Can a Computer Be a Person?

https://mindmatters.ai/podcast/ep212

Austin Egbert:

Greetings and welcome to Mind Matters News.

Recently our host, Robert J. Marks, sat down with Wesley Smith on the Humanize podcast. Special thanks to Humanize for allowing us to air their talk here on Mind Matters News.

Now, Wesley's already done a great job introducing his show, so I'll get out of your way and let him take it from here. Enjoy.

Wesley J. Smith:

Welcome to Humanize from Discovery Institute Center on Human Exceptionalism where human rights meet human responsibilities. We speak with writers, thinkers, and newsmakers on the controversial issues of human life and human thriving that impact our daily lives. We are exceptional as creatures in the cosmos, as equal members of the human family, and as ethical beings. Humanize explores some of the fundamental questions. How do we thrive? How do we live well and care for what we've inherited? How do we act responsibly with one another and in the wider world? And how do we conserve the good things of this life for the future? We matter. Our actions matter. Let's get into it. I'm Wesley J. Smith and this is Humanize.

In this episode of Humanize, we're going to focus on the emergent technology known as artificial intelligence. Are we on the verge of the era of machines? Is AI destined to supplant most human endeavors and activities? Can a computer be deemed a person? And if so, should that computer be granted rights as part of the moral community? Will we ever attain immortality by uploading our minds into computers as transhumanists predict? And what the heck is AI anyway?

My guest today, Robert J. Marks II, has the answers. An electrical engineer, Marks is the distinguished professor of engineering in the Department of Engineering and Computer Science at Baylor University. He has worked in the field of artificial intelligence for more than three decades. His research supporters include NASA, the Jet Propulsion Laboratory, the National Institutes of Health, the Army Research Lab, and the Office of Naval Research. He is consulted from Microsoft and Boeing and has authored several books and hundreds of peer-reviewed journal and conference papers. He is a fellow of the Optical Society and a life fellow of the IEEE Computational Intelligence Society. He is also my colleague at the Discovery Institute where he is director of the Walter Broadly Center for Natural and Artificial Intelligence. His most recent book is Non-Computable You: What You Do That Artificial Intelligence Never Will. Bob, welcome to Humanize.

Robert J. Marks:

Wesley, it's an honor to talk with you. It's an honor to be here.

Wesley J. Smith:

Thanks very much. Electrical engineering, that's a very technical endeavor. What got you interested in that field?

Robert J. Marks:

It's interesting. I believe that in my life there's been doors which have been opened and let me say by God as to what you want to do, and the following is true. I went to a place called Rose-Hulman. The first two years I was subjected to the same course as everybody was, and then on your third year you had to choose what your major was. And I didn't know it was the night before we were supposed to do it. And I thought, "I'm pretty good in chemistry. That'd be fun. Oh, I'm a gifted nerd. Maybe I should go into mathematics." And then I thought, "I'd like to be able to fix TVs and stereos, so I should become an electrical engineering." And then the irony of that, Wesley, is that I still don't know how to fix televisions and stereos, but that's what got me into electrical engineering and it's just been a bunch of doors opening since then.

Wesley J. Smith:

Oh, that's very interesting. My talent in electrical engineering is plugging in the toaster.

Robert J. Marks:

Yes, exactly. Yes, I understand. Everybody has their own gifts, right?

Wesley J. Smith:

Yeah, that's right. You've been working on AI for decades. What was it about artificial intelligence that attracted your attention? Because I hadn't even heard about it until let's say the last 10 years or so.

Robert J. Marks:

Oh my goodness. Artificial intelligence has been around, of course, in the science fiction literature for decades. We have Isaac Asimov and some of the classic sci-fi writers writing about artificial intelligence back when television was black and white. Bradbury wrote a famous short story that was aired on The Twilight Zone called I Sing the Body Electric, and it was about a robot that came over and was a nanny. And so it's been around for a long time.

But serious scholarship in artificial intelligence has also been around since the 1960s. And some remarkable things were done in the 1960s by Bernie Widrow of Stanford University and Frank Rosenblatt of Cornell. And Bernie Widrow is especially incredible. He trained a neural network. Now, this was in the '60s, mind you. He trained a neural network to do things such as recognize speech to forecast weather, to play the game of 21 blackjack. And he got it to the optimal playing level. And this was all done in the 1960s, believe it or not. I believe that a lot of people don't appreciate the depth of artificial intelligence and the history and what was done way back in the '60s. And today we're doing things which are variations of that, and of course, much more sophisticated. We have much more sophisticated computers, we have a lot more memory, and a lot of great things were happening. So that was the first tidal wave, if you will, of artificial intelligence in the '60s.

Now, there was kind of a turf battle between, let's say neural networks where you had connectivity and another area called expert systems. Neural network would take data and it would try to recognize patterns in the data. Expert systems, on the other hand, you would go to experts and you would say... Say for example, an expert in trading. And you would say, "Okay, what if the S&P goes up 10 points and the Dow goes down five points? What do you do with Boeing stock?" And it would say, "Buy Boeing stock." And so you would enter that in your computer and that was the expert system.

So there was a big battle. And yeah, it's kind of unfortunate, but there was a book called Perceptrons, which was published by one of the AI pioneers, Marvin Minsky, that criticized neural networks. And the purpose was to put down neural networks, this connective system. And it did, and the money dried up.

But Minsky, who wrote this book called Perceptrons which downplayed neural networks, was hit by the backfire. It also dried up funding for what he was doing, which was expert systems.

Wesley J. Smith:

You're talking a little bit in just a bit of jargon that I need to understand. What is a neural network?

Robert J. Marks:

Oh, neural networks. Neural networks are supposedly modeled after the way that the brain works. The brain has a bunch of neurons, and these neurons are connected by interconnects. And one neuron might be connected to a thousand, a million other neurons. And the brain does some special stuff that engineers have tried to capture in a computer.

One of the examples, I believe it was 1948, there was a guy named Hebb who said that neurons that fire together are wired together. And this is the reason that you develop habits. You have certain triggers that are associated triggers in one neuron that are associated with actions in another neuron. And as you build these up, this path between the two neurons increases. That's a reason when I smoked, it was one heck of a problem to quit smoking because my connection between these triggers and my desire to smoke cigarettes was so strong and I needed the ability to overcome it and diminish this path.

