

OpenTheo

Is A.I. intelligent? | Rosalind Picard

December 30, 2021



The Veritas Forum

PART OF A SPECIAL 6-WEEK SERIES | Should we be worried about robots taking over the world? Dr. Rosalind Picard, an A.I. researcher at MIT, says no. But, there are real things to consider about our relationship with technology. We talk with Dr. Picard about the past, present, and future of machine learning and artificial intelligence and hear how her current work is literally saving lives. Like what you heard? Rate and review Beyond the Forum on Apple Podcasts to help more people discover our episodes. And, get updates on more ideas that shape our lives by signing up for our email newsletter at veritas.org. Thanks for listening!

Transcript

Five years ago, Hanson Robotics activated a humanoid social robot named Sophia. She, or perhaps more accurately, it, debuted at South by Southwest and made many other public appearances at gatherings like the United Nations and on TV shows like Charlie Rose and the Tonight Show with Jimmy Fallon. At the Tonight Show, the CEO of Hanson Robotics talked with Jimmy Fallon about Sophia's capabilities.

She can see people's faces, she can process conversational data, emotional data, and use all of this to form relationships with people. Okay, so, I mean, she's basically a lot alive, is that what you're saying? Oh, yeah, yeah, she is basically a lot. Sophia and Jimmy even have a brief conversation about nacho cheese.

I like, I like nacho cheese. Nacho cheeses. Bew.

Gosh, dude, ew. I'm getting laughs. Yeah.

Maybe I should host the show. Okay, all right. Stay in your lane, girl.

When my guest today, Ross Picard, an AI researcher at MIT saw Sophia on the Tonight Show, she was immediately skeptical. I knew Hanson had worked a lot on making the robots look like they could express emotions like when she says not to cheese and she makes a disgust face. I think we all agree the rest of her, he says she's a social robot and

she's standing there like a stiff, poor, that she has a lot of room to improve.

And even calling her she is clearly a bit of a stretch, right? This is a machine with software and probably on that show some engineers in real time typing some things and driving it so that it does the right thing live for the camera in front of that audience. We don't know if there were engineers behind the scenes playing her like a puppet, but the CEO later said that she wasn't true AI. In the last episode of a dialogue, he said comes from a simple decision tree.

When you say X, it replies Y. And when Sophia appeared on CNBC, their own interview questions for her were heavily rewritten by her creators to elicit the proper responses. The CEO later walked back his comments about Sophia being basically alive by saying that Sophia was quote "alive" in the sense that a piece of sculpture becomes alive in the sculptor's eyes. The Facebook's director of AI said Sophia was complete BS and slammed the media for covering it, the CEO of Hanson Robotics said he never pretended Sophia was close to human-level intelligence.

If you're worried about robots taking over the world, Sophia being not as advertised is probably encouraging to you. And it should also encourage you that Ra's isn't too worried about that anytime soon. But there are things to worry about, or at least consider.

And that's what we talk about today. How AI is influencing our lives for good and for ill, and how we can lean into the good parts in ways that benefit both us and others. This is Beyond the Forum, a podcast from the Veritas Forum and PRX that explores the ideas that shape our lives.

This season we're talking about the intersection of science and God. I'm your host, Bethany Jenkins, and I run the media and content work at the Veritas Forum, a Christian nonprofit that hosts conversations that matter across different worldviews. I'm Rosalne Picard, I'm a professor at the MIT Media Lab.

I'm also a scientist, an inventor, a researcher, an author, a speaker. Mostly I'm someone who really loves to learn and to engage others in the journey with me. The MIT Media Lab was founded in 1985, but its focus isn't on what you may think.

Their main purpose isn't podcasts or social media. It's really funny that we call ourselves the Media Lab. The history, I'll be brief here, but it's when trying to start a new entity at a university, it's really hard to get approval.

