have greater speeds, you can entertain more complexity, absolutely.

Larry L. Linenschmidt:

Okay, but that's what's happening. It's still the same basic algorithm built on algorithm and that's the limitation of the computer, that it can't do non-algorithmic functions?

Robert J. Marks:

That's correct, yes.

Larry L. Linenschmidt:

Okay, and then so another term that if I could just talk about a second so we can go further with it, what is artificial intelligence? How do you define that?

Robert J. Marks:

Well, what is artificial intelligence? It's interesting today. I think that artificial intelligence today means any gee whiz things that come out of high tech and computers and gee whiz sort of thing that you get is called artificial intelligence. It used to be not like that. Artificial intelligence, as founded by Marvin Minsky and his lot at MIT, would correspond to trying to take the expertise of a human and having the computer program converse with the human and write down rules, and from those rules that you write in the computer, you have an intelligence which captures the human intelligence.

Robert J. Marks:

More modern artificial intelligence such as neural networks and reinforcement learning rather learns by example. That is, they are presented with a number of examples and after while they go, "Yeah, I understand the difference between a cow, a pig, a horse, and a snake," because you've shown them this a number of times and it begins to recognize it. In that sense, it adapts to the data set that you give it.

Larry L. Linenschmidt:

Okay.

Robert J. Marks:

Defining artificial intelligence is a difficult thing to do because, again, I think of the media, specifically, it's any gee whiz thing that comes out that involves computers.

Larry L. Linenschmidt:

Yeah. If it's something that sounds far out, sci-fi-ish, then it's AI is kind of what you're saying?

Robert J. Marks:

Yes, and unfortunately it's become kind of a buzz hype word that people talk about. There's a lot of hype out there about artificial intelligence and its capabilities.

Larry L. Linenschmidt:

Well, what is a neural network? How does that relate to AI?

Robert J. Marks:

Neural network is an attempt to take the human brain, or any brain, the mosquito brain, I don't care, any brain that you want to, and make computer code that follows the way that the human brain is modeled. We hear a lot about neural networks today, but shoot, neural networks go back to the 1950s. There was actually a professor at Stanford named Bernie Widrow that came up with a neural network in the 1960s that can predict the weather, it could listen to your utterances and have it typed out on a computer. Back then, they didn't have computer screens. They printed it out on kind of a typewriter. It was able to balance brooms with it and it was really, really exciting stuff.

Robert J. Marks:

There was also a gentleman named Frank Rosenblatt, I believe he was with Cornell, that developed something called the Perceptron and everybody got excited. Then, an AI winter came in and we started in... I guess it was kind of the 1990s there was a resurgence of interest in neural networks because people came up with a way to overcome some of the limitations that the early, early neural networks did. Today, there's a resurgence again. There has been the development of so-called "deep learning"

specifically deep convolutional neural networks and deep reinforcement learning that has been used, for example, win Go Games against the world champions. These style, I think, are the primary focuses of artificial intelligence and we hear a lot about these sorts of things.

Robert J. Marks:

The basic idea behind a neural network is that you have a bunch of neurons and these neurons in the computer are just a bunch of kind of dots. All of these neurons are kind of connected together with something called synaptic weights. In order to learn the difference between a horse, a cow, a pig, and a snake, you give the neural network the experience of seeing a bunch of these things. With supervised learning, what you do is you say, "Okay, here you go. This is a snake. This is a pig. This is a horse. This is a cow." I might be using examples that were not on my list, but you get the idea.

Larry L. Linenschmidt:

Sure.

Robert J. Marks:

Then, pretty soon after repeated trials, the neural network begins to say, "Ooh, okay, I've seen this before." At least that's what it seems to do. All that it's done is adapt its weights in such a fashion that it does exactly what the programmer told it to do. It says, "I want to look at these different categories of animals and I'm going to tell you what these animals are and you're going to do something. You're going to rearrange your synaptic weights." That's what training a neural network is. It's actually updating these weights, making them bigger or smaller. "I'm going to give you an algorithm to adapt these synaptic weights so that the next time you see a picture of a cow, you'll say that that's a cow."

