

Bingecast: Selmer Bringsjord on the Lovelace Test (<https://mindmatters.ai/podcast/ep101>)

Austin Egbert:

Greetings. I'm Austin Egbert, director of Mind Matters News. You're listening to another bingecast where multiple episodes are combined into a single program. This week, we talk with Dr. Selmer Bringsjord about the limitations of the Turing Test, AI's ability to be creative, and whether human minds are computable. Enjoy.

Announcer:

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

Robert J. Marks:

Greetings. The Turing Test, developed by Alan Turing in 1950, is a test of a machine's ability to exhibit intelligent behavior indistinguishable from a human. Many, including our guest today, think that Turing's proposal for intelligence, especially creativity, has proven inadequate.

Robert J. Marks:

Over 15 years ago, our guest, Dr. Selmer Bringsjord, proposed the Lovelace Test as a substitute for the Turing Test. Dr. Bringsjord is a professor at Rensselaer Polytechnic. He holds positions and appointments in three different departments: Professor of Cognitive Science, Professor of Computer Science, a Professor of Management and Technology. And he's the director of the Rensselaer AI Reasoning Lab. And that acronym is RAIR, R-A-I-R.

Robert J. Marks:

On a personal note, Dr. Bringsjord ranks with me, along with people like John Searle, and David Gelernter, and even Roger Penrose, in teaching me both the potential and limitations of computers and AI.

Robert J. Marks:

And I'm just honored to have you here today, Selmer. Thank you. Thank you for joining us.

Selmer Bringsjord:

Well, thank you for having me. I'm honored to be here.

Robert J. Marks:

Thanks. Thanks much. I wanted to talk about your Lovelace Test. Now, I think in order to talk about the Lovelace Test, we first have to go back to the Turing Test. Could you elaborate on that and build up towards what your Lovelace Test is?

Selmer Bringsjord:

Sure. We really only need two steps to do that, and I think it is wise to do it. One, we have to define the Turing Test. Most of your listeners will be familiar with it, but just in case. Lock a computer or a robot, an AI, what have you, in one room. Seal it off. Lock a human being in another and seal it off. Go get a judge who doesn't know which entity is in which room, but allow the judge to communicate linguistically, as Turing said, by teletype. But we could say in this day and age, you could use the full arsenal of purely textual social media. Email would be ideal. Send emails back and forth, ask whatever questions you would like to try to find out which room houses which of these two entities. And Turing said, "If you strike out as judge, then we will have reached the point in time when we have a thinking machine."

Selmer Bringsjord:

And there are contaminants out there for monkeying with and diluting Turing's original rules for the game. He said that this had to happen for at least an indeterminate or a long period of time. Clearly by context, you can't lay down an artificial period of time. If it was two seconds, meaningless, 30 seconds, probably a machine is going to, if not indistinguishable, going to have a chance. And no prohibitions on the nature of the judge, in terms, for example, of their background, their education, their expertise. Perverted versions of this test that have been played, generally restrict the judges to being lay people, not being experts in AI.

Selmer Bringsjord:

So, that's basically what the Turing Test is. As you already, I think, pointed out, he announced and defended this test in a 1950 paper, and he considered a series of objections. And one of those objections is the Lovelace objection, so this is step two for getting to the background of what we need to discuss the topics. And what Lovelace said was, "Well, wait a minute. I'm familiar with the nature of a machine." She was presumably, or Turing, maybe charitably on her behalf, was referring to the Analytical Engine of Babbage.

Robert J. Marks:

Could you fill in a little bit of the background about Ada Lovelace? She seemed to be a remarkable lady.

Selmer Bringsjord:

Yes. And these remarkable attributes continue to be, in many quarters, venerated, and in other quarters, questions and give rise to all kinds of debate, but, basically, I think she was the legitimate child of Lord Byron of poetical fame, and Lady Byron during the time when lords and ladies held great sway. She showed remarkable, this is incontestable, mathematical aptitude as a young age. And I believe that she met Babbage in school around 1830, maybe 1835. Babbage, of course, also a great amount of aptitude for such things.

Selmer Bringsjord:

And she, most relevant for our purposes today, got involved with Babbage and his Analytical Engine. And I believe this is where the controversy enters the picture, certainly wrote an account of what a process... It was never built, the machine, but what it would be like to run a process in the machine. She did that about 1840, maybe 1842. And this is canonical now, the publications out there. You can look at what she's done. So, she's regarded by many to be the first programmer.

Selmer Bringsjord:

And the controversy arises from two facts. She might not have really actually written the program, but rather just done a hand trace of the Analytical Engine. I take no position. And at any rate, even if she did write a genuine program, she might not have been first. But from my perspective and the perspective of the Lovelace Test, that controversy is entirely irrelevant and nicely so because no one can dispute the fact that Ada Lovelace had a deep, accurate understanding of what computation is, at least mechanical standard computation. She knew that and the publication of '42 certainly confirms that.

Selmer Bringsjord:

So when she complains about a machine being inadequate, even if it happens to pass the Turing Test... And that's what she did. We can talk more about that. She really knew what she was talking about. So, again, I don't care about the programming pedigree and the claim to being the originator of programming. I care about what she clearly thought was a big missing, and perhaps eternally, missing piece in a computing machine that it could not be creative.

Robert J. Marks:

I believe she said something like, "Computers cannot be creative. They only do what they're told to do."

Selmer Bringsjord:

Exactly. Exactly.

Robert J. Marks:

That was her basic objection.

Selmer Bringsjord:

Right.

Robert J. Marks:

And Turing was apparently aware of this and addressed this objection?

Selmer Bringsjord:

He did address it in his 1950 paper. Well, charitably put, he addressed it. I don't think he really dealt with the objection. That may be true of some of the other objections as well. I think the consciousness objection, which basically says, "Look, give me a break. The machine to be linguistically indistinguishable in a teletype interaction doesn't have to be conscious." And consciousness is a big deal, at least for us humans. It seems to be paired inseparably with intelligence. I think consciousness is part of the equation when you're talking about at least deep, robust creativity as well.

Robert J. Marks:

Yes, yes. Okay. I'm going to quote to you from your paper, and it's a quote I've shared with people. This is from your Lovelace paper from about 15 years ago or so. You say, and I love this quote, "Though progress towards Turing's dream is being made, it's coming only on the strength of clever, but shallow, trickery. For example, the human creators of artificial agents that compete in present day versions of the Turing Test know all too well that they have merely tried to fool those people who interact with the agents into believing that these agents really have minds."

Robert J. Marks:

I think that that's a wonderful quote and really shows the inadequacy of the Turing Test. Very succinctly put.

Robert J. Marks:

You familiar with Eugene Goostman at all? He was, supposedly, a AI agent that passed the Turing Test a few years ago.

Selmer Bringsjord:

Not intimately familiar with it, but I certainly followed it from a media point of view. Yeah.

Robert J. Marks:

Your quote comes to mind because Eugene Goostman was a teenager from Ukraine. So if you talk to him, and he didn't understand what you were saying, "Well, come on. The kid's Ukrainian, doesn't understand English. And if he gives a stupid answer, "Hey, the kid's only a teenager." So he passed the Turing Test, and that just reminded me of your quote about clever but shallow trickery.

Robert J. Marks:

So, let's get on to the Lovelace Test, a more robust test for intelligence in computers. How do we test whether a computer or artificial intelligence is creative or not? That's in your Lovelace Test. Give it a little description of that if you would.

Selmer Bringsjord:

Sure. I don't have perfect recall of the paper. In fact, it seems that you are much more familiar with it at this point than me. I know the basic ideas, for sure, and I still stand by them, so it's, to some degree, facetious. But the idea is really quite simple. It's to take Lovelace at her word. It would plausibly be a group of people who have built this amazing system that's supposed to be human level, conversationally. And if it's conversational, and it can write a sonnet, which was a specific ability Turing considered, it should be able to do things that we would regard as creative.