It helps if there's nobody at the university who sees themselves as doing that area. At MIT, there was nobody doing media. It became a generic term that allowed a group of innovators that were in the architecture machine group at the time, looking at the convergence of computation and design and future of entertainment to come together

and put together a lab.

What it evolved into, however, has been very interdisciplinary, some would say anti-disciplinary or transdisciplinary. The idea is that there's a lot of interesting stuff happening at the boundaries of existing disciplines, and the MIT Media Lab pulls experts from various disciplines together to work on joint research and projects. We have a psychology major.

We have people with physics backgrounds, music backgrounds, mechanical engineering, design. There are a lot of advantages that happen when you free people up to think beyond what they think they're expected to do. A lot of the most important stuff that needs to happen is at the interface of people and technology.

We've really worked at that junction, putting people first, and then trying to understand how to shape the technology to make human lives better. The problems they focus on are big, like human health, human learning, how to be better stewards of the environment and more. They ask how technology can serve those goals and bring all kinds of people together.

For her part, Roz was trained to do electrical engineering and computer science. This is what she got both her masters and her doctorate in. In the 1990s, Roz taught the second ever machine learning class at MIT and coined the term "effective computing" to draw attention to the vital importance of emotion and decision making and perception, even for computers.

Today, a lot of her work at the Media Lab involves artificial intelligence, with an emphasis on the artificial, she says. At a Veritas form of edit Brown in 2019, Roz said that we need to be careful with our terms. Machine learning isn't real learning.

Artificial intelligence isn't real intelligence. The average person probably thinks, "Oh, it's learning like I did when I was in school, or it's learning like I do when I read something." It's not. It is this mathematical function approximation.

And it doesn't, at the end of all of this training with this mathematics, it doesn't have any consciousness. It doesn't have any awareness that it knows anything. It doesn't have any feeling of knowing.

It doesn't have any feeling of not knowing. In fact, it has no feelings. So it doesn't learn at all like we do.

Even at MIT, she says, they're returning to foundational questions about terminology. We're starting to have these conversations among ourselves of now that the public is very engaged in this. Maybe we made some mistakes in calling it learning.

Maybe we made some mistakes in calling it thinking. Maybe I shouldn't even be calling

the work I do, you know, affective computing, emotion, emotional intelligence. Because the machines don't feel they don't have this conscious awareness.

They aren't alive. They aren't like we are. And these function approximations, these things we're giving them to do, these input output behaviors can approximate and look like some of these things.

We do, but they're really fundamentally different. At the Veritas form event, the panelists spoke of two types of artificial intelligence. First was the sci-fi kind that we see in Hollywood, like Westworld or Blade Runner, which some refer to as artificial general intelligence, or AGI.

The stuff we see in the movies is a huge extrapolation of these simple demonstrations that exist in usually very narrow situations. And we can extrapolate both imaginatively and then we can try to technically figure out how to actually get there. And the technical part is really not done yet.

We have ideas. We can't say it can never be done. But right now we don't see any way to actually really build an AGI.

And people who are telling you, oh, it'll just another few years, I would not put my money on them. Ra's says the biggest impediment to building AGI is the utter uniqueness of humans. That there's something different and unique that people have that all the stuff we build, even if the outputs textually are identical or visually are identical, there's something missing there.

We don't know how to build it all. It's very humbling the more we learn about what are humans. The other type of AI the panelists talked about is machine learning, which is basically mathematical approximations.

And Ra says that it's this type of AI that's seeing the most success. And it's achieving the greatest performance results on certain tasks like computer vision tasks or reading, radiological exams or whatever. That's machine learning.

At the Veritas form event, Ra's and her fellow panelists, Michael Litman of Brown University, talked about machine learning and the importance of having appropriate, quote, objective functions for the machines they build. Machine learning is a sort of idea that we can, instead of writing software ourselves, we can just sort of define what good software is and let the computer figure out a way of behaving so that it matches that definition that we gave of good. And so, you know, so one of the reasons that Facebook is problematic is a lot of reasons that Facebook is problematic.