Robert J. Marks:

That's what a neural network does, and the interesting thing about the neural network is that it doesn't rely on this idea of getting rules from people. Rather, it's something which is presented with examples. It used to be you would have to give features, which I won't go into, but now you can actually give them pictures and actually look at the pixels and that is pretty exciting stuff. That's what deep convolutional neural networks are able to do. That's what a neural network does. It's the attempt to capture the way that the human brain works. In terms of the algorithmic aspects of learning and adaptation, that's pretty successful and it's pretty useful.

Larry L. Linenschmidt:

One of the things we talked about earlier were algorithms and what computers can do and some things that maybe they can't do. What are the things that maybe computers will never be able to do?

Robert J. Marks:

Well, I think maybe the biggest testable thing that computers will never be able to do is creativity. Computers, they can only take the data which they've been presented and interpolate. They can't, if you will, think outside of the box. If you look at the history of creativity, like great scientists like the Galileo and Einstein and such, they actually had to take the data that they were given. They had to discard it and they came up with something which was brand new. It wasn't just a reshuffling of the status quo, which is basically what a computer can do, it was actually a creative act outside of the available data.

Robert J. Marks:

I think another thing a computer can't do, I have a whole list of them, I think another one is qualia, which is a word I just learned. Qualia is kind of the subjective experience that one has in themselves. Imagine, for example, having a big, delicious red apple and you anticipate taking the bite out of it. You take the bite, you feel the crispness, you feel the tart sweetness, you feel the crunch as you chew it and swallow it. That is an experience and the question is, do you think you could ever write an algorithm to explain that qualia experience to a computer? I don't think so. I think that that is something which is unique to the human being.

Robert J. Marks:

That creativity, qualia, understanding I think that you mentioned that Jay talked about in his book, I'm not sure about this interview, but he talked about Searle's Chinese room. John Searle was a philosopher and he said that, "There is no way that a computer understands anything." He illustrated with the Chinese room, but the basic idea was you slipped a little slip of paper in which there was something written in Chinese through a little slot. Inside the room, somebody picked it up and they looked at it and they wanted to translate it to something, say, like Portuguese.

Robert J. Marks:

There's a big bunch of file cabinets in the room, and maybe people still know what file cabinets are, but what the person in the room did is took this little sheet that had this Chinese on it and he did a pattern matching. He went through it, he looked through all of the file cabinets and he finally found something that matched the little sheet of paper that he had. With that little sheet of paper was the translation into Portuguese. He took the little translation in Portuguese, he wrote it down, he refiled the original things. Went to the door and slipped out the translation into the Portuguese.

Robert J. Marks:

Now, externally, the person would say, "My gosh, this guy knows Chinese, he knows Portuguese. This computer is really, really smart." Internally, the guy that was actually going through the file cabinets, doing the pattern matching in order to find out what the translation was, had no idea what Chinese was, had no idea what Portuguese was. He was just following a bunch of instructions. That is basically the concept of understanding. Let's see.

Larry L. Linenschmidt:

The computer processes, it turns out work product based on how it's directed, but in terms of understanding, as we think of understanding like you would expect one of your students to understand what you're teaching, they don't understand. They compute. They process data. Is that a fair way of putting it?

Robert J. Marks:

Absolutely. There was a great columnist about 20 years ago that commented on Deep Blue beating Kasparov, who was the world champion at chess at the time and Deep Blue was trained to play chess and Gelernter wrote a great column. He said, "Look, the idea that Deep Blue has a mind is just totally absurd. He feels nothing, knows nothing, and he plays chess for the same reason that toaster toasts." That was the way he was programmed to do it. He isn't happy he wins, he isn't sad he loses. If he wins, he doesn't," I love this phrase. He says, "If he wins, he doesn't plan a night on the town with Pink," obviously being his girlfriend. No, there's no adulation. There's no mind there. It's just like Searle's Chinese room. You also might be familiar with IBM Watson beating-

Larry L. Linenschmidt:

Yes.