So this is actually the way that modern neural networks are trained. It's kind of a Pavlovian sort of thing. You present it with data again and again and again. And these interconnects, they build up and they diminish in accordance to how much information is and what type of information is fed to the neural networks. So neural networks are a big part of artificial intelligence.

Wesley J. Smith:

But they'd all be within one computer or would it require several computers?

Robert J. Marks:

Well, maybe...

Wesley J. Smith:

Or both?

Robert J. Marks:

Both, I would say. Yeah, you can do it on parallel computers. But today's computers are so incredibly fast that most of it is done on a single computer, and it can take for difficult problems. It can take up to a week, two weeks, even longer, to train an artificial neural network because you have to build up these paths and knock them down. And it's kind of a Pavlovian thing. It's like you're trying to teach your dog to get you potato chips or something. And the dog comes back with Doritos and you smack him in the head with a newspaper and you go, "Bad dog." So he goes and gets the Cheetos and he comes back and you go, "Bad dog." And then finally brings the potato chips and you feed him the rest of your Slim Jim or something like that.

Wesley J. Smith:

There you go.

Robert J. Marks:

And eventually he learns how to bring you the potato chips. And that's how neural networks basically learned is this Pavlovian sort of process.

Wesley J. Smith:

We'll get this into a little more detail later. But while it's in my head, I think the difference between you and an artificial intelligence is that you are able to stop smoking. You are able to actually overcome that neural network in your brain that helped keep you because of the nicotine or whatever wanting to smoke, but you were able to actually surmount that, where an artificial intelligence would not be able to.

Robert J. Marks:

Exactly. In fact, this is... Michael Egnor is a friend of Discovery and he's also a neurosurgeon, and he talks about something called Libet's experiment. Libet was able to show that your brain actually knows what you want to do before you do. Now that sounds kind of weird, but your brain knows what you want to do before you do. And when I smoke, my brain knew I wanted a cigarette, and so it told me I needed a cigarette. And Libet also found out that we have the ability to stifle that feeling. He didn't call it free will, he called it free won't. And the idea is is that we can have these impulses to smoke cigarettes, to drink alcohol, whatever your vice is, and you have the ability to exercise free won't and therefore you kill this path between the neurons and it begins to diminish. But I notice even though it's diminished, I still have little triggers that... Gosh, I've quit smoking for... Let's see, 1970, so I don't know, 50 years. I quit smoking 50 years ago. It betrays my age, but that's whatever it is.

Wesley J. Smith:

That's all right. We're all there.

Robert J. Marks:

Yeah, we're all there. But I still get triggers today. I go on an airplane, they used to smoke on airplanes and there used to be a little ding when the no smoking light went off and everybody would take out a cigarette and smoke in the airplane. And even today when I go in the airplane and there's a little ding when the fastened seatbelt sign goes on, I have this little curious feeling that I want a cigarette. It's really interesting, even after all these years. So that path is still there, but it is really diminished.

Wesley J. Smith:

And you're in control of it. It's not in control of you. You're the director of the Walter Bradley Center for Natural and Artificial Intelligence at the Discovery Institute. What's the center's purpose and what work do you do there?

Robert J. Marks:

Well, I am the director and the purpose of the center is to advance the idea that artificial intelligence is an incredible, wonderful tool. But that's all that artificial intelligence is. It's like electricity or thermonuclear energy or a bulldozer. It's a tool which enhances the capabilities of people. But there are those, and it's primarily from a materialistic foundation, materialistic ideology that people come along and they say, "Well, this artificial intelligence is going to take over. It's going to surpass the intelligence of humans." Ray Kurzweil calls this the singularity when this point happens. And one of the purposes of the Bradley Center is to push back against this with science, philosophy and solid computer science, which shows that there are certain brick walls that artificial intelligence will never go through.

Wesley J. Smith:

Yeah. Our friends, the transhumanists, who believe that human life is inadequate, but at the same time they want to live forever, look to artificial intelligence as the secret to immortality. And we'll get into that. It almost becomes a religious belief system in technology, doesn't it?

Robert J. Marks:

Yes, it does. In fact, there was literally an AI church that was founded in California by a guy that was working on self-driving cars named Anthony Levandowski. And the first thing you do when you find a church is you apply to the IRS for tax exemption, right?

Wesley J. Smith:

Yeah.

Robert J. Marks:

So he wrote a little letter, an epistle, if you will, to the IRS saying that literally that the purpose of the church was to worship AI and recognize that AI would someday be our master, someday be our gods. This was something which was reinforced in a book called Homo Deus by Yuval Harari. And he was granted that.

But curiously, the AI church that he started had no equivalent of the 10 Commandments because a few months after he founded the church, he moved from Google self-driving car to Uber self-driving car, and he took with him 14,000 files. He was convicted of technical theft, theft of intellectual property. And he was sentenced to pay big bucks. Google sued him and it was millions of dollars that he was supposed to pay Google. He declared bankruptcy. But curiously, the last day of office, he was granted a pardon by Donald Trump. Probably by the pressure of some of the big tech donors to the Trump campaign. And so he was pardoned on the last day of office. So the punchline is that in Christianity, forgiveness of sins comes from God and Jesus, whereas in the AI church, forgiveness comes from Donald Trump.

This guy was very serious about it, by the way. Levandowski was serious about the idea that AI would take over and that he could literally have an AI church. And he was thinking of writing the equivalent of a gospel. He was thinking about getting some real estate and setting up shop as a church. But of course, the theft in the trial and the bankruptcy kind of put him back. And I think it's totally defunct now.

Wesley J. Smith:

Your newest book is Non-Computable You: What You Can Do That Artificial Intelligence Never Will, I found out a very interesting read. What made you think that that topic was important?

Robert J. Marks:

Because it was a push of, like you say, the transhumanists and those that embrace artificial intelligence as a future God, that we wanted to indeed push back on that. And I give expressions, or not expressions, but discussions of the book of things that AI will never do. There's kind of the obvious list of love and compassion and emotional sort of things. But even more deeply, artificial intelligence will never have sentience, it'll never be creative, and it will never understand what it's doing.