Selmer Bringsjord:

Now, if it's already got sonnet typed in by the programmer, and when an appropriate context is detected in a conversation along with input, it spits back the sonnet. That's not going to cut it.

Selmer Bringsjord:

And most importantly, and this gets at the nature of the test, the developers of the system would know upon seeing the sonnet emerge, "Oh, great. Our guinea pig over here wasn't exactly sure why the context was there and that it triggered the sonnet being spit out, but we know what's going on. So, in a case where the developers, literally, have absolutely no idea how sustained production through time of artifacts that irresistibly imply, in our minds, that this is a highly creative agent, that is what passing the test amounts to.

Selmer Bringsjord:

Now, this does assume that you can't take a developer who's been on the fringes of the team, so it's really an ideal observer view. And if I had to do this over again, or if I had to update it, and I do think it

needs to be updated in light of machine learning/deep learning these days, and maybe we can get to that in a minute or two.

Selmer Bringsjord:

If I had to update it, I would appeal to an ideal observer more rigorously than I did. There should be some formal or mathematical way to say, "Look, every iota of what's going on here may not be assimilated by any one developer, or even the team of developers, that built this remarkable AI."

Selmer Bringsjord:

But, hypothetically, let's assume that the ideal observer has total command of everything that's going into this system. And it sits back, and now lets it start its operation. Lets it work and generate these artifacts that are going to, by hypothesis, imply, at least in the observers, that there's something remarkably creative going on. If that ideal observer is in the dark, now we know we're dealing with something truly astonishing.

Selmer Bringsjord:

It is, of course, roughly analogous to what we have in every computational neuroscience lab on the face of our planet, no matter what they might say in public, if you get any of those folks in a room, saying, "How actually did John Updike come up with that page there in the Witches of Eastwick? We're a little bit confused about it. Can you please explain that?" They have, literally, no idea, zero, how that passage comes out of the brain. So the test says when we reach that point, we better start paying attention to these machines, at least if we want to measure them against the human mind, which was exactly what Turing was all about in his 1950 paper.

Robert J. Marks:

Excellent. I think some people that would argue against that would point to probably the match between Alpha Go and the world champion. I believe his name was Sedol. And there was an incredible move at one point where Alpha Go made a move which was totally contrary to convention, and some people point to that and they say that this move was creative. Any comment on that?

Selmer Bringsjord:

Yeah. Yeah. Oh, I have many comments on that. I think that's easily refuted or that attitude is easily shown to be irrational, the attitude that this is true creativity or the source is true creativity. If the problem is such that it's winnable by a search strategy, that humans are able to lay out, that is certainly the case with respect to chess. There isn't enough time in the universe to run this algorithm.

Selmer Bringsjord:

Chess, checkers, Go, are all in the same complexity class. Little known secret, they're all in EXPTIME. Your listeners might not be familiar with that. They're really equally hard, but if the machine running this search algorithm is able to work fast enough, no human's going to be able to comprehend how it does anything. Forget about one move. It's going to be invincible, and it's going to do things time after time after time that are completely mysterious and efficacious as well. And yet it's running a simple algorithm.

Selmer Bringsjord:

So, the counter example is... Look, it's not a subjective piece of behavior at this level. We have to start with some kind of task that we don't have a garden variety search algorithm for.

Selmer Bringsjord:

In the case of chess, we've found ways to take search and make it more efficient, alpha-beta minimax search. And the case of Go, it's a bit different because machine learning plays a significant role, and the machine is approximating the function by running a gradual process that we also can't follow as it goes along.

Selmer Bringsjord:

But again, the bottom line is we know the task is intrinsically easy, and we know the algorithm does exist. However, Alpha Go, or any other such machine arrives at its ability to play really good Go. We know that if there was enough time and energy available, this is absolutely not a difficult game because we already have perfection. We can define it, so that really doesn't count.

Selmer Bringsjord:

A novel, or even a short story... A novel would be rather a different affair. I just listened in audio book form to Flannery O'Connor's *The Dead Bear It Away* or *The Violent Bear It Away*, and what struck me about this novel is that, over and over again, she makes you feel what a character's feeling with just a few sentences. So, her consciousness is present and exploited, and then she comes up with these amazing sentences, two or three in a row that just get the job done.

Selmer Bringsjord:

We can't say that the production of sentences like that is even in theory possible to carry out by following an algorithm, because we don't have the algorithm. We don't have what it would take to do it. So, returning to your original... And I think it's an interesting objection, but the original objection, "Hey, wasn't that creative?" when we look at Alpha Go, no, that's not a true contender here.

Robert J. Marks:

I think that the Alpha Go was trained to play Go. That's what it did. It did exactly what the programmers designed it to do. So let me ask you this. In your monitoring of the Lovelace Test... You proposed this in 2003. Have you seen any place where your Lovelace Test has been passed by AI?

Selmer Bringsjord:

No, but I would imagine... I've never had a specific, explicit conversation regarding this with any devotee of black box machine learning, specifically deep learning, which was all the rage, although it's now conceded that it needs to be changed, even by machine learning theorists. I've never had any conversation with anyone in that camp about the Lovelace Test, but I'm pretty sure that if I did that, they would say, "Well, look, you yourself have criticized our systems for being black boxes. So if they're black boxes, they satisfy one of the criteria that you've given us, which is the developers don't know how it happened."

Selmer Bringsjord:

Now, they'd be to some degree dishonest, but it's true. These are black boxes, and that's why there's been much concern about explanations being absent as to how they do what they do. But in the case of

all these machines, we get back to what I said about Alpha Go in the case of Go. We know the functions are relatively simple that they're trying to learn. In fact, we, we know where they fall mathematically.

Selmer Bringsjord:

A car, learning how to navigate roads with road signs and stay within the path that it's being given and refined. Come on. Just because we don't know how the system arrived at its ability to approximate the computing of that function in the real world doesn't mean that we as human beings don't know that function is extremely simple and it is simple. So the things that the Lovelace Test, if passed, would present us with and confront us with are not relevant to those kinds of tasks.

Selmer Bringsjord:

So, I'm sure I would hear this from machine learning folks. There's no question. In fact, I'm shocked I haven't heard it explicitly aggressively, but I'm sure that day is coming.

Robert J. Marks:

I can tell you that in monitoring the literature, having learned of your Lovelace Test, I usually apply it to what I see in the news, and nothing that I've seen really passes the Lovelace test, which is really, I think, fascinating. Basically, computers do what they're trained to do, even the deep learning. It's trained to do what it's supposed to do, and it does that and does it quite well. You might get some weird results or you might get some surprising results, but surprising results do not equate to creativity.

Selmer Bringsjord:

No, that's exactly right. And the pre-engineering or the setup, the engineering that provides the setup for the process of "learning to happen," is just as you say, highly premeditated and very rigid and nontrivial. But the point here is, again, it's understood from the get-go and provides a rather artificial context and precludes creativity.

Robert J. Marks:

Let me ask you a question. If somebody came to you and said, "I had AI do this," what would it take to convince you that AI had passed the Lovelace Test? An example.

Selmer Bringsjord:

Yeah, I don't know enough about the visual arts to go in that direction, so I would pick something from the space that I know and spend a lot of time experiencing. I would say if the machine can produce a novel of the right sort, that's going to get my attention. And then I have to verify all the background information to make sure that it satisfies the constraints on the test. I can't have a pre-stored novel. I can't have passages that have been pre-stored with someone banking on the fact that I'm not familiar with Proust because their variance on... All that stuff verified in the background, and yet it still gives me a novel of the right type, yeah. That's something worth writing home about.

Robert J. Marks:

Let me tell you an example I came up with. I'd like to get your reaction. If Alpha Go, without additional programming, was able to play checkers. I think that that would be creative.