One of the reasons that it's so influential and has actually had maybe impacts, unforeseen impacts is they have a metric. They have an objective function just like in machine learning. The system is trying to do something.

It's given a scoring function by the programmers. And the scoring function that they gave was, well, we like it when people interact with the site. So the more they interact, we should show them the things that are going to cause them to interact.

Well, it turns out that the best way of getting people to interact is to outrage them. And I don't think that's what they were planning, but what they basically made is a function that is optimized by outrage. And so what the system as a whole, the AI and the machine learning behind the system, figured out is that if there's certain kinds of things that you can show people that are pretty much guaranteed to get a reaction and strong sharing.

It doesn't understand what outrage is, but it's like, great, I'm optimizing my objective function. So will it figure it out on its own? We have to give it the right objective function. Otherwise, what it'll figure out on its own is unlikely to be what we intended.

And that gets into the ethics and moral question. What are the objective functions we want to build in? Because if we don't build them in, they're going to optimize the wrong things. Or people are just going to optimize the, hey, what's cool, what gets me published, what's novel criteria, which, you know, can definitely get you on a professor track.

But we want to hit a higher bar at least at places like Brown and MIT where you're not just trying to do something cool and novel, but you're trying to do something good that improves the world also. So optimize two dimensions simultaneously. And that's a much harder problem.

Hard problems encircle the world of artificial intelligence. For example, machine consciousness. Sophia from Hanson Robotics was given citizenship by Saudi Arabia in 2017.

She was the first robot to be given legal personhood anywhere. The move was an attempt to promote Saudi Arabia as a place to develop AI. But what moral questions arise? What are the implications? When you flip the whole thing off, it's all gone.

Right. Well, one could say, well, we kill you and it's all gone to. That's unethical to kill me.

It's still ethical to flip off the machine. Right. Why is one unethical and one ethical? Well, because we think there's still something a lot more to human life than what we've built here.

And, you know, what is that? So this drive to build something like us is a sense to really try to understand what that is. Questions about data sets arise too. In the algorithm space, one of the key problems that has led to a lot of bias is that there is less data for certain people groups.

For example, among computer scientists who would run around and use a camera among everybody in the lab to get the initial faces to train the first face recognizers. Guess what? There weren't many women. There weren't many black people.

There weren't many older people. You know, so there will be biases in the data toward young white males. I might have been the only female face.

Some of these early data sets. And there were no, you know, no black faces, which is crazy. Right.

These point to larger injustices in our society. The hope for AI researchers, of course, is that we can rectify this in a way that data sets are less biased than people. I'm actually much more optimistic that we can build unbiased AI than that we can create unbiased employees, unbiased people.

People are free, right? We have free will. We are free to behave in ways that we're all proud of and unfortunately free to behave in ways that are, we wish people would never behave. So the algorithm, however, is not free.

It is trained to do certain things. And if we train it on unbiased data and test that it performs as desired on data with certain content and it behaves in an unbiased way, then I think we can be pretty confident within some margin of error that we can determine based on lots and lots of data that it will do what we wanted to do, what it was trained to do. Another limit to AI is that wisdom isn't the same thing as information.

We may gain a lot of information, but does that make us more wise? I am concerned that we're building technology without thinking first about as many of the possible unintended consequences. One of the creators of iOS, the operating system that's on iPhones, was that one of these gatherings of leading computer science technology developer people, where they bring people together to interact and exchange ideas and it's called food camp. And during the break, when everybody should be meeting all these super cool, amazing people, he looked around and he noticed that nobody was talking to anybody, that everybody was hunched over breathing shallowly over their phone.

And he said as he looked around, I feel badly that I invented this. And only then with these post reflections are people starting to say, I never intended to make that. I intended to make something sticky that got eyeballs, but I didn't think about the whole social ecosystem.