Wesley J. Smith:

Oh, that's a key point I think. And we're going to get into all of those issues. But it does not understand what it's doing. It's doing what it has to do. So let's define a few terms some more. What is artificial intelligence and how does that differ from just say the computer that I'm recording this interview on?

Robert J. Marks:

Well, in my world, if you go to different conferences, people tease apart the idea of computational intelligence, machine intelligence, artificial intelligence. But that's not the way that's presented in the media, is it?

In the media, it's kind of presented as anything that a computer does that you look at and you go, "Wow, that's really incredible." And I think that that is probably a good enough general definition for the discussion that I like to have because that definition includes a computer and my attention is on the computer and the limitations of what a computer can do. So that's my definition of artificial intelligence. Anything that a computer does, which you are just astounded with.

The interesting thing is that we have artificial intelligence that surrounds us. I think of, well, for example, GPS. I kind of use it daily because I have no sense of direction. I get lost going home about once a month. And it's something that we're known by. It's just here and it's something that we're used to. Same thing with cell phones with Uber and Lyft. I use Lyft last weekend. What a remarkable service. I use Alexa quite a bit.

Wesley J. Smith:

Yeah, I was going to ask you about Alexa because she may come on in behind me. Of course, she's not a she. That's just a female-sounding voice. But I can ask, we'll call her A, so she doesn't come on, what time it is and she'll tell me immediately. I can tell her to play certain music and she'll play it immediately. How does that operate? I mean, she's not... That program is not intelligent, is it?

Robert J. Marks:

No, it isn't. We go back to something that John Searle mentioned that the reason that AI can't understand. He said that, I don't know Chinese, but imagine me in a room with a bunch of file cabinets and through the doors slipped a little question that's written in Chinese. And he said, "I'll go to these file cabinets and I'll look until I get a match. I don't know what it's saying, but I'm going to look until I get a match and I'm going to copy down the answer and I'm going to slip that answer in Chinese through the door."

Now, from the outside, it kind of looks like whatever is inside that room knows Chinese, understands Chinese, and it's just astonishing. But Searle in the room doesn't understand Chinese. It's the same thing with Alexa. When you ask Alexa a problem, it is in this humongous room, which probably includes all of Wikipedia. We have the memory and we have the computational resources to do this now. It is in this humongous room. It does some language recognition on your voice, and it goes to this big humongous room and it looks through all these file cabinets until it finds the response, which it thinks that you want.

Wesley J. Smith:

It thinks. Is that a correct terminology there? It thinks that I want?

Robert J. Marks:

Well, maybe thinks isn't the right word, but I will say that Alexa screws up a lot.

Wesley J. Smith:

Yes, but it's programmed to give the response that its prior programming, its prior experience would seem to indicate is correct. Is that a good way to put it?

Robert J. Marks:

Yes, yes.

Wesley J. Smith:

Because we're going to get into the difference between what a computer does and human thinking. And I just wanted to keep that pretty clear. I mean, we obviously use the term it thinks, but it really doesn't, does it?

Robert J. Marks:

It doesn't. It's in this incredible Chinese room, except it has billions of file cabinets and it looks through all of those and tries to find the best match. And when it looks, it often finds not the match you're looking for, but what it identifies. I was going to use the word think there...

Wesley J. Smith:

Yeah, because we naturally anthropomorphize things. It's a natural human tendency. I'm not trying to be a jerk, but I'm just saying I want, for my understanding and our audience's understanding, to keep certain terminologies accurate in the sense of how we're going to distinguish between what a computer does and what we do as human beings when we think.

Robert J. Marks:

Exactly. In fact, I talk about this in my book. When artificial intelligence is hyped, they use what I referred to extensively as seductive semantics. They are used extensively in artificial intelligence media to talk about artificial intelligence. You hear words about artificial intelligence being sentient, about artificial intelligence being conscious, about artificial intelligence thinking. And these are words which are used without definition and what needs to go back and define these words before they apply it.

What do we mean when we say that artificial intelligence is sentient? It's kind of like political talk. Great politicians can talk and use these seductive semantics, and everybody in the audience from the left to the right thinks that they should support this candidate because of the seductive semantics that's used. It's the same exact thing in the hype of artificial intelligence. All of these fuzzy words are used without proper definition. And that's one of the things I'd like to try to do in the book is to appropriately define words before we discuss them.

Wesley J. Smith:

Right. And it's important for us so that we can understand the distinctions. Human exceptionalism is my passion. And one of the dangers of humanizing artificial intelligence is that it actually undermines and corrodes our understanding of the important unique dignity of human life. And which is I think a predicate to human liberty. And that's one of the dangers I think that this transhumanism movement has and this drive to... It's a reductionist drive in essence to say that all that matters is our ability to have these synapses that fire rather than the full depth of human experience.

Robert J. Marks:

Yes. Yes, absolutely.

Wesley J. Smith:

So what makes AI so powerful?

Robert J. Marks:

What makes AI so powerful? It is number one, the algorithms. And we should define an algorithm since we're going to use it.

Wesley J. Smith:

That's right. That's my next question.

Robert J. Marks:

Okay. An algorithm is a step-by-step procedure for doing something. If Google Maps, when they tell you how to get from point A to point B is giving you an algorithm. You go down I-35 for two miles, take Exit 32A, take a right at the light, et cetera, et cetera. So that's a step-by-step procedure for doing something.

An algorithm is nothing more than a recipe. In the book, I use the example of a German chocolate cake. You have the input, which is all the ingredients in the cake, but then you have the algorithm, the step-by-step procedure that you have to do to make the cake. You put the mix and you stir in the milk and do all of this stuff. So that's an algorithm.

Now, it turns out that the only thing computers can do are algorithmic. If something is non-algorithmic, it is not computable. And if it's not computable, it's something that computers can't do. And the interesting thing is that this is true not only for today's computers, but yesterday's computers and computers of the future. Doesn't matter how fast, doesn't matter how incredible they do, they are still not going to be able to do non-computable things.