Selmer Bringsjord:

It would be creative, but there, in my mind, is a continuum of types of creativity. That would be some form of creativity. I would say that's problem-solving. And probably the world's leading authority on musical creativity, David Cope, does say explicitly that if the machine can do problem-solving that catches people by surprise, he would stick to his guns and say, that's creative. I absolutely reject that notion.

Selmer Bringsjord:

I think the next step up from that is MacGyver creativity, or what I've called M creativity fairly recently, and that is all the humans put their minds together and get a wonderful artifact intended to be used in a particular domain for a particular set of tasks, the machine takes it and does something completely different with it. So that's beyond what you're talking about because, not only is a game, but the two games you're talking about have the same formal structure. That's why we know how difficult they are. And then the history of AI general gameplay, which has petered out, but which was a brilliant idea, I think, of Michael G, tells us that that jump is not ... From Go to checkers or whichever it is, it's not that large. But even the MacGyver creativity, which I don't know that there are cases of that out there, but even that is a long way from origination. Something that is at the level of theory, different than what we've seen, so I didn't get to specify the type of novel that it would take but I'm heading in that direction.

Selmer Bringsjord:

Yeah, I do see what you would be willing to say is creative as creative, but it's a problem solving type of creativity and that doesn't cut it. That's not genuine origination. And that was the complaint to Turing was ... Lovelace was like, wait a minute, we originate things, Alan. To paraphrase. We originate things, our computer doesn't originate anything.

Robert J. Marks:

Yes. Excellent point. You're familiar with Ray Kurzweil's singularity, which assumes that AI can write better AI that writes better AI, et cetera, et cetera and pretty soon we're going to have the capabilities of humanity exceeded. But that assumption assumes that the AI can be creative, if it can write better and better, more creative AI. And since computer programs can't be creative, in the sense of your test, that's never going to happen. Would you agree with that?

Selmer Bringsjord:

I would. That's something I hadn't thought about, but you're certainly correct. I have a little piece on the formal requirements of one computing machine or information processing machine creating another. And I think there are mathematical problems with that. I hadn't thought about that in terms of creativity, but you're certainly right. That's a creative act and that's a big, probably indeed, as you indicate, fatal problem for what he believes is going to happen.

Robert J. Marks:

There are some who say immortality is available if we can upload our minds to a computer. This, though, presupposes our minds are computable and can be duplicated by a computer. David Gelernter said, "I doubt if there is any kind of human behavior computers can't fake." Do you have a response to that quote? Do you think computers can fake just about any human behavior?

Selmer Bringsjord:

If the time parameters and a few other background, contextual pieces are fixed favorably for the machine, I would agree. And I suppose my first book, *What Robots Can and Can't Be*, can be viewed as a book length of fence of the proposition that David advanced.

Robert J. Marks:

Excellent. By the way, we'll give a link to Dr. Bringsjord, his book, on the podcast notes. There's a difference between faking something and actually duplicating it. Are there things humans do that can't be duplicated by a computer?

Selmer Bringsjord:

Yeah, I think, apropos, of some of our earlier interaction, things in the area of creativity, anything in creativity or creative behavior that leverages consciousness, self-consciousness, since I don't think consciousness of the subjective variety can be possessed by any computing machine. That would be impossible to duplicate. But I suppose the machine, over a short amount of time, could generate sentences, say, for an opera in the libretto, and maybe the music accompanying it, which when performed, just knocks people over because it's so incredible. So faking there would be accomplished and be measured by the reaction of the audience. But it would, going back to your question, it would be faking it so it's not duplication.

Robert J. Marks:

Well, one of the things that I see in terms of measuring creativity of the arts is the subjectivity of the assessment. It seems that different people are just going to have different views and it's not something itself which can be measurable, except maybe in a fuzzy sense, and wouldn't that be problematic?

Selmer Bringsjord:

It is. I think it's intensely problematic and has repeatedly surfaced in the domain of painting. In connection, specifically, with the AI artist, Eric. If a machine is doing painting and it's all fuzzy and maybe it's praiseworthy because it is in terms of its assessment because it's modern art or extremely abstract and conceptual, you get a lot right out of the starting blocks toward many people counting it as creative. That's not going to happen in sending the manuscript of 400 pages off to the literary agents in Manhattan. After two paragraphs, I don't care how many agents you send it to from this AI, the result is going to be the same condemnation and probably it's going to be tossed aside and shredded.

Robert J. Marks:

Agreed. You wrote a paper called *Cognition Is Not Computation* and I wanted to talk about that with your permission. Could you, first of all, define cognition and then we'll talk about why it's not computable.

Selmer Bringsjord:

Yeah. Cognition, at least in large part, I define in keeping with how the relevant scientists have characterized it and studied it. And that generally would be cognitive science and maybe even more narrowly cognitive psychology. So fixing your beliefs on the basis of evidence, reasoning, decision making on the basis of evidence, or perhaps not on the basis of evidence, but semi systematically. Emotional States driven by the environment, hoping that something will happen, fearing that something will happen. All these phenomena are part of cognition.

Selmer Bringsjord:

I think sometimes it helps to insist that we're talking about higher level cognition here, these are the kinds of things that are solely the province of human persons. Animals cannot lay claim to doing any of these things, at least to any appreciable degree. A mouse may be able to navigate a maze and so it's "making decisions," but I'm talking about decisions that imply a rationale that can be articulated. Linguistically, as a matter of fact. Proving things, discovering things, all these mental verbs constitute, for me, what cognition is.

Selmer Bringsjord:

I don't think more than a hair's width of these things have even been simulated in computational cognitive science and in AI. And the question is which route does one fall into? I remember asking James Moore, the Dartmouth professor who's written quite a bit, at least a while back on AI, "Jim, you really are a true believer in this stuff, but can you tell me how much time you're willing to give these AI people? I mean, if we give them another thousand years and we still don't have cognition, as I've characterized it, captured computationally with the relevant artifacts and outputs produced, are you going to be skeptical now?"

Selmer Bringsjord:

And he was, as an academic, I suppose predictably clever and evasive, but bottom line is we don't have this cognition captured. So you can either say it'll eventually happen as a true believer, or you can start taking seriously the possibility that there's something deeper going on. There might be some mathematical gap that we just don't get. There might be some kind of qualitative gap in the case of consciousness, because a lot of what I enumerated requires consciousness. You can't hope deeply that something's going to happen without being a conscious creature.

Selmer Bringsjord:

So that's, of course, where I stand. I think the math is ... people don't want to confront it, but it's very stark. I think we do things that are beyond what Turing machines are capable of doing, and we're also conscious. And we exploit our consciousness to do things every day. Every day that we're healthy and neurobiologically normal, we do all kinds of things that exploit the fact that we're conscious.

Robert J. Marks:

Some of this posturing for things happening in the future is pretty old stuff. I actually saw a video of Claude Shannon, I believe it was around 1960, who said, "I believe in a few years, that artificial intelligence would actually be able to duplicate the mind, in the sense of the robots, that we see in science fiction films." So even Shannon, back then, was posturing about something happening in the future. And in a way, it's scrutiny avoidance. We always say something is going to happen in the future, like Ray Kurzweil says, that singularity is going to be here in 2049.

Robert J. Marks:

Well, goodness, 2049, Kurzweil will probably have passed. I'll probably have passed. And he's just going to be some blip, very dim image in the rear view mirror that nobody's going to care about. But boy it gets a lot of attention now. So I think that this delayed scrutiny is something that has been happening for a heck of a long time. In your paper, you claim that consciousness is a special case of cognition. That's the first time I've heard that sort of claim. Could you elaborate and unwrap that a little bit?