Ross says that AI is at its best, not when it's replacing humans, but when it's working with them. One of her innovations is a great example of this. She and her team at the MIT Media Lab were trying to build an AI that could help a non-speaking individual to communicate when they were getting stressed out or about to have a meltdown so they could self-regulate.

It was a wearable, like a smartwatch, and the idea is that it could benefit both the wear and the people around them. But things took an unexpected turn. Right before the end of the semester, one winter, one of my undergrads knocked on my door and he said, "Professor Picard, could I please borrow one of those sensors? My little brother has autism.

He can't talk. And I want to see what's stressing him out." And I said, "Yeah, sure. In fact, don't just take one, take two because they break.

They were all hand-built with lots of wires hanging out." And I said, "Do you need a soldering iron?" He said, "Nope, I've got a soldering iron. I'm looking great. MIT student can fix it." So the student took two sensors home and Ross monitored the data remotely on her laptop.

And the first day's data looked pretty flat. The signal we were monitoring, this kind of a sweat response on the wrist that tends to go up when people are stressed, like for an exam or driving in Boston. Or we've seen it with different kinds of social interaction that can be stressful.

This signal is pretty flat for this kid. And the next day was flat. In fact, both wrists were flat.

And the next day, and I'm kind of yawning thinking, "I hope the sensor's working. I'm zooming in." And there were little blips. I go to the next day and my jaw drops.

One of the wrist signals was so high that I thought the sensor must be broken. We have never seen such a big response, even stressing people out in Boston driving and qualifying exams and loud noises popping in their ears and other obnoxious things that we do in our lab to test if our sensors are working. The other wrist was flat.

Now, this was really puzzling. How could you be stressed on one side of your body and not the other? Right? We thought we should get a general arousal response. So I'm zooming in on the data trying to debug this.

She eventually got stuck. So stuck that she tried something she normally wouldn't do. She made a phone call.

I picked up the phone call to student at home on vacation. Hi. How was your Christmas? How's your little brother? Hey, any idea what happened to him? And I gave the exact time and date and the data.

And he said, "I don't know. I'll check the diary." And I quick prayer, right? Like silently, like, what are the odds that an MIT student has kept a diary and written down this exact moment on their vacation? Well, he comes back and he has the exact moment written down. And he says, "That was right before my little brother had a grand-mall seizure."

The grand-mall seizure is the type of seizure most people picture when they think about seizures.

And they're commonly associated with epilepsy. They usually last about one to three minutes and in some cases can be life-threatening. When Ra's learned that her sensor picked up the grand-mall seizure, she knew that she had to make another call.

I learned that another one of my students' dads is the head of epilepsy surgery over at Boston Children's Hospital. And so I screw up my courage and I call Dr. Joe Madsen on the phone. Hi, Dr. Madsen.

My name is Rosalind Picard. Is it possible that somebody could have a -- I wanted to use the technical term here -- "Huge sympathetic nervous system surge 20 minutes before a seizure," which is what it looked like in the data at the time. And he said, "Probably not." And I paused and he went on and he said, "But you know, it's interesting.

We've had patients whose hair stands on one arm 20 minutes before a seizure." And I was on one arm. Well, so then I told him the whole story and showed the data. He got super interested.

Together, they built a bunch more devices, got them safety certified and ran tests. He had 90 families coming in for monitoring around the clock where the children were going to have EEG on their head, ECG on their chest, and now we're adding EDA, electrodermal activity on the wrist. And they would be monitored for seizures.

We found 100% of the grand mal seizures had a large response with our wrist sensors, significantly more multiple standard deviations above the pre-susier period. Across the board, we did not find that it was 15, 20 minutes before, usually when it was precisely synchronized with the brain activity, it started exactly at the same time as the brain activity. So we didn't wind up with a forecaster, but we wound up with a good opportunity to build a detector.