Robert J. Marks:

The world champions at Jeopardy. If you think about it, that's just a big Chinese room, except that in this Chinese room, you have all of Wikipedia and all of the internet available to you and you're given some sort of question on Jeopardy and you have to get the answer. You look around. There's a pattern matching and then you bring back the answer to the question. The Watson beating the world champions in Jeopardy is exactly an example of a Chinese room, except the room is a lot bigger because computers are a lot faster and can do a lot better things.

Larry L. Linenschmidt:

Sure, and Watson had a built-in advantage, then, by having... It might not be infinite, but it would seem virtually infinite information available to it to do those matches.

Robert J. Marks:

Yes, and, oh, I learned this. I read a book by, and I would recommend it, it's called The AI Delusion by Roger Smith. He pointed out, and I never knew this before, that the people that did IBM Watson were a little bit concerned because sometimes when you present things to the computer, there's ambiguity. I'll go back now to Fred Flintstone to illustrate that. There was a Fred Flintstone cartoon where he got his fingers glued inside a bowling ball. He told Barney Rubble, he said, "Okay, we got to get this off."

Robert J. Marks:

They tried pulling it and everything. Barney got a big hammer and Fred said, "When I nod my head, hit it." The idea there was, of course... Of course, the idea was that the vague pronoun "it", and so when you present things to a computer that don't have context, like when I nod my head "hit it," by the way, Barney did hit Fred in the head, but when you present things like that to a computer, they don't know how to respond, as Roger Smith pointed out in his really, really great book.

Robert J. Marks:

One if the things that the Watson programmers did, according to Roger Smith, is they said, "Look, we don't want any questions asked of the Jeopardy contest that are confusing like this." The Jeopardy people says, "Yeah, but you don't want to fix the game by removing questions like that," and so they actually arrived at the compromise that they would actually go back and look at old questions that were from Jeopardy programs that hadn't been asked yet. This wouldn't be a topic when these people put together the questions for the contest between Watson and the other participants. Even there, we see the inability of computers to do things that humans are able to do, at least in the case of Watson for that example.

Larry L. Linenschmidt:

Well, there's one other game example that comes up quite a bit in the literature, and that's the game Go, and apparently Go is the most complicated game and a computer did very well. Is that just an extension of the same idea that it was able to match possible outcomes and evaluate the best of those? Or what? How do you [crosstalk 00:36:45] look at that?

Robert J. Marks:

Yeah. Go was a remarkable computer achievement. I don't want to degrade this at all. They used the concept called reinforcement learning and this reinforcement learning was used in the chess, it was used in Go. It was actually used to win the old arcade games that just by looking at the pixels in an arcade game such as, oh, I don't know, Pac-Man, for example, that the computer could learn how to win. Now, in all of these cases, of course, there was the concept of the rules. You got to know the rules and the fact that Go was mastered by the program is an incredible accomplishment of computer science. However, notice that the computer is doing exactly what it was programmed to do. It was programmed to play Go, and so very narrow application of artificial intelligence.

Robert J. Marks:

I would be impressed if the computer program would pass something called the Lovelace Test, which is the test that computer programs are given for to test their creativity. The Lovelace Test basically says that you have seen creativity if the computer program does something that can't be explained by the programmers. Now, you might get some surprising results. There was some surprising results that Alpha Go used when it played the master, but surprising doesn't count. It's still in the game of Go. If Go had gone on to do something like... Let me make the point by exaggeration. If the Go program went on to give you investment advice or to forecast the weather without additional programming, that would be an example of AI creativity.

Robert J. Marks:

Now, there's some that think that so-called "general AI" in the future is going to accomplish that, but we're just kind of waiting and seeing. We see a lot in AI, the so-called "algorithm of the gaps." We say, "Yeah, we can do this now, but someday we will have an algorithm that can do all of this other stuff." This algorithm of the gaps is used to substantiate the future of artificial intelligence and not really realizing that this algorithm of the gaps, that if it actually is used to augment what computer programs do will actually be the result of human creativity. Algorithms in computers are the result of human creativity.