Selmer Bringsjord:

Sure. Let me just say first, I couldn't agree with you more about the delayed scrutiny, and the antidote to that, by the way, is taking a bet. But nobody wants to take the bets, so if I'd been around when Herb Simon ... So we're talking 50s and he said, "Well, in a few years, we're going to work it out." Or when Marvin Minsky, it's a few summers, or maybe even one summer, well don't worry, we'll bring you back to say ... What I don't get about this is, really, well here's \$5,000, here's 50, here's 100. Here's my entire 403b that says right now, right here, you can take it and we'll deal with ... My descendants will sort this out, let's work out the contracts, I'm willing to make a bet. You define the test, you take the money if it happens, or your descendants, otherwise my descendants get the money from your estate. I've said this, no one's willing to take the bet. So that's the only antidote I know to delayed scrutiny, that's for sure.

Robert J. Marks:

Yeah, there's an old story about the farmer that was selling his tomatoes to this trader and the trader bought them for \$3 a pound and he went to market and he sold them for \$2 a pound, and he was losing money, of course. And when he took to his accountant, the accountant who wasn't too smart, said, "Well, I have the solution, you need a bigger truck." All of these delayed scrutiny things remind me of that, as an algorithm of the gaps that someday we're going to have an algorithm that does this sort of thing. Yet it's been promised since 1960, what is that? 60 years ago? And nothing has really happened, and nothing has passed, as I see, the Lovelace test as you proposed about 15 years ago. So consciousness being a special case of cognition.

Selmer Bringsjord:

Well, right, I certainly will agree that it is not in any way universally affirmed. And some people steer clear of consciousness and try to prevent consciousness from entering the scientific discussion. Whether it's neuroscience or even sometimes things that are more formal, like treatments and decision theory. Whether it's normatively correct, that is whether the agents are good at it or whether they're making bad decisions. No, no, let's come up with an account of decision making that doesn't take the consciousness associated with desires seriously. So we don't have to worry about what desire really is and the consciousness associated with that, let's keep it separate and look at the formalisms.

Selmer Bringsjord:

I'm not denying people do that, some people do. But unfortunately, when we look at the list of the mental verbs, sometimes philosophers, I think, call them attitudes or propositional attitudes. When we look at these verbs that are clearly part of cognition, because they're in the textbooks, they're in the cognitive psychology textbooks or the cognition textbooks. If we're just honest with each other, we'll see you can't instantiate these things in an agent, at least of the human variety, unless that human has feelings. Unless that human is such that there's something like to be the human, unless the human feels pain less, unless the human feels pleasure. All the things that are wrapped up with consciousness or that perhaps constitute being conscious are really there, so we just artificially separated.

Selmer Bringsjord:

So really, that would be my way of defending it. Let's just write down the activities that we say are part of being a cognitive agent, as opposed to just an agent. Because in AI, in the textbooks say that an agent just computes a function from percepts of the environment to actions, so even something as simple as a program that computes the square root is technically an agent in AI. But if we say a cognitive agent, we

can't suppress consciousness rising up before our faces and we have to deal with it. But again, some people can try to dodge it and suppress it.

Robert J. Marks:

So if we have these things, cognition and consciousness, which are attributes of humans, your claim in your paper is that they are non algorithmic. That is, you can't write a computer program to simulate them. They are not computable. What is your argument that cognition is not computable?

Selmer Bringsjord:

Well first, to be careful, some of them are not computable. Clearly playing checkers is a computable process, provably so, by definition. If we want a simpler case applicable even young children, then tic-tac-toe. Very young child can learn an infallible algorithm for tic-tac-toe, and when they make those decisions, when they do that reasoning, that thinking, and even the perceiving of the board or the symbols, if it's written out on paper. I'm not saying they're doing something which is not computable.

Selmer Bringsjord:

And notice we have machines that can do those things. And even skeptics like me would have conceded immediately, maybe even centuries ago, maybe a millennia ago. I figure Aristotle would have said, "Oh, that's an interesting game, checkers, but I'm pretty sure that is ..." and he would have not been able to attach a detailed mathematical sense to the term, but said, "It's mechanical, it's calculative." But I'm talking about things that distinguish homo sapiens sapiens as a species, things that distinguish the human person.

Robert J. Marks:

So what would be some examples of cognition which were not computable. Clearly chess and checkers are computable.

Selmer Bringsjord:

Right, well, remember at the top of the list, although it can come across as elitist, the top of the list is discovery and confirmation. Conjectured discovery and confirmation in the formal sciences, or to use what probably is good enough, mathematics. Doing mathematics where you're making discoveries and conjecturing them, and then making discoveries and confirming them, is untouchable. So I'm working a lot currently because I have a book, eternally I'm done, but getting quite close now. Fortunately .. On Gödel's great theorems, and look, I mean, if you just look at one little piece of his career where he proves that the continuum hypothesis, basically there's no number, there's no set of a size between the natural numbers and the reals.

Selmer Bringsjord:

Okay, continuum hypothesis, got to be one or the other if you find that infinite set and there's either got to be equivalent to the naturals or equivalent to the reals. This is astounding, this is positively astounding work. So when we talk about AI doing all this stuff, it doesn't really do anything in mathematics. I mean, how does an AI generate, by the way, a new paradigm, a new branch of ... Math has expanded, so we don't think we've got even all the branches of mathematics before us. So that would be at the top of the list. The great thing about that one is that we can inspect the output produced by humans who are playing in this space.

Selmer Bringsjord:

So it's not like they just give us vague reports about doing these amazing things. They write the results down so we can look at the results and then we can say, "Now, can a machine generate something like this?" And the answer is flat out, with a ring of iron, no. We're nowhere. And this would be my number one. Number two, I think, would be creativity, and I'm reluctant to go away from literature because of ignorance, it's harsh to the visual artists and the painters. But things in the literarily creative realm, as much as he was a maniac, Wagner. I mean, how does one human being create the storyline, the music, the libretto, which is essentially poetry, and weave it all together and produce that out of whole cloth? I mean, this is crazy. So number two would be that.

Selmer Bringsjord:

And because those are elitist, I guess I have to round it off by saying, let's just think about what it's like to love. I mean, what is love and what's it like for one person to genuinely love another person and be loved by that person? Now we can't mathematize that, we've got no account of what it is. In fact, the leading formal account of human emotions, the so-called OCC account, which I like very much, has come up totally empty on any kind of formal account of love.

Selmer Bringsjord:

And yet we love people, and we want to be loved, and we know what we're talking about, and we know what the sensation is associated with real love or like. I'm not talking about mere pleasure, I'm talking about love. So everyone who's a human being on the face of the planet can just see there's a major problem here. From love comes, arguably, the other two things. Why do people insist on locking themselves in a room in Princeton and working on the continuum hypothesis? Because they probably love their interaction, even with the formal space that they're involved with.

Robert J. Marks:

Did Gödel prove that while he was at the Princeton Institute?

Selmer Bringsjord:

Yes. Well, he proved one half of it. The total independence of the continuum hypothesis from set theory, from what we assume to be the basis of all classical mathematics, was accomplished in two pieces. Gödel being the first piece, and then Cohen, who, as I understand it, actually took a road trip to Princeton. The god of the land of calculus told me a story about this, that Cohen came to Gödel and said, "Oh, I proved the other half. I'd like to meet with you," and Gödel said, "Well, today's not the day when I meet with people. Come back." And I guess he found some place to stay for a couple of days and he came back and he handed him the other half. Which, by the way, introduced the technique that only a few have heard about called forcing, which is another example in the arsenal I'm talking about to shoot down any notion that AI is amazing.

Selmer Bringsjord:

This forcing technique that Cohen came up with is, in and of itself, ridiculous. And anyway, Turing certified it, came back with it and said, "It's correct, thank you very much," and they didn't share any other words. But yes, Gödel did that. I believe in '30, I forget the date of the first part. It took about 25 years for Cohen to do the other side. So we really don't know what the status of this continuum hypothesis is. I once remember John McCarthy, in an interaction I had with him as one of the founders of AI, trying to endlessly argue with me saying why it was true. And it was unforgettable, but it was

entirely intuitive, and there is no classical proof that's possible on the basis of what we take mathematics to be. We know this is just impossible.