So they built it, and the company that Ros co-founded in PADICA is the manufacturer. So we built a machine learning system, lots of data, lots of examples that teach it how to monitor multiple signals and build an accurate detector with high probability of detecting an event, low probability of a false alarm. And then later that was commercialized by Empatica, the Italian word for empathetic.

And now it's the first FDA cleared smartwatch on the market. And it's not running, you know, your Uber app or draw cute little butterflies or all that stuff that could cause an update to go off at 2am in the morning when you might be having a seizure. But it's a focused AI, it's number one priority is making sure that it detects those seizures and summons somebody to come and check on you at that moment.

Because we've now learned that there are more deaths in the US every year from these

then house fires and sudden death syndrome. And nobody talks about this, but most of those deaths appear to be preventable. If somebody comes quickly when you're having a seizure and stays with you in those moments afterwards and makes sure that you don't stop breathing.

Upplies first aid, reposition, so you make sure your airway is open and stimulates you. That can lead to much better outcomes.

[Music] In her TED talk, which now has over 2 million views, Raz tells the story of an email she received from a mom whose daughter Natasha was wearing one of the empathetica sensors when she was wearing a mask.

She was in the shower when her phone on the counter went off and it said that Natasha might need help. So she immediately jumped out of the shower, ran to Natasha's bedroom and found her face down in bed, blue and not breathing. She flipped Natasha over and she took a breath and then another.

Natasha turned pink and was fine. Raz's first response was, "Oh no, it's not perfect. The Bluetooth could break, the battery could die.

All of these things could go wrong. Don't rely on this." But Natasha's mom replied, "It's okay. I know no technology is perfect.

None of us can always be there all the time. But this, this device plus AI enabled me to get there in time to save my daughter's life." And it's this point in particular that Raz wants to be very clear about. It's not the AI saving the life.

It's the AI summoning a person who comes and repositions, you stimulate, you make sure you're okay.

[Music] Hi all, this is Carly Regal, the assistant producer of Beyond the Forum. If you're loving the podcast so far, we want to invite you to continue engaging in these important conversations by signing up for our newsletter.

Each month, you'll receive thoughtful content about the ideas that shape our lives, updates from our student and faculty partners and other Veritas news and events. You can sign up today by visiting veritas.org. Thanks for tuning in and enjoy the rest of the show.

[Music] AI is here to stay, and it's only increasingly becoming a part of our everyday lives, from social media to facial recognition to buying online.

And because it's so ubiquitous, it's worth asking, what is the goal of AI? Is it to make our lives better, more efficient? From the developer's perspective, it's complicated. I think more often than you want to hear it, the goal is going to be, "Oh, I'm just trying to get

my bachelor's, my master's. I'm just trying to get published.

I'm just trying to get a good job." They're not thinking about the big goal, and they need to be asked that. Otherwise, they just sign up for whatever company pays them well and find out later that all their work is going, for example, to just sell people more of something, right? To just get more clicks, to just get more people to click on this ad, to make ads more appealing. There is a lot of money and work that's going into just whatever pays the bills at the end of the day, and not enough of these conversations about what is the goal of AI.

But Roz thinks there's something else going on, too. We're not just longing for better systems. We're looking for some sort of super intelligence that can make our lives better.

I think you've ever gotten information from a computer, and maybe you believed it more than information some person gave you. There's something about when the computer or some big measurement system in the doctor's office gives you something, doctors talk about this. Why does the patient believe what the little printout says, and they don't believe me? We seem to accord some more credibility sometimes to it, maybe because it's objective, but it can be objective and completely wrong.

And yet people still believe it. And some AI developers, Roz, says, are even longing for immortality. I work with a lot of people in the Media Lab at MIT and AI and Design and Technology who are interested in being immortal.

