Bingecast: George Montañez on Intelligence and the Turing Test (https://mindmatters.ai/podcast/ep88)

Austin Egbert:

Greetings I'm Austin Egbert, director of Mind Matters News. Welcome to another Binge Cast where multiple episodes are combined into a single program. This week's Binge Cast features an interview with George Montañez about the famous Turing test, artificial intelligence, and the origins of natural intelligence. Enjoy.

Introduction:

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

Robert J. Marks:

Alan Turing wrote a seminal paper in 1950, entitled Computing Machinery and Intelligence. There, he asked the intriguing question can machines think? That was in 1950. The details of his question, according to our guest today, is too vague to admit an exact answer. Our guest today is Dr. George Montañez. He's the Iris and Howard Critchell assistant professor of computer science at Harvey Mudd College. Dr. George Montañez earned his PhD in machine learning from Carnegie Mellon University. Before Harvey Mudd, he worked on AI software for Microsoft. At the international joint conference on neural networks, Dr. Montañez was co-recipient of the best poster award and the INNS/Intel best student paper award for the LICORS Cabinet: Nonparametric Light Cone Methods for Spatio-Temporal Modeling. We cannot let the LICORS cabinet pass without content. What was the LICORS cabinet? It wasn't a cabinet full of Jim Beam and other stuff?

George Montañez:

No, it was not, but it was a cabinet that contained mixed liquors, moonshine, and 100 proof, LICORS being an acronym. L-I-C-O-R-S. And I don't remember what the acronym is for at this point.

Robert J. Marks:

You don't remember what it said? It's always important to have a good acronym, even if you don't remember what it was. He previously was a co-recipient of the best paper award at CIKM for a paper on cross device search. Most recently, he was awarded the best student paper award at the IEEE international Conference on Systems, Man, and Cybernetics. Today, we'll be discussing with Dr. George Montañez his paper, entitled Detecting Intelligence: the Turing test and Other Design Detection Methodologies. We'll provide a link to the paper on the podcast notes for those that are interested. Dr. Montañez, welcome.

George Montañez:

Thank you. It's a pleasure to be here.

Robert J. Marks:

For those who are not familiar, can you summarize what the Turing test was, and the Turing test was actually done in that paper that I referenced? Is that right?

George Montañez:

Yeah. So, there was actually three versions of the Turing test that Alan Turing gave, but at a high level, what the Turing test is, is a test for trying to distinguish whether or not a computer could reliably essentially fool a human judge into thinking that it's a human. So, a human would interrogate a machine via a system of text messages if you would, and the machine would give responses, and the goal is to see whether or not the machine could reliably imitate a human being. So, this was a variant of a popular party game called the imitation game where you had two people in different rooms. One of them, a man, one of them, a woman, and party guests would try to ask questions of these various parties and guess which one was the man, and which one was the woman. So, essentially, both would pretend to be a woman and you would try to differentiate which one was actually a woman based on the responses given.

Robert J. Marks:

Now, in the party game, how did the man change his voice? Or did the woman change their voice?

George Montañez:

That is, an excellent point. So, because they were in different rooms, what would happen was the person who was running the game would pass written messages. So, this was to remove side information. In Turing's version of the test, he was very adamant that all the interactions had to be done via just a message passing. He didn't want to bias the results. For example, if the answer came back, yes, I am a human, he didn't want that to color the judge's expectations say, Oh, that's a machine, obviously. So, he wanted these to be based only on the responses. So, in the imitation game, it's done by slips of paper with things written on them. In the Turing test, it's based on text messages on a computer screen.

Robert J. Marks:

It always struck me as interesting that the converse of the game is very easy. It's very easy to determine if who you're talking to is a computer. You just ask them to compute the square root of 30 or something, because a human would take a while to get the square root of 30.

George Montañez:

Depends on the human. We have a professor here, Art Benjamin, who's the math magician, and you can ask him what the roots are of very large numbers and he can give you an answer. Not as quick as a computer, but pretty quick. Quicker than most humans would be able to.

Robert J. Marks:

So, can he give the twelfth root of two?

George Montañez:

Oh, I don't know if he can do multiple. I think it's mostly doing very large numbers. So, you say what is the root of some number with seven digits and he'll be able to tell you.

Robert J. Marks:

We're familiar with a bunch of chatbots. Chatbots come in different flavors. There's the ones you can actually hold conversations with. There's the single exchange chatbots, like Alexa and things of that sort. Do you think that the chatbots that actually perform the exchange conversation are passing the Turing test?

George Montañez:

No, because I don't think that people are actually thinking that they're speaking with a human, unless they're really small children. So, my children happen to think that Google is a person. Yes. Yeah. So, they'll ask Google, "Google, what is your favorite cartoon? Google, what is your favorite color?" And I have to explain to them that Google is not a person. Google is a system. But most adults would realize that they're not speaking with a human when they're talking to Alexa, I would think.