It was shown back in the 1930s by Alan Turing that there were some certain problems that were non-computable. Most undergraduate computer scientists are introduced to this through something called a halting problem. Without getting into the details what it was, but Turing showed mathematically that this was non-computable. There have since been a number of different things which have shown to be non-computable, which means it can't be done by a computer. Now, if that's the case, we have to ask ourself, are there things which humans do, which are non-computable? And those I would claim include sentience, consciousness, understanding, and creativity.

Wesley J. Smith:

Yes.

Robert J. Marks:

All of those are beyond the capability of an algorithm or a computer program to do. You can't program creativity.

Wesley J. Smith:

So when we talk about software being developed and so forth, and this is coming from a place of total ignorance, are they actually building different algorithms in a logical step-by-step process so that the desired function will operate?

Robert J. Marks:

Well, here's the basic scenario. There was a AI computer program, I believe it was last year, that discovered new antibiotics. Now, any time that you do anything creative, it's iterative. It takes a number of different trials. You try something, you get something back and you say, "Woo, it has a few bugs here, so let's fix those." And then you do it again. And you can do it a number of times. In fact, you probably have evidence of this in your household. Formula 409 is so named because it took 409 iterations in order to perfect the formula. WD-40 translates to water displacement system based on the 40th trial.

Wesley J. Smith:

Ah, I didn't know that.

Robert J. Marks:

So all of these are iterative. Yeah, it's very interesting.

Wesley J. Smith:

Yeah.

Robert J. Marks:

So both of these guys probably perfected this in a wet lab, but imagine we could do all of this in a computer, and that's exactly what the computer does. It can do not only the design, but it can say how good the design is. Where does this information come from? The human programmer.

And so some of the great artificial intelligence is based on search that if you look at these antibiotics, I think it was something like searching through a hundred million different molecules or whatever, different antibiotic configurations before they found these new exciting antibiotics. This is something that would've take a long time in a wet lab, but they had computer models to do it.

But where was the creativity in here? The creativity was in the programmer deciding what sort of problem they would solve. It was in the computer program to decide what solutions would be sought. And all of the creativity comes from the computer program. The computer just sits there and chunks through the different solutions and...

Wesley J. Smith:

What about the idea that a computer can self-program? Does that change what you're saying here?

Robert J. Marks:

Well, no, it doesn't. Computer cannot self-program. Now, there are computer programs that write other computer programs. But this is one of the things that people like Ray Kurzweil and Yuval Harari talk about is computer AI writing better computer AI, writing better computer AI, and eventually reaching the so-called singularity and so-called reaching super intelligence. But there is not any evidence to date of computer programs being creative, meaning that a computer program cannot write another computer program, which is beyond the intent of the original programmer.

In fact, this is the way that Selmer Bringsjord of Rensselaer defines creativity. He says that we will experience creativity if a computer program does something which is beyond the explanation and the intent of the computer programmer or somebody of equivalent skill. And to date, there has been no computer program, which has done that. And artificial intelligence writing better computer programs with results and attributes that are beyond the intent of the original computer program have to date not been distributed or demonstrated. I like to say that artificial intelligence has written computer program languages like Python and C++, where super intelligence and so-called artificial general intelligence is written in PowerPoint slides and press releases. You don't see any of this in practice, a computer program writing a computer program which is beyond the intent.

Now, this does not mean that the computer program results are not surprising. Often the computer program results are surprising and unexpected. But if the computer programmer looks at that, and I've done this many times in my AI research, I'm astonished at some of the results. I say, "That wasn't the solutions that we gave them. We gave them a whole bunch of different solutions," just like the antibiotic example. And yeah, this turned to be one of the solutions. So it did what I told it to do, even though the output was surprising and unexpected.

Wesley J. Smith:

And that's not the same thing as thinking. How does human thinking differ from an AI following algorithms and finding perhaps some amazing discoveries? But it's still not thinking, right?

Robert J. Marks:

Well, it isn't thinking. Let's talk about thinking in the terms of creativity, because I think that that's the result of successful thinking is creativity. It turns out that humans have this ability of having flashes of genius. We see this historically with people like Nikola Tesla. He was walking along the beach and he had the idea for the brushless motor, which we all use today. And he said, "It hit me like a flash of lightning." And in his biography says that he smoothed out some dirt and kind of drew a schematic in the dirt so he wouldn't forget it.

A flash of genius was from the great mathematician Friedrich Gauss who said, "I woke up one morning, I had the idea and I had the solution for this problem. It wasn't like anything that I talked about before." He also, coincidentally, without knowing about Tesla or Tesla knowing about him, called it a flash of lightning, which hit him.

We also see this in the creative arts. Paul McCartney of the Beatles, for example, woke up one morning and he had this melody going in his head, and he went over and he worked on it in the piano, and he didn't know where it came from. His father listened to a lot of jazz. He thought maybe this was something he heard before. So he began to shop it around and he said he was going to treat it kind of like somebody who finds a wallet. You take it to the police, and if the police don't find the owner in a couple of weeks, you got to keep the wallet. So he was going to shop it around, see if anybody had heard about it. And he went over and it was the melody to Yesterday, which turns out to be one of the most covered songs of all time.

So we see this in the arts. We see this in science. And I don't know if this is thinking, but it certainly is in the area of creativity. Creativity in the human side often results from this flash of genius. It used to be that the US Patent Office required all patents to display a flash of genius. That was their policy, but that has been abandoned as you can tell ... What's that?

Wesley J. Smith:

You're talking about epiphanies?

Robert J. Marks:

Yeah, an epiphany. An Archimedes epiphany, where he sat down in the bathwater and he recognized how to solve a problem he'd been working on. He popped up and the myth says he ran through the streets naked crying, "Eureka. Eureka." That was his flash of genius. So we see that...

Wesley J. Smith:

And AI is not capable of an epiphany?

Robert J. Marks:

No, because you cannot program a flash of genius into artificial intelligence. It's never been done. If you're going to get a result, it's within the scope, within the silo, if you will, of the original programmer's intent.

Wesley J. Smith:

It's interesting. In your book, you use the example of jazz when you're talking about creativity, and how humans, the jazz musician, particularly when they're riffing, right? They just go off and they do their riff, and it's based on the creative feel of the moment. And what they're doing in that particular moment may be completely different than what they do the next night's performance. And you basically say, "AI can never do jazz." And I thought to myself, that's a really brilliant example.