Robert J. Marks:

That is not a controversial viewpoint. Nadella, who is, I believe his name is Nadella, he is the current CEO of Microsoft. He says the same thing. He says that, "Look, computers are never going to be creative. Creativity will always be a domain of the programmer." That's probably more than you wanted to know, but that's my explanation of Alpha Go and what that is doing.

Larry L. Linenschmidt:

No, I think that this is absolutely fascinating and I don't think there's anything that could be more contemporary, more current for us to deal with because AI comes up in and literally I've been in exchanges, particularly on Facebook, with non-Christians and they just say AI is going to wipe out your faith. I'm delighted to think about what AI is, what AI isn't, and these limitations.

Larry L. Linenschmidt:

In thinking about art particularly, in The Guardian yesterday there was an article about art that's being created using AI and there are apparently two networks and they interact with each other and they develop a lot of images, thousands of images, and then the person working with it picks the ones that they like. They've sold some for thousands of dollars at some of the auction houses. While I guess this

has the appearance of "creativity", but yet if I'm understanding correctly what you're saying, these programs, these networks have been designed to make art, to make pixels together. Now, they may not understand what they're doing, but if you teach them more about what you want, their program will grow in some fashion. Is that fair? Is that a good way to put it?

Robert J. Marks:

Well, it will grow as long as it's guided by the human programmer.

Larry L. Linenschmidt:

Okay.

Robert J. Marks:

You noticed that the human programmer is the one that controls what's happening. Let me comment on the ability of computers to generate art. Let me begin with the illustration of music. Let me tell you a typical scenario of training like a neural network to actually do a composition. Usually a neural network is presented a genre of different types of musics. Say, for example, the works of Bach. Then, it looks at all of that and you ask it to generate some music, and guess what that music sounds like? It sounds like Bach.

Larry L. Linenschmidt:

Surprise.

Robert J. Marks:

Surprise, exactly. It's the same thing with these... Well, let me address the creativity before I go into the art. It creates music that sounds like Bach, so that program will never be creative. It will never create music which is written by a Wagner or a Schoenberg or some of my favorite people, Schnittke or some of the modern composers like Charles Ives. It won't. It is constrained, if you will, to think inside that box. It doesn't throw away old conventions to do something new and creative. It is constrained to be that way.

Robert J. Marks:

I would maintain, although I don't know the details of the art except what I've read in the literature, that the same thing is happening in this creation of art. It is being fed a number of pictures from a number from genres, and guess what it generates? It generates something that is part of that genre. It isn't creative. It doesn't go from the work of a da Vinci to a Picasso, which certainly involved creativity, so the AI will never be able to be trained in works of paintings of a da Vinci or a Michelangelo and then spit out works that are done by a Picasso because that is just too big of a jump. That's creativity, and creativity is something which computers will never, never accomplish to a high degree.

Larry L. Linenschmidt:

As you're explaining it, I'm thinking that a computer is as good as its programmer, it's good at matching, it's good at putting things together, but true creativity, what the entrepreneur Peter Thiel refers to... A lot of people can take us from one to infinite, but it's that zero to one that is creativity in the tech world, in the business world that sets us apart. That verbiage, then, a computer can't take us from zero to one. It needs instructions, doesn't it?

Robert J. Marks:

It does, and in fact, Thiel's book, I think it's called From Zero to One, he talks about the requirement of creativity. His philosophy is parallel to that of some other people, Jay Richards, for example, who you've interviewed, and George Gilder, who look at business in a very different way. It turns out that some people look at a business as kind of a Darwinian competition. Who survives? It's going to be either rich person. Peter Thiel and Jay Richards and George Gilder say, "No, what drives entrepreneurism is creativity. You come up with a new idea like a PayPal or a Facebook or something or Uber.

Robert J. Marks:

You come out, you develop that as it's that creative aspect that is mandatory in order for you to have thriving economies and great entrepreneurship. That creativity in business is never going to come from a computer. A computer would have never come up with the idea... well, would never have come up with the idea of Uber unless the programmer programmed it to look in a set of different things. That was something which was creative which was above and beyond the algorithmic.