Robert J. Marks:

Well, I think some of the chat-bots were not actually written to pass the Turing test. I think that they're at a lower level. Certainly, Alexa and Google Home isn't. So, let me ask you this, George. RPI Selmer Bringsjord, over a decade ago, said something very interesting about the Turing test. He said, "The progress toward Turing's dream is being made. It's coming only on the strength of clever, but shallow trickery. For example, human creators of artificial intelligence that compete in present day versions of the Turing test know all too well that they have merely tried to fool the people who interact with their agents into believing these agents really have minds." Do you concur with that?

George Montañez:

I think so. So, I think if you look at maybe the best known example of this being Eugene Goostman, the creators of that piece of software actually made the personality of that software be a 13-year-old Ukrainian boy, because they felt that any mistakes that were made either in factual knowledge or maybe in grammar would be more easily forgiven if the person thought they were chatting with an adolescent versus speaking with an adult. So, this was a specific design decision that was made in order to fool people and give their system a leg up. So, I think in some sense, yes. That's occurring.

Robert J. Marks:

The other thing about Eugene Goostman, I think he was Ukrainian. Did he talk in a Ukrainian accent?

George Montañez:

This is all text messages. Textual messages. So, the systems don't actually have speech vocalization. It's just based on things that are typed.

Robert J. Marks:

Yes. Well, I'm just wondering, if he responded incorrectly, could they chalk it up to the fact that "Gosh, it was Ukrainian. English isn't his first language?"

George Montañez:

That's exactly right. So, for the messages that they read, they could say, "Oh, if there were grammatical mistakes or maybe he didn't understand something that was asked it's because English is a second language versus because I'm speaking with a machine."

Robert J. Marks:

Yeah, yeah. You have a couple of quotes from Eugene Goostman in your paper. One is where Eugene Goostman was asked, "Why do birds suddenly appear?" Isn't that a line from a Carpenter song?

George Montañez:

I believe it is. Yeah.

Robert J. Marks:

And Eugene required, "Just because two plus two is five, by the way, what's your occupation? I mean, could you tell me about your work?" It seems that he's kind of deflecting the problem by asking questions back to the person that's querying them.

George Montañez:

Yes.

Robert J. Marks:

And that's an easy way to deflect things. If I was interviewing somebody and they started asking me questions and I'm trying to interview them and they started asking me questions, I would dismiss them because I want to learn about them. I don't want them learning about me. It's the other way around. Okay. Here's another one. The judge says, "is it okay that I get sick of sick people? How is your stomach feeling today? Is it upset, maybe?" Eugene responded, "I think you can't smile at all. I bet you work for a funeral agency." I don't know. To me, that doesn't seem very respondent. I wouldn't say that that was an intelligent give and take.

George Montañez:

I wouldn't either. So, I put these quotes in the paper to show instances of where the system's responses are actually seemingly independent of the question asked. So, I think this is clever on the part of the programmers in that they built in fall back mechanisms when they couldn't actually answer a question well. They gave just a random response, which again, if you're speaking with a 13-year-old, 13-year-olds do this from time to time, but this isn't an answer that's actually directly answering the question that was asked. So, I would agree with that assessment. Yeah.

Robert J. Marks:

I don't think somebody that was really educated would think that Eugene Goostman, based on this short excerpt that we just read, would really think that he had any intelligence at all.

George Montañez:

Based on that one, no, but you have to figure that in order for the claims of Eugene Goostman passing the Turing test, he had to have fooled greater than 30% of the people. So, it did something right. Whatever the system was designed to do, it beat that criteria that was set up for it.

Robert J. Marks:

George, what do you think the importance of passing the Turing test is to the establishment of mainline artificial intelligence? Is it an important test? Was it as important as when Turing proposed it back in 1950?

George Montañez:

So, I forget the name of the principle, but there's a principle that once you start trying to optimize for a metric, it ceases to be a good metric. What is the name of that principle?

Robert J. Marks:

Goodhart's Law.

George Montañez:

There you go. So, I think this is a perfect example of Goodhart's law in the sense that now that people are specifically trying to beat the Turing test, the thing that it's supposed to measure is no longer a valid test for that because people are trying to build systems just for the point of passing the Turing test versus building systems that actually have artificial general intelligence. So, it would be nice if people were just working towards the goal of developing artificially general intelligent systems, and then those systems happen to pass the Turing test. Then, it would be a good metric, but because we're essentially trying to optimize towards that metric, it ceases to be really meaningful in that sense.

Robert J. Marks:

Yeah. Goodhart's law, I believe, can be stated that any time you establish a metric as a measure of something you want to do it no longer becomes an effective measure for what you want to do. There's a follow-up and that's called Campbell's Law.

George Montañez:

Okay. I haven't heard of this.