Robert J. Marks:

Okay. It can't do real jazz.

Wesley J. Smith:

Real jazz and the riffing part. Yeah.

Robert J. Marks:

Yeah, the riffing part. Yeah. There are those that claim that artificial intelligence is creative in that it can compose music. But here's what goes on if you want... Let's not take jazz. Let's take Bach for example. If you want to have your artificial intelligence compose Baroque music, what do you do? You gather a bunch of Baroque music of Handel and Bach, and you present it to the artificial intelligence. That defines a silo in which the artificial intelligence can work. Once you train it on this artificial intelligence on Bach and Handel, guess what it generates? It generates stuff that kind of sounds like Baroque music, but never will it generate something that sounds like Wagner or Stravinsky...

Wesley J. Smith:

Paul McCartney.

Robert J. Marks:

... or jazz. It's never going to do that. It always has to think inside the box, inside of what it was programmed to do.

Wesley J. Smith:

It's never able to color outside the lines to use that old cliche.

Robert J. Marks:

Okay. Yes, yes.

Wesley J. Smith:

Yeah, that makes sense. That makes sense. And that gets to the issue of free will. I mean, human beings, and I know this is contested by some, but human beings have free will. That's why we can be held morally accountable for our actions. AI can never have free will, can it?

Robert J. Marks:

No. It can't. It's going to go to what it was programmed to do or what it has learned to do. We have, for example... Oh, this is kind of interesting. In self-driving cars, for example, there's always the question of AI ethics and whether a self-driving car, if it's programmed. And it has two choices, it can either mow down and kill Albert Einstein or it can kill a mother with her three kids. And the AI has to make a decision, a moral decision, if you will. That decision will not be made by the AI. That decision will be made by the computer programmer. So it isn't the AI's responsibility for that morality. It is the programmer's responsibility.

Wesley J. Smith:

It cannot be a moral agent because it really doesn't know right from wrong and can't even, as we discussed, think it through. It just does what it's programmed to do.

Robert J. Marks:

Yes, exactly.

Wesley J. Smith:

Let me ask the question in a different way. Can an AI ever be irrational?

Robert J. Marks:

Yes. Irrational in the sense of being irrational from the point of an observer. A classic example, and this happened a number of years ago, was that the Soviets during the Cold War developed a high technology to decide whether the US was being attacked by... I'm sorry, whether the Soviet Union was being attacked by the United States. And so they had these missile detectors. And there was this false alarm one day when the Oko, it was called an Oko system, O-K-O, when it was stimulated and sirens went off and it says they were being attacked by US missiles. And the protocol was to launch a counter strike.

But fortunately, the person in charge, if I remember, he was Lieutenant Stanislav, I believe. He said, "This doesn't make sense. Because if the US was doing a military strike, they would just not send one or two missiles over. They would do just a preemptive strike of just launching a number of different missiles towards the Soviet Union." So he called his superiors and they canceled it off. So this was something, I believe, an example of AI, or at least high technology being irrational and coming to the incorrect conclusion. Later, they found out that the Oko system had mistaken the reflection of sunlight off the clouds as a US missile. And so it was something which was totally incorrect. But that one guy, that one Lieutenant Stanislav, he saved possibly a nuclear exchange, really.

Wesley J. Smith:

Because he had something you talk about in your book called common sense...

Robert J. Marks:

Yes, he has common sense.

Wesley J. Smith:

... that AI doesn't have.

Robert J. Marks:

Yes, it doesn't have. And I used to believe that common sense might be able to be programmed into a computer. But I'm starting to believe with more and more results that this might be something which is also non-computable. One of my favorite ideas of ambiguity and the inability of artificial intelligence to understand and resolve ambiguity are so-called flubbed headlines. And I have a great collection of them, by the way, over a hundred flubbed headlines. These are headlines where the author means one thing, but if you read it, you can also read a humorous interpretation of it.

For example, helicopter powered by human flies. Now, if you think about that, you think, "I know what the guy meant, but I also know the incorrect interpretation of it." And you know that immediately, but the AI will never figure that out. It doesn't know which interpretation it is. Hospitals sued by seven foot doctors. Is it seven podiatrists or is it these really tall doctors?

Wesley J. Smith:

Seven foot tall doctors. Yeah.

Robert J. Marks:

Farmer bill dies in house. Was it a farmer or was it a piece of legislation? So there are a number of those. And one of the challenges with artificial intelligence is there's been these claims of this certain narrow type of ambiguity called Winograd schema. But a lot of these Winograd schema are actually answered online. And so if it's on Wikipedia, if it's on the web, it can resolve it because it's seen it before. But just by its own...

Wesley J. Smith:

You can get that data very quickly, right? That's one of...

Robert J. Marks:

Yes, exactly.

Wesley J. Smith:

... the great strengths of AI for us is that as a tool, it can gather a lot of information much faster than we can looking in a library book.

Robert J. Marks:

Yes, it's kind of like a Searle's Chinese room in a way. It looks up the ambiguity and resolves it because the answer to the Chinese question is the resolution of the ambiguity of the flubbed deadline.

Wesley J. Smith:

Recently, a Google engineer claimed that a chatbot had attained personhood and was entitled to rights. First off, what is a chatbot? And then why did the engineer believe that?

Robert J. Marks:

A chatbot is a piece of artificial intelligence that communicates with you. And you type in a question, it'll try to answer, for example. And there's some chatbots that can carry on extended dialogue.

Wesley J. Smith:

Again, based on its programming only, right?

Robert J. Marks:

It based on its programming. But after a while, you can kind of tell that there isn't the appropriate depth that is needed.

So this guy at Google, I think his name was Lemoine, L-E-M-O-I-N-E, he said that he had talked with a chatbot named LaMDA, which was an acronym for something, and that he swore it was sentient. In fact, he even said that LaMDA became afraid and said, "Do I need a lawyer," which is really a curious sort of response.

Well, it turns out that a lot of this hype that comes over artificial intelligence hits the media, boom. But then when it's discredited, you don't see a lot of backpedaling, if you will. And it turns out, if one backpedals on the Lemoine thing and looks at the media that came out after the big report, you'll find out that it isn't so impressive. Because this LaMDA was trained with...