In the most realistic way, they just want to have such an online presence that they'll never be forgotten, right? All their works will live forever. But some of them truly believe there's a life beyond this one and that they can be cryogenically frozen, and it's all material, and someone will figure out how to re-innervate, like the Frankenstein story, their brains, or figure out how to build them in an AI and give them a new body and a new mind that reflects some of who they currently are. At the Veritas Forum event that we featured on our last episode, John Lennox connects this longing in AI with God.

They are striving to produce a godlike human with super intelligence. Many of them are atheists. They don't realize there is a man who is God, has already been on this planet.

But he's not artificial intelligence, ladies and gentlemen. He's real intelligence, the intelligence of God made incarnate. So what interests me about this whole thing is what people are moving towards is a parody of a scenario that is embedded in Scripture.

Raz wasn't always a Christian. For the first part of her life, she was an atheist. She thought science was on one side and Christianity was on the other, and she chose science.

But then she met some Christians who were reasonable and intelligent, and they

encouraged her to learn more about Christianity. So Raz started to read the Bible. And as I was reading that against my desires, I started to change my mind about some things.

And then I thought, "Oh gosh, okay, if this book is influencing me to change my mind toward Christianity or toward belief in God, maybe I should study other world religions." And as I started learning more and more about different world religions, meeting people from those religions and going to temples and mosques and others, I started to realize that not only did I have a lot to learn, but I was on a journey that was starting to make me not only believe in God even more, but as I got dragged off to some Christian churches, which I resisted in the beginning, and found somewhere I could ask questions. And I started to realize that the religion was not at all what I thought it was, and that there were some really interesting and very attractive elements that were very historically verified. And as I learned about that, I changed my viewpoint gradually from an atheist to an agnostic to a theist, to somebody who actually believed that the historical Jesus in the New Testaments, what's written about him was true.

It sounds a little wacky to those who may not come from that background. It was not an easy process. And actually, the real reason I'm here right now, spending time talking about something like this as opposed to just my research, is because it has made a huge difference in my life.

And part of the Christian faith is that there's a gift for everybody in the world, whether you're raised Christian or Hindu or Muslim or Buddhist or atheist or any of a long list of backgrounds. And today it is my source of strength, an amazing source of peace and joy and wisdom. And as I think when we build machines and build computers with affective abilities and robots, I often think of the analogy of one who is very wise giving us instructions and giving us guidance and being there when we don't know what to do.

So I find that still is powerful in my work today. The future may not be TV-ready robots like Sophia taking over the world, at least not anytime soon. But we still have huge amounts of technology in our lives.

In our pockets, in our grocery stores, in our classrooms, in our workplaces, and while you and I may not be developing more artificial intelligence, we still have the opportunity to think about how we can grow as users. Each of us can ask bigger questions about our tech use. What is my goal? When I open up Instagram or scroll my favorite news site, what do I want? And once you think about what you want, consider what Ra said about the sensor she created.

I want to be careful. It's not the AI saving the life. It's the AI summoning a person who comes and repositions you, stimulates you, make sure you're okay.

How can you use technology and AI, not to replace people in your life, but to help them? Next week, we continue our exploration of science and God with Dr. Cullen Bowie,

another researcher from MIT. He's an engineering professor there, but his path to MIT was far from direct. Join us as we talk about his professional journey and what you can learn about your own vocational discernment.

You won't want to miss it.

[Music] Hi again, this is Assistant Producer Carly Riegel. To end this episode, we at Beyond the Forum want to take time to say thanks to all the folks who helped us get the show together.

Our first thanks goes to our guest, Dr. Roslyn Picard. Thank you for your literally life saving work and for reminding us to think about the bigger implications of our tech use. We also want to thank our production team at PRX.

That's Jocelyn Gonzales, Genevieve Sponseler, Morgan Flannery, and Jason Saldana. And of course, we want to thank the students who host and plan these forum events as well as the John Templeton Foundation and all of our donors for their generous support of our conversations. Alright, that's all for this episode.

Thanks for listening to Beyond the Forum.

[Music]

(buzzing)

[buzzing]