Robert J. Marks:

Yeah, if a listener is a mathematician or a nerd in general, the continuum hypothesis is just an incredible piece of work. So I think the bottom line is, don't upload yourself. Don't pay anybody to upload you while you're on your death bed, because they're not going to get the non-algorithmic you. They're only going to get the algorithmic you, and that upload is going to be very, very incomplete.

Robert J. Marks:

AI is starting to write things and some of the copy reads pretty well. Other writing is garbage. You wrote an article called, Artificial Intelligence and Literary Creativity, that was a long time ago. In fact, it was last century, 1999. And more recently in 2013 you wrote, What Robots Can and Can't Be. Now in this work, you claimed that AI will never write creative prose. What's your basic argument that AI will never write creative prose?

Selmer Bringsjord:

Yeah. I'm not sure I can distill it to a single argument, but there's certainly the one that leads the way. And I'm talking about novels here. Because I remember at day camp, my mother sent me to this expensive ritzy day camp when I was younger, it was wonderful. But I remember we had some session where the counselors were trying to convince us that by rotating a wheel we could all be story generators. We could all make stories. And for whatever reason, the counselor leading the way took this very seriously. You change the character names, you change the actions, but essentially every story, maybe not literally, probably true, every story had the same number of sentences. And I remember saying, "Oh man, I'm so sorry, but could we go out and swim in the lake? Could we do something else?" And, no, no, no, no. My gosh, there's no real creativity. It's a wheel and it's boxed us into a corner. So at the other end of the spectrum from this kind of plug and play way to generate fiction, we have great exemplars, so we have...

Robert J. Marks:

You have to elaborate on that a little, you spun the wheel and it stopped at a place, what were some of the examples?

Selmer Bringsjord:

Well, it was conscious and deliberative. You could change the character names, so you could make the protagonist, instead of Betty, be Billy, so now we've got a gender change. With that would automatically come appropriate changes for the pronouns and the gender switching there. Basic actions were set for movement along a storyline. And there was an arc to it as I remember, so you've got a genuine story, but it was always the same frozen space of stories. I don't know what the size of the space was, but I think some of these things, years and years ago, I'm dating myself, I think some of these things were actually even in cereal boxes.

Selmer Bringsjord:

And from the standpoint of the history of AI, actually they go back to Ramon Llull, who had an impact on Leibniz, who was of the opinion that if you built such a wheel big enough, you could figure out how to reply to any skeptic when trying to convert them to Christianity. And so it was a small space of stories

and very rigid with minor variation, and that's basically what we're talking about when we're talking about fanfare being stimulated by some of the things AI is doing today. I'm talking about things at the other end of the spectrum, I'm talking about...

Selmer Bringsjord:

So to your question and to the leading argument, I don't care who you are, if you're an author and you're aspiring to do something that's significant, you referred earlier to the Nobel Prize, the Pulitzer Prize, if you're aspiring to do something like that, win something like that, you are really out of luck if you don't have characters. If you don't have characters, you lost before you got started. You can say, "Well, I want to do a prose poem". No, I didn't say prose, I said novel, I'm talking novel. Sustained narrative needs characters.

Selmer Bringsjord:

All right great, so now we know that, maybe that's just being reminded of something totally obvious, but if we need a character, where does that character come from? How do we get the character? Do we steal the character from some movie that's already been done, so we've got the... Yeah, maybe we do that. Let's take the character from... Well wait, but where did the person come up with the character that was in the great literary movie? Where did it come in the case of the characters in fiction? You have to create the characters and in order to do that, you have to put yourself in the shoes of the character. And if you put yourself in the shoes of the character that you're creating, you're exploiting the fact that you are a conscious creature, you're doing a form of creativity that requires consciousness. And that, since machines can't be conscious, is a real problem for the machine.

Selmer Bringsjord:

So, that's argument number one, argument number two is... I'm not saying who this person is, but I was at an event, a social event, and the very topic that you have engaged me on here, which clearly I'm probably overly passionate about, but was active, and the assertion was made to me; "Well, writing is very mechanical. The syntax is pretty simple, writing's mechanical, we've got Shakespeare emulators and Shakespeare parties where people just on the spot confabulate and sound just like Hemingway", sorry, not... Hemingway, I don't know if I said... I don't want to say Shakespeare.

Selmer Bringsjord:

And my response was immediate. "Really? Have you read any Proust? Have you read all of Proust? Have you read a 427 word sentence by Proust that was created from scratch? So wait, excuse me. There wasn't a 225 word sentence that was reminiscent of the 427 word sentence, no, he wrote the 427 word sentence from scratch, okay. So don't tell me the syntax is... yeah, you can hand pick the syntax that you say is simple, but I'm going to tell you right now we have machines that are supposed to understand natural language, translate natural language that they have "understood" into some other natural language, and also generate from underlying content, natural language. And I'm just going to tell you..." And it was a cheerful exchange, although there was an edge to it I won't deny, "I'm just going to tell you, I'm familiar with that AI technology and you try giving it not the longer Proustian sentences, just take one that's in the top 30% for complexity, or at that level, let me know what your parser does with that sentence, I'd love to see that result".

Selmer Bringsjord:

So problem number two is that the language itself is too complex. And the dominant approach today to natural language understanding with the advent of machine learning was painfully reflected in the movie *Ex Machina*, where I think it's kind of a dark movie and it relates to the Turing...

Robert J. Marks:

What was the name again?

Selmer Bringsjord:

Ex Machina. *Ex Machina*, it's really about the Turing test. So I think given your interests and your knowledge, and maybe even some of our prior discussion, you will find it... Again, I'm telling you it's dark and I don't think it's going to win, it's never going to win any placement in a top 10, but it's about the Turing test, it's about androids that can pass the Turing test and at one point the protagonist says, "Wow", I'm paraphrasing, obviously, "Wow, this AI is unbelievable. How did this AI arrive on the scene with the capacity to talk back and forth conversationally in a way that is this good?" And it passes the Turing test, unrestricted Turing test, and the answer comes back, "Well, we listened in on everyone's cell phone conversations and got all that data and then we unleashed machine learning on it".

Selmer Bringsjord:

I was allowed to go to a prescreening, I kind of panned it, so I went to a prescreening and I get to that point in the movie and I'm watching and I hear this and I'm like, "Oh my goodness". So now we know that anyone with any degree of skepticism about how language and AI works together and how hard it is, is not going to bother taking the movie seriously. So some consultant told them to put this in, you should have left out the details. The minute you put that in, you've got a big problem, it's that Proust isn't talking. There aren't 10 million Prousts out there Twittering, or tweeting, or texting, or writing emails or letters back and forth so that the data would support the machine learning how to do something that's truly impressive from a literary point of view.

Selmer Bringsjord:

So this approach is not going to work. There is no data on this by definition, because if the human is originating new paradigms, new ways of writing, stretching the boundaries of syntax, this technique is screwed. You're dead, you can't do it. So that's reason number two why this is a tough challenge. By the way, both of these challenges, if you put your mind to it, I suppose, in AI and you have enough money backing you, can be modulated by saying, "Well, please let us try to do it for a two paragraph long story. Don't make us do a novel". And just to be clear, I'm talking about a novel. They've been around for a long time, we know what they are, I suppose, maybe *Don Quixote*, that's a novel, so that's what I'm talking about. So those are the two arguments, I could keep going, but those are two.

Robert J. Marks:

I just listened to a great translation of *Les Misérables* and it was 60 hours. It took me 60 hours to listen to it, but wow, could Victor Hugo turn a phrase. I mean, just the incredible way that he was able to communicate in his... I don't know if we've lost that today, but he is just so incredible. Now the length of the article kind of brings me to the next thing I wanted to talk about. And this is a report from the media, from pretty high credentialed places, about AI being able to write. The first headline is from *Forbes*, it says "Artificial intelligence can now write amazing content and what does that mean for humans?" And included in that article is the statement, "Many financial institutions are churning out 10 to 15 page financial reports in an instant by using Narrative Science's NLG platform Quill". So that's one

of the claims that these people are making, that these financial reports can be generated by this AI. I guess in one sense, that's really not being creative, is it?