Well, they asked so-called crowd workers to respond to randomly selected samples in evaluating data sets. So questions were asked and answers were given. And it turns out that they had over 6,400 dialogues and 121,000 different responses that was used to train the artificial intelligence LaMDA to give nice answers. And the crowd workers, these people that did this, they were explicitly informed to reply in a safe, sensible, specific, interesting, grounded and informative manner. And so guess what LaMDA learned? It learned how to respond in a safe, sensible, specific, interesting, grounded and informative manner. So it was trained to do exactly what it was doing.

There was also a Washington Post response on Lemoine's suspension. He worked for Google and he was suspended. The Washington Post said that some chat logs were leaked from the Washington Post were leaked from Lemoine. And there were disclaimers that, and this is a quote, "This document was edited with readability and narrative coherence in mind." This is something which happens a lot in claims that artificial intelligence is creative. That there will be a number of responses which are cherry-picked, and you only see the cherry-picked responses. Sometimes it's also edited, and we see a lot of that also. So no, there was no sentience there.

Wesley J. Smith:

By sentience, do you mean consciousness? Because a fly is sentient. A fly can tell... If a fly lands on my arm and I'm go to swat it, it can sense my hand coming to swat it and will fly away. That's kind of what I always think of sentience versus consciousness and rationality. When we're using the term sentience, are we talking about consciousness and rationality?

Robert J. Marks:

Yes, yes. And you're right. Sentience sometimes is equated with the idea of self-aware. And shoot, I get in my car and I back up and my car can tell me that I'm getting close to bumping into another car, and it

goes, beep, beep, beep. It's aware of its environment. Does that make it self-aware? Well, in seductive semantics, yeah, I guess it does. But I think in the general definition that we are using though would not be self-aware because the sentience also requires a degree of understanding. I don't know that much about flies, but I wonder if they understand why they're supposed to.

Wesley J. Smith:

Right. But they can sense that there's something coming at them and they fly off. I mean, it's just that there's a... not a sentience, but an ability to experience, to feel, to have an awareness. And what we've discussed is actually an AI is not aware. It is just performing algorithms.

Robert J. Marks:

Look, let me give you an example of a type of sentience called qualia. Qualia is what you do experience. I'm looking at your shirt right now, and I see it's red. I'm having an experience now. I'm experiencing red. And imagine me trying to explain to a man who has been blind since birth my experience of seeing red.

Wesley J. Smith:

Right. You couldn't.

Robert J. Marks:

I could tell him the wavelength. I could tell him the frequency of the red. I can tell him apples are red, the blood is red. But he will never be able to duplicate that sentience, that qualia that you and I experience. Now, you and I can talk about red because we both experienced red, but for that blind man, it would be impossible for me to communicate the experience that I have.

Now, if I can't communicate my experience to that blind man, how the heck am I going to write a computer program that has any sort of ability to experience qualia or what it means to see red or taste hot buttered corn in the cob or bite into a lemon or pain even. All of these qualia things are things that we experience, which are beyond the capability of a computer.

Wesley J. Smith:

And one of the reasons for that, it seems to me, is that we have bodies. And a computer, even if it ends up like being a cyborg with some organic materials, isn't really an organism. And an organism, if you suddenly said boo to me and it scared me, I would have a shot of adrenaline. That shot of adrenaline would give me that feeling of fear. If my wife comes in the room and I'm feeling particularly romantic, I might have a very warm sense of warmth in my chest because of the feelings I have for my wife. I don't see how a computer could ever have that experience of "feeling" because it's not an organism. Am I onto something or am I a little off phase?

Robert J. Marks:

Absolutely. And I think that that was one of the premises we kind of started out with is that these emotional things, love, compassion, fear, all of these are non-computable. They are something that are... These are things which are uniquely human and have no place in artificial intelligence.

Now, with artificial, I went to France one time, and I met a guy called a nose. You ever heard of a nose?

Wesley J. Smith:

No.

Robert J. Marks:

They work at a perfume factory.

Wesley J. Smith:

Oh, yes. Okay, sure.

Robert J. Marks:

Their thing is to come in and smell perfume, and they say, "Ooh, I sense a little bit of citrus, a little bit of a cedar in there." And these are guys like the taste wine too. They have really high tuned senses.

Now, can we duplicate that with technology? The answer is yes. We can have artificial tongues, we can have artificial noses that have little sensors that look at for certain chemicals that are in the air. And they can detect the presence of the smell of lemon or the smell of other citrus. But do they experience the qualia? No, they don't.

Wesley J. Smith:

Yeah, yeah. And different noses might approach the same perfume differently, right?

Robert J. Marks:

I don't know that much about noses, but I imagine yes.

Wesley J. Smith:

That the human subjective experience is not anything that I would assume that an AI could duplicate.

Robert J. Marks:

One of the things I do know is the US Patent Office will not patent aromas or things or taste.

Wesley J. Smith:

Oh, that's interesting. Yeah.

Robert J. Marks:

Yeah. It just isn't something, as you mentioned, that can be quantified.

Wesley J. Smith:

You talk about AI and swarm intelligence in your book. What do you mean by that?

Robert J. Marks:

Well, it turns out one of the most interesting areas of artificial intelligence to me, because it's an area that I've done research in, is swarm intelligence. This is the ability of dumb bugs to do smart things. If you think, for example, about ants building an ant hill, all of those ants don't understand that they're building an ant hill. They have no idea of the end objective.

One of the simple algorithms, for example, is if I bounce a bunch of Skittles on the kitchen floor and they all bounce around making their little clickety-clack sounds, and I set a bunch of bugs, these little robots out, and the purpose of the bugs is to, anytime that they run into a Skittle, they pick it up, and then they

go around randomly. And anytime they hit another Skittle, they put it down. This is all that they know. They have no idea what they're doing.

And in fact, we think about it, "Okay, you hit a Skittle, you pick it up, you walk around. You hit another Skittle, you put it down. You walk around into a Skittle, you pick it up, walk around. Then you put it down when you hit another Skittle." It turns out that what they're doing is stacking the Skittles. And it turns out that this is exactly what happens, for example, in termites, clearing wood, if you will, saw dust and kind of stacking it also as an algorithm, which is used to describe the way that ants clear their dead. They end up stacking them up.