Robert J. Marks:

We see this idea of the inability of humans, I'm sorry, the inability of computers to create, manifest itself not only in artificial intelligence but also in business and entrepreneurship and in a number of other different places.

Robert J. Marks:

We are fearfully and wonderfully made, as the Bible says.

Larry L. Linenschmidt:

Yes. In Jay Richards' book The Human Advantage, he has countless examples of entrepreneurs seeing a need and then filling that need, and that's totally against the idea that capitalism is just about greed. He made the case that capitalism or free market enterprise is really altruistic, that the best entrepreneurs actually fill in a need. That's reality, isn't it?

Robert J. Marks:

Yes it is, yes it is.

Larry L. Linenschmidt:

Well, let me ask the question about AI a little bit differently. Self-learning, a computer teaching itself to do something different, a way that the programmer's not foreseeing. There's a program called Deep Patient and it's a way of managing information on the medical side and a couple of other programs that I read about and they solved the problem, but they aren't doing it in a way that the developer of the network can explain. Now, does that imply that there's a learnability going on in there? Some way that they're doing it? Or is everything that they're doing, even if it's not fully understood by the developer, still subject to the way that the developer set up the network?

Robert J. Marks:

Well, one of the things we have to differentiate here is the difference between surprise and creativity. I have certainly written computer programs that have the element of surprise in them. I look at them and I say, "Wow, look at what it's doing," but then I look at the program and say, "Yeah, this was one of the

solutions that I considered." One of the ideas, especially in computer search, is to lay out thousands, maybe millions or billions of potential different solutions, and you don't know what the effect of those solutions are going to be. It would be almost like putting out a bunch of different recipes for cake. You had different amounts of batter, different amounts of milk, a number of different eggs, the amount of oil that you put in, et cetera, and what you want to do is you want to figure out what the best one is.

Robert J. Marks:

If you have no domain expertise, if you want to walk around in the search space and try to find the best combination, you might get something which is totally unexpected. We did something in swarm intelligence, which is modeling social insects. We actually applied evolutionary computing, which is an area in electrical engineering, and we evolved dweebs, it was a predator-prey sort of problem and our prey was the dweebs and our predator was the bullies and the bullies would chase around the dweebs. We would evolve and try to figure out, what was the best way for the dweeb colony of the colony swarm to survive the longest? The result that we got was astonishing and very surprising.

Robert J. Marks:

What happened was that there was self-sacrifice that the dweebs learned. One dweeb would run around the playground and be chased by the bullies and self-sacrifice himself, and then, I guess, dweebs are males because I said himself, so they would kill the dweeb and then there would be other dweebs which would come out and individually they would self-sacrifice themselves. By using up all of the time in order to survive, the colony of dweebs survived for a very, very long time, which was exactly what we told it to do.

Robert J. Marks:

Now, once we looked at that, we were surprised by the result, but we looked back at the code and we said, "Yeah, of these thousands, millions of different solutions that we proposed, yeah, we see how this one gave us the surprise." Surprise can't be confused with creativity. If the surprise is something which is in the domain of what the programmer decided to program, then it really isn't creativity. The program has just found one of those millions of solutions that work really good at possibly a surprising manner.

Larry L. Linenschmidt:

Well, it seems in some way that God is the ultimate Creator and He's given us license and responsibility to be creative with this world and universe and all that's in it that He's given us. We may surprise Him with something that we do, but He knows the capabilities that we have already. How much is that an analogy to the developer of a neural network or any other type of computing capability?

Robert J. Marks:

Well, the neural network is referred to sometimes as the black box. You get responses out, you're really not sure where they come from, but if you get a response out, it's because you trained the neural network to do that, so certainly that's the case. In terms of the Lovelace Test, this is something I've been thinking about. Remember the Lovelace Test, which was a test of creativity where the computer program does something that the programmer cannot-

Larry L. Linenschmidt:

Explain?

Robert J. Marks:

Doesn't understand. Yeah, they can't explain it. I think probably everything that we can do in a theological sense God can explain, but I was thinking about our creativity. God is our "Programmer." We're more than a computer program, of course, but as our Programmer, could look inside of us and see what happens and see what the consequences of our programming is going to do. That's an interesting problem, one that I haven't explored in depth and one I need to think about a little bit more.