Selmer Bringsjord:

We've talked, I think, in the past about grades or levels of creativity. I wouldn't be unhappy saying that some simple form of creativity is satisfied in the case of these NLG systems, but clearly we're not talking about creativity that's necessary to write a literary or genuinely belletristic novel. The reason is, in the case of finance, the data, the knowledge, a sentence about the current market cap of Apple Computer, if it's a nice polished sentence that the machine has generated, wonderful. Well, where did the knowledge come from that the sentence expresses or communicates? It was already there. It was pre-stored. In the case of creative literary activity, you've got to create the stuff from which you generate to begin with.

Selmer Bringsjord:

So you can't... I guess I wouldn't consider Elmore Leonard's fiction as entertaining as it is creative, but look, even he didn't just start down and write "stream of consciousness", do the generation of natural language. What he did was get some characters in mind first and a couple of situations, then he was off to the races, so he created all that background knowledge himself. In the case of finance, all of that is pre-stored. Now, it's a wonderful AI strategy for making money. This is, and I was of this opinion, a friend of mine had a startup company that did rather well in this domain, it was acquired. I want to do natural language generation where the content is already pre-established and kind of rigidly represented under the hood. I want to do the generation to put it into natural language for humans to digest. Great idea. Great idea. It's an engineering problem. It's wonderful. But if you have to get the knowledge itself from scratch, this is really not relevant.

Selmer Bringsjord:

Now, sports would be similar, I am a sports addict and my wife is very tolerant, but I'm a huge Patriots fan, I follow the NFL and I still maintain it's... The sport gets you injured, there are going to be ethical problems with it, but from pure fan pleasure it's just amazingly complex, et cetera. And sometimes I listen to games and I'm looking at people commenting, and they have a six, seven, maybe, I don't know, eight figures, they get paid a lot of money and they're doing the running commentary, et cetera. And sometimes I'm listening, I'm thinking, "Wow, you're lucky that you live in 2019 and you don't live in 2025 or 2030, because you wouldn't have a job doing what you're doing".

Selmer Bringsjord:

Now, there are other people who are doing commentary, who are doing what I'm talking about, where they aren't just in command of everybody's jersey number, everybody's basic stats, everybody's reported personality and tendencies before they show up and then do play by play, they're doing something generative at the knowledge level. And as much as I'm not a Cowboys fan, I have to say Tony Romo, he's done amazing things right out of the box cognitively, because he said, "Oh no, did you see that? That's inaudible and here's what's going to happen". Now that is quite different, it's heading in that direction of being knowledge generative. He's making inferences to generate new knowledge on the spot and then he's tapping into the background knowledge that he has as a good broadcaster because he's done his homework, and then he's generating it in English. So we don't see that in the case of finance, we don't see that in the case of summarization either, because summarization by definition is piggybacking on prior content that's already in play. So you can imagine I'm not super impressed by some of the things that you rightly cite as interesting.

Robert J. Marks:

Yeah. There's other sorts of places that generate short commentary, USA Today uses something called Wibbitz, and News Tracer is used by Reuters to actually comprise short reports, and BuzzBot was originally designed to crowdsource reporting from the Republican National Convention. So yeah, all of these are actually conglomerating information that already exists there. You talk about football, it strikes me that much of sports radio could be probably replicated with artificial intelligence. It seems to me that they say the same things over and over again like, "Frank, it's all about getting points on the board. It's going to be all about whether we can move the ball. We need to stop their passing game", all of this stuff. And I bet you that could be replaced by AI very easily, but that's my takeaway.

Robert J. Marks:

Now there's some other paper generators, and this is really cool. One of them is called Sci-Gen and it actually generates computer science papers and they're really weird, but they've actually been accepted for publications in places like IEEE and Springer where their peer review process wasn't at their top, they were later found out and removed. And another one was MathGen, and if you go to mathgen.com, you can write your own math paper. And I actually entered our names and we have a paper now, Selmer Bringsjord and Robert Marks, called "Positive Pseudo Independent Finite Submarines and Rational Algebra," and this great name was generated by it. But I read these and you're right. If you go to any length, you detect that there's something phony, especially the more expertise you have, the greater domain expertise you bring to the paper, the quicker you recognize it. I don't know if you're ... do you remember the movie, the Stepford Wives, or Invasion of the Body Snatchers?

Selmer Bringsjord:

Yes. I mean, not scene by scene, but certainly.

Robert J. Marks:

In both of those there were robots, or in one case it was an artificially generated human by these pods from outer space in Invasion of the Body Snatchers. But you looked at them, they looked human and they acted human, but if you were with them for a bit, you realized that they were phony. And that's kind of been my experience with these writers, like the MathGen papers and the Sci-Gen papers, they look good when you look at them, but you get a little bit in-depth and you realize they're totally phony.

Robert J. Marks:

And I suspect that the same thing you're talking about in terms of AI writing great novels, it's going to be exactly the same thing. In fact, there was one, this is from SingularityHub, it said the first novel written by AI is here and it's as weird as you'd expect it to be. Now I have not read the novel, but I think weird means that it really is very incoherent, it comes up with very very strange phrases. So even though we have AI writing all of these things, your claims, you believe, still stand up about AI not being able to write good prose.

Selmer Bringsjord:

I think so, at least that's my opinion. And it's interesting you mentioned the domain expertise, the variance of domain expertise and what role that plays here, that if the person judging is familiar with the domain has some bona fide knowledge, perhaps even expertise, they know that something's off and they're not going to give it high marks. So yeah, that's true and it does go back to the Turing test, and I think maybe something we have discussed in the past, that as far as anyone can tell, Turing was of the

opinion that you shouldn't prevent judges from showing up who have any particular expertise. And people have shied away from that when running sort of pseudo Turing tests in the "real world". Here, yeah, Forbes telling me it's weird, I could not wager enough money that what you suspect is true, that weird means this is really not very good.

Robert J. Marks:

Exactly. Oh, one of the big things too, is to collect everything in an overarching theme. It's kind of like AI can only look at its feet, it can't look at where it's going. And I think because it always looks at its feet, it always gets this kind of bad rep as you walk a little bit, because again, it doesn't know where it's going. In your talk a few minutes ago, you used a word that I had to look up and that's the word belletristic, and it was the first time I've heard the word. Could you define belletristic? Because that seems to be a key word in your discussion of what AI can and can't write.

Selmer Bringsjord:

Yeah. I think to be truly literary means... it's got to be truly, when applied to fiction, it's got to be truly literary fiction. It's got to have great depth. And I think ultimately, and I'm partially operationalizing this and going beyond what dictionary.com undoubtedly's going to say, but I need the people living at the time, who are people of letters, to apprise this fiction and allow it into their community to affirm it. So that's necessary. Now, we know though what the attributes would be that would be necessary for genuinely literary fiction, as another way to say belletristic fiction here, if people of letters say, "But you only have a character that had one continuous emotion, it didn't seem like a robust character. I don't understand, why do you think this would be a good...?"

Selmer Bringsjord:

Step number one is what they tell you in film school for three dimensional characters in the software that humans use to write screenplays; you can't have belletristic fiction unless you have a truly rounded, well-developed, coherent character. So belletristic, truly literary fiction that would be regarded or classified that way and given the suitable accolades that you would expect, that's what it would take.

Robert J. Marks:

In a way, it goes back and talks about domain from the credential people that you talk about, who are looking for aesthetic values rather than specific content, they're looking at things from a higher level.