So this is the example of swarm intelligence. It's dumb bugs doing smart things. Now, where does that come from? That has to come from a super duper programmer, a computer programmer that actually controls what these bugs do and the so-called emergent behavior of stacking saw dust or stacking dead ants or whatever. And there's numerous examples of this. We were hired by the Navy to do some work in doing the opposite problem. If we wanted swarm to do something, what are the simple rules that we should tell the insects to do? And we used some artificial intelligence to actually do the inverse problem, which was a very, very difficult problem. So we wanted a result. We wanted the swarm to do something, but we had to figure out the simple rules that the bugs would go by.

Wesley J. Smith:

And so we're not going to get into intelligent design here today. But you're basically saying that in nature, there's a cause and effect that goes beyond what the bugs do. Is that right?

Robert J. Marks:

Yes, yes. And this is all controlled by the programmers. And in fact, swarm intelligence doesn't have to be bugs moving. If you think about the human body, for example, in the lungs, there are a certain set of cells that they all work together. Do you think one cell knows what the other one's doing? No, they don't. They're all there doing their own individual thing, and there's this incredible emergent behavior which allows the lungs to work. So even swarm intelligence would be applicable to human organs, for example.

Wesley J. Smith:

Is the analogy that AI does that same kind of swarm intelligence when it does its computations and its various computing efforts?

Robert J. Marks:

Well, there are swarm intelligence algorithms to do artificial intelligence. I think the bigger point is that this is akin to a computer program not understanding what it's supposed to do. So we can do a computer program to do swarm intelligence, but that program doesn't know what it's doing.

Wesley J. Smith:

Computer doesn't know what it's doing. It's just doing...

Robert J. Marks:

No, exactly.

Wesley J. Smith:

... what it's programmed to do and it can't act... Again, a computer cannot act outside its programming. And of course, human beings are not programmed in that sense and we have free will, and we can.

Robert J. Marks:

Yes, exactly.

Wesley J. Smith:

We're almost out of time. But I want to get into our friends, the transhumanists a bit. Their great desire is immortality, and they keep trying to work out various means by which at least indefinite life could be maintained in this corporeal world. And their favorite one seems to be that at some point come to singularity, we will be able to upload our minds into computers and therefore live forever in cyberspace as either an individual consciousness or even part of a group consciousness. What do you make of that notion?

Robert J. Marks:

Well, first of all, you have to go back to their foundational assumption, which is materialism. They believe that our brains and who we are is all controlled by us being meat computers, if you will. But everything which goes on in the brain is materialistic. We talked previously about the idea of non-computable. Now, if these things such as emotions, creativity, sentience and understanding are non-algorithmic, that means a computer program and a computer in general cannot duplicate these characteristics. So that means that if you upload yourself into a computer, you can only upload the algorithmic part, the computable part, and you are not going to be able to touch the non-computable part. And I maintain that if we just had the algorithmic part of ourselves, that would be a very boring person. And so that is not going to happen.

Wesley J. Smith:

And you're saying it would be incomplete. It wouldn't be... If I tried to upload my mind into the computer, whatever the computer spitting out isn't really me.

Robert J. Marks:

No, it isn't you. You would still be able to add. You would be able to go through routines. You would be able to do things algorithmically. But as far as experience, the things that we talked about, emotion, qualia, sentience, et cetera, that would not be possible with the uploaded mind.

Wesley J. Smith:

It would be mimicking what it was able to receive, assuming we could ever do this. But it would only be mimicking me. It wouldn't be me.

Robert J. Marks:

Yes. I don't know if you've ever seen the movie Invaders of the Body Snatchers.

Wesley J. Smith:

Yes.

Robert J. Marks:

And it kind of reminds me of Invasion of the Body Snatchers, where these big pods came in from outer space and they began to replicate human beings, but they replicated them non-algorithmically. And if you looked at these replications from afar, you would go, "Oh yeah, that's Frank. I recognize Frank." But if you go up, and for people who know them, that have familiarity with them, they say, "Man, something about Frank isn't there." And that is that non-computable aspect. That's what would happen if we were non-computable you. We wouldn't all be there. But anyway, I think that...

Wesley J. Smith:

It would just be facsimile.

Robert J. Marks:

It would be facsimile. Exactly. And it would be very, very boring facsimile.

Wesley J. Smith:

So AI is a powerful tool, you say. It is something that can benefit the human condition. So how do we go about accomplishing that while at the same time, not unleashing the dystopian dangers that we've sometimes seen reflected in movies or that could have happened with that Russian false alarm on the nuclear attack?

Robert J. Marks:

Yes. The answer is that we have to be careful with the AI tool. We had to be careful with the tool of electricity. In fact, Edison, at the time that Tesla was pushing alternating current, was trying to show the world how dangerous it was by electrocuting animals with alternating current at state fairs. Most famously, there was a elephant named Topsy that was in New York City at one of the fairs and they electrocuted them to show how dangerous alternating current was. And there were a lot of people that were afraid of it. Well, this has been mitigated. We use electricity all the time. Is it still dangerous? Well, heck yeah. You can have frayed wiring that burns down your house. You can have an electrician, one of these line workers that touches a live wire. So it still has the dangers, but we have mitigated it.

So therefore, we have to be able to, in the development of artificial intelligence, to mitigate how it is used. We have to pay attention to the ethics. And there are two types of ethics. There's number one, the design ethics. This is the engineer's problem. And the design ethics is to make sure that the AI does what it is supposed to do and doesn't do anymore. This is true of any engineering design. We have to be very careful in the final product that it does what it's supposed to do and nothing else.

Then there is the end-user ethics, which says what the AI is going to be used for. This, for example, would be a commander in the field deciding whether he wanted to deploy artificial dangerous killing robots that were foundational in terms of its use of artificial intelligence. So we would have to be very, very careful about that. But that's the end-user ethics and it's different from the AI ethics.