Larry L. Linenschmidt:

Sure. Thank you. Well, I'd like to ask you about ethics and AI. China has a massive data collection system and the data that they gather is ultimately used to determine who does the bidding of the state and therefore how they're treated for employment and loans and so on. The power of the state there is kind of like Big Brother on high-tech steroids, or to borrow from C.S. Lewis, it's like man's power over nature is the power of some man over others. Should we be concerned that AI can be misused? How would we prepare for that?

Robert J. Marks:

Absolutely. I think any new technology is not good nor bad, it's how you use it. You could talk about fire, electricity, and a number of other technologies. They're neither good nor bad. It's how mankind uses it, and so that's the decision that we have to make. In so far as ethics, AI has no ethics. The ethics belong to the computer programmer. There is the old story about the trolley problem, where you see a trolley going to impact a bus which is full of three babies and two mommies and you have the ability to throw a little switch to derail the trolley to go on another track, but on that track is a guy that's fainted and he turns out to be the world's greatest researcher in cancer curing. The question is, do you throw the switch or not? That's a question which we can debate back and forth ad nauseum.

Robert J. Marks:

I don't have the solution for that. That's for people that like to argue and go back and forth about such things, but if that switch were programmed to be controlled by an artificial intelligence, the decision would be that of the programmer and the ethics of the programmer. It would not be of the AI that controlled the switch. This is something that we're going to have to be able to look at and examine a little bit closer. The ethics is always going to belong to the programmer. We're seeing this, for example, in self-driving cars. Self-driving cars are always going to kill people, always, because they have an intelligence and that intelligence is going to let it survive, but there's always going to be contingencies or cases where that artificial intelligence hasn't been trained to anticipate.

Robert J. Marks:

The question is, how many people do you let the AI kill before you adopt the self-driving car? That's a very nice ethical question, and the answer is probably when the AI kills maybe a lot less than human drivers do. Maybe we can adopt it then, but again, this is a task that the computer programmer, the people that do this AI sort of programs. Ethics always belongs to the computer program. Always belongs to the human and if it does something unexpected, well, that's something that you have to fix. If you have a dumb programmer, the AI will do dumb things. If you have a smart program, hopefully they will do smart things which embraces all contingencies or all probable contingencies, if you will, so that the self-driving car is indeed safe.

Larry L. Linenschmidt:

Well, it goes back to C.S. Lewis. The man without chess. If we don't teach ethics and doing the right thing, then it will come out in the program just like it comes out in our personal life, doesn't it?

Robert J. Marks:

Yes, yeah. That's a good point.

Larry L. Linenschmidt:

Well, Bob, this has just been fascinating. I thank you so much for our time. I'm afraid we need to stop and I do hope we can visit again sometime. I wish you all the best with The Walter Bradley Center for Natural and Artificial Intelligence and I do hope that that center will help Christians, non-Christians, the whole world to think through the kind of issues you were just talking about.

Robert J. Marks:

Well, that's what we hope to do. Thank you, Larry. Appreciate it.

Larry L. Linenschmidt:

Thank you. God bless. Thanks to all of you for being with us today. I hope you will visit our website, hillcountryinstitute.org, to listen to podcasts of our previous programs and on iTunes and other podcast providers as Hill Country Institute Live. We have audio and video for past conferences on faith and culture issues and we're also dealing with lots of topics and concerns such as the environment, scholarship, stewardship, and fighting human trafficking.

Larry L. Linenschmidt:

We ask for your financial support for this program so we may pay the radio stations and continue to be on the air. You can donate on our website, hillcountryinstitute.org, and by calling 512-680-7993. That's 512-680-7993. For donations of a hundred dollars or more we have a book, The Human Advantage: The Future of American Work in an Age of Smart Machines by Jay Richards, which continues the discussion and looks at it from a few different angles of how we deal with artificial intelligence and technology. Thank you again for being with us for Hill Country Institute Live. We encourage you to share the love and wisdom of Jesus Christ wherever God calls you.