Robert J. Marks:

Young Kurt Gödel toppled a tall tower of mathematical reasoning with his publication of his work, showing that no formal system of math could be both complete and consistent. Gödel lived the last part of his life in the Princeton Institute where he became buddies with Albert Einstein, indeed Einstein accompanied Gödel when Gödel was sworn in as a US citizen. Interestingly, Gödel also gave a formal mathematical proof of the existence of God, and that's what we're going to talk about today on the program. Before we talk about your paper that you presented in October 2019 in Romania entitled The Argument for God's Existence From AI, let's talk about the book that you're writing. You're writing about Kurt Gödel and his proofs, could you tell us what your book is about? Because Gödel was just a monster genius in terms of mathematics.

Selmer Bringsjord:

Right. Well, the book's a little different in that I've set myself the goal of covering all his great theorems. I may miss one or two that some people would say are great, but...

Robert J. Marks:

Oh, so this would include, like we talked about the continuum hypothesis?

Selmer Bringsjord:

Yes, indeed.

Robert J. Marks:

And probably stuff he did... he did things in cosmology too, didn't he?

Selmer Bringsjord:

There is indeed the Gödel time travel theorem, is what that chapter's called, his present to Einstein, his birthday present to Einstein showing him that general relativity courtesy of Einstein, or exactly as Einstein's equations defined it, is consistent with a form of backwards time travel. He proved that. There is a chapter on that, yes. So there's a chapter for each of the great theorems and there is a chapter on what I call the God theorem.

Robert J. Marks:

And that's what we're going to talk about today. What's the history here? What's the history of this so-called God theorem? Gödel never published this, did he?

Selmer Bringsjord:

No. No, he didn't. He worked on it... And there could be varying degrees of speculation, or how speculative the hypotheses are about how long he worked on it and how seriously he worked on it, but it seems pretty much undeniable that he worked on it for a long time and put a tremendous amount of concentrated effort into it. And the first time to my knowledge that it left his consciousness, and his wife wouldn't have probably been able to assimilate what was going on, was when I think he thought he was dying, he showed his notebook or a notebook to Dana Scott, who was not one of his students as in your advisor of the descendant, he'd be... your descendant is your advisee, et cetera. But he was a student of Gödel so that means he was learning from Gödel. So he gave his notebook to Dana Scott and Dana Scott copied out a proof of God's existence that Gödel had worked on and that is the first... It's now been reproduced it's in the literature, it's in the AI literature in formal form. That's I believe the first transmission of this discovery from Gödel to another human and perhaps any other externalization of the ideas. And even if one doesn't think the argument ultimately succeeds, it's still a discovery for reasons we can perhaps get to.

Robert J. Marks:

You know, I actually have arguments with people about whether God exists and I think it's very fun to just show them the page of Gödel's proof. Which is just for those not aware of the mathematics, it just, it looks like hieroglyphics, but say, yeah, here's the proof that Kurt Gödel did. So take a look if you have any more questions you can ask me about whether God exists or not it's... Was there any evidence that Gödel was a deist, a theist or a Christian at all, that you're aware of?

Selmer Bringsjord:

It's certain, I'd like to meet the person who denies it, it's certain he was a theist. It's also certain, I would claim that he was maybe a non-organizational Christian. So he clearly ... toward... As things got worse he wasn't able to socialize in any environments. Just having a dinner without thinking someone was going to kill him. He died under very, very distressing psychological circumstances. So it's not like he's the ultimate barometer of whether someone could tolerate organized true religion, but he wasn't an organized or organizational. To use my adjective here on the fly Christian. But yeah, I believe you can find some stuff regarding his wife's observations and his relationship also with his mother who was very devout.

Selmer Bringsjord:

And on a personal note, having looked at not just the argument, but a lot of the life's work and the picture that we have emerging you pretty much would have to be an idiot to take pot shots at any such proposition, negative proposition as, Oh yeah. He wasn't, in some sense a believer whether it's fully Orthodox and organized. You'd have to be crazy really to say that because across the board, his views on the mind, right, his views on the future it's really a pretty compelling picture, but again, for people expecting if it was John Calvin who showed up to interview him, that would be big, big problems for some systemic thinker theologian this... But no, he clearly was in that direction to put it mildly.

Robert J. Marks:

Well, if at least he wasn't a deist or a theist, what motivation would he had for offering such a proof of the existence of God, certainly?

Selmer Bringsjord:

Right, the only countervailing what you would get is a... Well, he really liked the thinking or thoughts or work of Leibniz rather than the topics and that's absurd.

Robert J. Marks:

Well let's yeah let's unpack that a little bit. So the origin of Gödel's proof actually goes back to alignments as you mentioned, is that right? Or unpack some of the history of the proof.

Selmer Bringsjord:

This is an unsanitized politically incorrect history. I'm going to tell you that right up front. And then if we can tell you why it is. So the year, 1000, approximately the idea that you can in your armchair by considering whether or not God is possible determined that God is possible and then infer validly from the mere possible existence of God, that God must exist that's 1000 or 1078, Anselm, who was a very, very subtle thinker comes up with this idea it's so-called ontological argument for God's existence.

Robert J. Marks:

Now, Anselm, A-N-S-E-L-M.

Selmer Bringsjord:

Correct.

Robert J. Marks:

I think that's how you spell that. And when did he exist?

Selmer Bringsjord:

Well, that was 1078. So he's turned the millennium now. A whole thing is kind of eerie to be honest with you, but that's the year, 1078. And then we start quickly getting into political incorrectness on the history of this. And by the way, if anybody wants a history, there's a wonderful, I don't know how focused it is on the history of it certainly includes it, in the Stanford Encyclopedia of Philosophy has a treatment of the ontological argument and the history.

Selmer Bringsjord:

But the problem is you have to go about two centuries before anyone does anything significant with this idea. And it's not clear what the heritage and transmission of the ideas but that's Aquinas and Aquinas doesn't like these ideas. So, no. So that's like 1260. And then you have Descartes showing up in the sixteen hundreds, 1640 approximately saying, Oh yeah, actually I do need to prove God exists because otherwise to modernize it, you might be living in the matrix.

Selmer Bringsjord:

And I don't want that. So what I'll first do is prove that God exists. And since God exists in is good and all powerful, he'll ensure that you're not living in the matrix. Therefore you can trust your senses, et cetera. So that's 16, I think 1635. And he claims that this is original. He didn't read it. This is original to him. So from a theological perspective, in terms of organized religion, and we already have a bunch of serious issues arising. But at any rate, now we have Leibniz and this is where everything changes. And some people believe Descartes, some people don't. I tend to believe him. And so Descartes says basically the same thing, but it's fancier, it's better. He argues that the mere possibility of this amazing being existing, you can see that by definition, if it's possible that it must exist after all the being is unpreventable, uncausable, the greatest thing.

Selmer Bringsjord:

So, but then Leibniz shows up and he writes his own version, the turn of the next century, just after 1700. And this is where everything gets much, much more serious, at least from the standpoint of modern, formal logic, and even first rate, I think rigorous philosophy and also AI, because what Leibniz says is, well, I'm a Christian, and he is, and I'm convinced God exists. And you only publish one thing in his... Yes, we have theodicy early in his life where he's just aggressively Orthodox, right? And talks about Jesus and so forth. So he is not happy with Descartes' work. And when Jones on the streets not happy with your work that's one thing when Leibniz isn't happy that's... He's probably got a reason for being not happy. And he comes to the conclusion that Descartes just kind of assumes that God is possible after all.

Selmer Bringsjord:

How do we know that the conception of God, that we have omniscient, omnipotent, omnipresence, put all these attributes together, how do we know they're consistent? So he sets himself the task of demonstrating this and giving another version of the argument. And then we jump all the way to the 20th century. And so another source of political incorrectness says there are lots some other people here. When we go to Gödel we skipped over the modern advocates of this argument, Norman Malcolm, James Ross, all these... It's harsh. I'm just going to say it from the standpoint of someone who's reasonably well versed in formal logic. I think it's kind of a bit of a doldrums, despite some of the

attention until Gödel does his thing. And Gödel does it formally. And, and then some folks in Germany doing automated reasoning verify it a few years back.