We have a National Science Foundation grant right now that we're working on that deals with artificial intelligence, ethics. And my parent society, the IEEE, you mentioned in the beginning, IEEE, is also interested in the ethical use of artificial intelligence. And the guy that I work with, Pablo Rivas, is on the board there that is adopting standards in electrical engineering for the use and deployment of artificial intelligence.

And so yes, just like electricity, just like thermonuclear power, it's going to be dangerous. But both electricity and thermonuclear power are only tools. So is artificial intelligence. It is a tool. And if it is used for evil, it's either going to be due to carelessness or it's going to be used because of an evil person

behind it. And that's true of all tools. It will not be because the AI has taken over in the sense of the Terminator movies, for example.

Wesley J. Smith:

Right. So it's like a car. I mean, if I drive a car to the store, it's one thing. If I use a car to plow into people, it's another thing. But the car is just the car. It's a tool. And the outcome in that, the evil, if you will, would be me. So an AI is not a moral agent and cannot be evil. Humans can be evil, but machines cannot be evil. And it seems to me, I'd like your opinion on this, that AI is a very valuable piece of equipment, but it has no inherent moral value because it is not a living being, it is not a moral agent, it is a tool. So even though it could be a multimillion dollar machine, it has no intrinsic value any more than a broken toaster does because it's inanimate.

Robert J. Marks:

Yes, exactly. And in the development of the ethics that is required to control AI, an interesting example is the self-driving cars. You're familiar with the Uber a few years ago that killed the pedestrian? And I believe that there's been literally hundreds of Teslas that have killed people or resulted in accidents when the Tesla was on its so-called autopilot.

Here's the interesting thing. The courts have determined that Uber and Tesla were not responsible for those accidents. Why? Because they were ethically designed. One of the things that Tesla and Uber required their drivers to do is to pay attention all the time. It turns out the Uber driver, for example, at the time of the accident, was watching a video, and I think it was called the... I forget what the name of it was called. But she was watching a video when she killed this. It's the same thing with these Tesla accidents. These people are not going and not obeying all of the algorithm that the self-driving cars would do. So many of them stream videos or doing something else.

But the cars are very, very robust in this. I was in IKEA the other day, a self-driving car, and the driver demonstrated to me that if you didn't put your hands in the steering wheel, after a while, the car would let you know, they said, "Hey, you're not paying attention. You have to do this." So part of the AI ethics in that case was to recognize, yes, there are certain capabilities that the self-driving car can do, but it also includes in the operation of the AI for a human to be in the loop. And when a human isn't in the loop, we will have so-called self-driving cars.

I can't wait for that. I think that it might be possible where I can get my backseat and program myself to go from Waco, Texas to my place in West Virginia, go in the back and take a nap. I mean, that would be really cool. But that is a long way off. And right now, the AI is such that AI ethics requires that a human be in the loop at all times. And as a rule, that's always something that should happen if possible. There are places where it's not possible. But any time a human can be in a loop, it needs to be in that loop. We're not to the point yet where we can go to totally autonomous self-driving cars.

Wesley J. Smith:

But even if we came to an autonomous self-driving car where, we've seen those in science fi too, you just get in the car and it takes you where you want it, it doesn't know where it's going. It's just going where it's programmed to go.

Robert J. Marks:

Oh, absolutely. And it's programmed with a bunch of rules to keep you safe.

Wesley J. Smith:

Yeah.

Robert J. Marks:

And the big problem is, oh, are there going to be some contingencies that happen, which it hasn't been programmed with? So that possibility is always going to be there.

Wesley J. Smith:

And sometimes you'd need the creative human mind to work your way out of a circumstance that was not anticipated.

Robert J. Marks:

Well, you do too. But humans have that same problem, right? You don't anticipate all of the problems that you might have if you go out on a drive.

Wesley J. Smith:

I don't anticipate everything I might have in interviewing. Well, we're out of time, but I'd like to know, is there something I should have asked you about AI that I have failed to do so that might be worth discussing a little further?

Robert J. Marks:

Yeah. One of the things that you can mention and which I'll mention is our sister centers at Discovery Institute. Mine is the Walter Bradley Center for Natural and Artificial Intelligence, named after one of my heroes, Walter Bradley. But we have a website called mindmatters.ai. Politicians say to mention the website three times. So let me do that. It's mindmatters.ai, mindmatters.ai, mindmatters.ai.

And we have great people working for us there. We have not only journalists, but we also have people such as attorneys. We have a neurosurgeon that writes for us. We have economists that write for us and a number of PhD computer scientists that write for us. So the expertise is at a very high level and it does... Well, it has the common denominator of pushing back against some of the claims, which are made about artificial intelligence and trying to keep us on a more grounded playing field of what AI can and can't do.

Wesley J. Smith:

And we'll put mindmatters.ai in the program notes. You also have a podcast? What is that?

Robert J. Marks:

Oh, we do. It's Mind Matters News and it's also available at the same website.

Wesley J. Smith:

Wonderful. Well, we're out of time, but I want to ask what next for Bob Marks?

Robert J. Marks:

That's a good question. This is something which I've been thinking about. Maybe I'd like your opinion. I'd like to write a book on supply-side academics.

Wesley J. Smith:

Supply-side academics.

Robert J. Marks:

Supply-side academics. I have been in the field of academia for, I don't know, 40 some years. And where it's going is kind of disgraceful and we have a degradation of a number of different things which are happening in the academy. And so I was thinking of writing a book about that.

Wesley J. Smith:

Well, when that happens, maybe we'll have you back.

Robert J. Marks:

Okay. That's right. That'll be great, Wesley. Thank you. This has been a great conversation.

Wesley J. Smith:

Thank you for being on Humanize, and we'll talk to you again.

Robert J. Marks:

Okay, thank you.

Wesley J. Smith:

Thanks for listening to Humanize from Discovery Institute's Center on Human Exceptionalism, where human rights meet human responsibilities. Discover all the good work of the Center on Human Exceptionalism by visiting discovery.org/human.

We can only do this work speaking on behalf of human life, human thriving and our exceptional place in this world and our cosmos with your support. We invite you to make a one-time gift today and to consider starting a monthly gift to support the Center on Human Exceptionalism and this show. Wherever you're listening to Humanize, please take a moment to write and review the show. You matter. Your actions matter. Be bold, be exceptional, and be back soon.