Robert J. Marks:

Campbell's law is that the Goodhart's law will often lead to trickery and deception.

George Montañez:

Oh, there you go.

Robert J. Marks:

So, that's exactly what is happening here.

George Montañez:

Yeah.

Robert J. Marks:

So, the Turing test has been proposed to support artificial intelligence. Again, as we mentioned previously, artificial intelligence and intelligent design, both contain the word intelligence.

George Montañez:

Well, variants of it.

Robert J. Marks:

Variants of it, yes. You talk about this in your paper. What is intelligent design?

George Montañez:

So, intelligent design is the theory that there are certain features of nature, either the cosmos or biology, that show empirical evidence of having been caused by an intelligent agent rather than just undirected causes. I like to differentiate in two ways that intelligent design is thought of. So, there's the causal theory of intelligent design, which is that between unintelligent processes and mechanisms and intelligent agents, there are causal differences in their abilities and that these differences can be picked up on in empirical, observable ways, right? So, that's the causal theory. Then, there's the historical theory of intelligent design, which is that when we actually look at nature, when we look at the cosmos, we see that there are signs of this intelligent having acted in history. So, this paper is not concerned with the latter form of intelligent design. It's more focused on the first form, of whether or not there are discernible differences between the activities of intelligent agents and unintelligent unguided systems.

Robert J. Marks:

There seems to be, at least to my knowledge, at least three different flavors of intelligent design. One of course is from the Christian perspective, who believes that God is the creator. Another one is directed panspermia where the complexity that we see on earth is too complex to understand or to evolve, and therefore, must've been planted here by aliens. Panspermia's just the general way of doing that. In other words, the life could have come here on a meteor.

George Montañez:

Sure.

Robert J. Marks:

But directed panspermia means there was an alien life form that actually came and planted it here. The third one is actually, I would propose due to Elon Musk. He says, things are so complicated here that, boy, we must all be simulations. That we can't have been generated by chance, so therefore we must be simulations by some super being. We've talked about the Turing test and artificial intelligence and about intelligence design. Your paper makes a connection. Could you elaborate on that?

George Montañez:

Yeah. So, I came to the realization when thinking abstractly about what the Turing test is doing. So, essentially, the Turing test is looking at artifacts. In the case of the Turing test, these are words on a computer screen, and you have a person who, based on these artifacts, is trying to determine whether or not the cause of them was an intelligent agent or was some sort of unintelligent process. The Turing test presupposes that based on the artifacts and the observations alone, you should be able to make that determination. So, if the Turing test is passed, we want to say that the system's actually intelligent. So, for the Turing test to be meaningful, you need for there to be some sort of difference in these causal powers and for these differences to be empirically observable, empirically verifiable.

George Montañez:

I realized that if you can do that, this is a major core assumption of not just the Turing test, but also of intelligent design detection, which is that we're supposed to be able to look at artifacts and be able to, from those artifacts, determine whether or not the cause was an intelligent agent or some unintelligent process. So, it's this interesting tension because a lot of materialist reductionists would say, especially if there are believers in strong AI, that perhaps the Turing test is a scientifically valid test of whether there was an intelligent cause or not. But when this is applied to things like biology, there can be no scientifically valid test for whether or not the cause of a system was intelligence.

George Montañez:

So, the paper essentially draws out the parallels between these two methodologies and shows that they rise and fall together. So, you can't have your cake and eat it too, essentially. So, you're forced with a choice. You can say that the Turing test is not a scientifically valid test and neither is intelligent design theory, or if you say that the Turing test is scientifically valid, if its assumptions are scientifically defensible, then you cannot a priori rule out methods of design detection such as those used by intelligent design theorists.

Robert J. Marks:

Fascinating. I think the SETI, which is the search for extraterrestrial intelligence, isn't that kind of a Turing test in a way? A variation of it, I would say.

George Montañez:

Yeah, in the sense that SETI is looking at artifacts which are radio signals primarily, and trying to, from those observations alone, determine whether the cause was an intelligence or was something else that was unintelligent. So, to the degree that SETI could actually work in the real world, this is based on that same set of core assumptions.

Robert J. Marks:

Excellent. Well, thank you, Dr. Montañez. By the way, we've been talking about Dr. Montañez's paper, Detecting Intelligence, and we're going to provide a link to it on the podcast notes so those more interested can view it and read all of the nitty gritty detail. So, thank you again, Dr. Montañez.

Robert J. Marks:

We've been talking to Dr. George Montañez. He's the Iris and Howard Critchell assistant professor of computer science at Harvey Mudd College and we really appreciate his participation at Mind Matters News. So, until next time, be of good cheer.

Conclusion:

This has been Mind Matters News with your host Robert J. Marks. Explore more at mindmatters.ai. That's mindmatters.ai. Mind Matters News is directed and edited by Austin Egbert. The opinions expressed on this program are solely those of the speakers.