Selmer Bringsjord:

They verify the version that Dana Scott copied out of the notebook. But that is what they verify is that there is no doubt. It is a formerly, it's machine verified proof. So now we're left with just the truth of the promises and how we judge them. So that's really the history. I guess the subsequent chapter is the Gödel's own version, accurately transcribed from his notebooks was not formally valid. So Dana Scott made some decisions in the transcription that were well, if one is a theist, quite fortuitous, because now we have an invalid version and we have a valid version. And then the valid version post Gödel has been improved and repaired. Some of his supposed deficiencies have been repaired. So this goes beyond what I discuss in the book. And it's not what my book is about. I'd get in trouble if I went to these more modern modifications, but this is basically the history and I think the argument will never die.

Robert J. Marks:

Yeah. Well let's do that. Let's go to the argument that's a great history. Could you kind of walk us through the reasoning behind, I guess Anselm's ontological proof?

Selmer Bringsjord:

I might not be the authority to do that. I'm going to give it a shot, but all bets are off. And I will say to your listeners, if you want not only a different exposition, but a thoroughly modern one and one that matches also what a machine has verified as at least valid reasoning. You can go to the work of the editor of the Stanford encyclopedia Ed Zalta. He's got work with I think, Oppenheimer. They have a wonderful paper on this.

Robert J. Marks:

Could you spell the last name please. Ed Zalta.

Selmer Bringsjord:

Yeah, Zalta is real easy because it's Z-A-L-T-A, Zalta, Ed Zalta and absolute force in online content. That's academically first rate because the Sanford Encyclopedia of Philosophy has centuries and seems on everything. So it's no longer just philosophy. But anyway, right. So to your question I'm going to give up my version of it and with great debt to my professor now deceased at Penn, university of Pennsylvania. So, all right.

Selmer Bringsjord:

What kind of things are there, out there from the broadest perspective? Well, there should be things that exist and don't have to. Like you and me, like Baylor, like RPI, like New York City, it's great they exist, but they don't have to. Then we have things that don't exist, but could like another city over 7 million people in New York state, we don't have that, but we could. Then we have things which don't exist and can't like a man, seven feet tall and six feet at the same time in place. These are things that are impossible. And then we have things that must exist. Things like probably the quantities and the natural numbers. It's hard to imagine how the concept of a set composed of two things could ever not exist.

Selmer Bringsjord:

That concept is indestructible the laws of formal logic, modus ponens if P then Q, P therefore, Q of everybody died that's not going to stop being true, it's not going to stop governing the universe it's true. Maybe the basic constituents of primitive or undisputed mathematics and sets, relations, properties, functions. So these four categories are nice. They seem to cover everything. So where do we put God? And I think what, again, with a debt to Ross, I think what Anselm realized was, wait a minute. If I look at the attributes that God has the greatest. It can possibly be conceived, that seems like a perfectly respectable coherent concept. So, God can't be in the category of things which don't exist. And can't because all those things involve contradictions. I cross out that category. Now, wait a minute, I'm left with three other categories, the category of things that don't exist, but could that can't be God, because if God's in that category, what causes God to come into existence?

Selmer Bringsjord:

God is supposed to be also eternal, uncausable, unpreventable. Oh, by definition God's not in that category. So I'm down to could exist, sorry, does exist but needn't and must exist. Well, the first one isn't going to work at all because if God exists and he's in that category, something could zap God out of existence, like kill him off. And that doesn't match the definition of God. So we ended up by what logicians call disjunctive syllogism, P or Q in this case, P or Q or R, S not P not Q, not R with S. And I think this is the best way to view what Anselm did in fancy, somewhat hard to understand language. Now he does not extract for his readers out of the greatest being that can be conceived. The attributes that we, at least in the Judeo Christian tradition and probably in Islam as well to a high degree but historical monotheistic religions, ascribed to God.

Selmer Bringsjord:

But, obviously Anselm was and Anselm and anybody who knows any... And that's what he's talking about. So he's talking about the best thinker, the most powerful agent, the most knowledgeable agents, we just make a long list. And we end up with pretty much the Orthodox list. And he says pretty much what I just said about them. It's real. Once you've got these attributes laid out in front of you, good luck trying to fit it to the screen train to any of the other three boxes. It just isn't going to work. And so I view the history of this thing as he makes the discovery, and we just have to painstakingly figure out how to express it in a way that is formally valid and at this point by machine. So it has to be verified of so much above reproach that there's no longer a question the reasoning steps are valid and then get enough people to agree the premises are plausible.

Selmer Bringsjord:

Everyone is... No matter what you do into the future, whether the year 3000, because the conclusion of the argument is so expensive, right? Because it can change lives, instantly if you accept it. People are always going to want to deny something. That's how... It's just how human beings unfortunately are structured, right. So we do the same thing in physics, no matter what you tell someone who's a proponent of the multiverse interpretation of quantum mechanics, I having seen this, I don't care what you say. They're going to keep pointing out that it is the most reasonable explanation and it is consistent and they are going with it. And anyway that's the broad, that's how I would do the broad stroke, a cheater encapsulation of what Anselm originally did, by the way, you probably know that Anselm did write the, I forget the title.

Selmer Bringsjord:

It would be something that he wrote a wonderful piece on the logic of the incarnation. And we know from that, that is why did there have to be an incarnation from the Christian perspective for redemption and anyone who studies that, I don't care whether they're Christian or an atheist, just if you think the guy didn't know how to reason carefully and rigorously and powerfully, just read that and then come back to the ontological argument and you come back realizing, Oh, wait a minute. Is it pretty sharp, dude? I want to think about what he's saying here, who knows?

Robert J. Marks:

Excellent. Excellent. So I think, yeah, I think unpacking Anselm's, he's was a pretty smart guy, I guess, to be so ancient. There were a lot of smart people back there in history. It's kind of amazing. I heard of Cliff's notes, I don't know, see I like your reaction to this. Is that one of the axioms of Anselm is God is greater than anything we can imagine. And if indeed that was the case, then God exists. Existing is more wonderful than him not existing. So therefore he must exist. Now that leaves out a lot of details. But does that touch on the, his ontological argument at all?

Selmer Bringsjord:

Uncontroversially, it touches on it. But I'm not sure that it's consistent with what we do know is a linchpin, which has God is literally the greatest. Now we have to say agents or probably person in my opinion, but we have to have in his argument the premise, and maybe there's some subsidiary argumentation supporting it, but we have to have the premise that God is the greatest agent that can be conceived, or there's not... In such that there's no agent that you can conceive that is greater. We've got to have that premise because we have to max out, we have to. And this is a common, everyone has marched in accordance with this. Alvin Plantinga, for example, I said that some of this was politically before he does, he's got a version, he may have backed off from it now. But he explicitly says, we got to talk about what makes something great.

Selmer Bringsjord:

And we pretty much have to max out what makes something great as we analyze it. So we do need to get that premise of Anselm in the argument. I wasn't clear that you have that in your encapsulation.

Robert J. Marks:

I did not. That's a great correction. Thank you. So I think the bottom-line of today's podcast, even if you don't understand all the little Johnson tittles is that Kurt Gödel offered the proof that there is God, that God exists and we will in the podcast notes actually make a posting of that proof that you can just go and show people if they say prove me that God exists, you can say, well, here it is Kirk Gödel did it. Our guest today has been Dr. Selmer Bringsjord of Rensselaer Polytechnic. And it's been a wonderful time having conversation with him. And we appreciate his participation in the effort of Mind Matters News. And so until next time be of good cheer.

Announcer:

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. Mind Matters news is produced and copyrighted by the Walter Bradley Center for Natural and Artificial Intelligence at Discovery Institute.