

The Non-Physical Nature of Being: More with Dr. Selmer Bringsjord

<https://mindmatters.ai/podcast/ep289>

Pat Flynn:

Hello everybody and welcome back to the podcast. I am your host today, Pat Flynn, and I'm joined once again with Dr. Selmer Bringsjord to continue and complete our discussion on his wonderfully provocative article titled Mathematical Objects are Non-Physical, so we are too. You can find this in the excellent volume minding the brain. I'm sure the relevant links will be provided in the show notes. If you have not already secured a copy of that, I would highly encourage you to do so. It is a truly excellent volume and I've had the great fortune to discuss the contents of that volume with many different authors. And I have to say I was really excited when I first got this book because I saw this article in there and as I mentioned in part one and part two, I'm a huge fan of James Ross, who is one of the contemporary philosophers who really brought this argument back onto the scene.

Pat Flynn:

I say back onto the scene, I think it is an argument that has very ancient roots. And so I was very excited to see another development and defense of this argument because I've always found it to be just exceptionally forceful. It's a little technical, but in the first two parts of this conversation, we did our best to simplify the argument and summarize it. And of course I'm going to encourage everybody to go back and listen to those. I think they're necessary listening if you want to get the most out of this final part of our conversation. But to just quickly summarize what we've been up to, we've followed James Ross in arguing that all formal thinking is determinate. And by that, we mean that it is unambiguous in its conceptual content, but no physical thing can be determinate.

Pat Flynn:

And so, Ross is going to argue and many others following him that formal thinking, it's not material. It's non-physical. And then what we're going to argue here is that this argument can be extended not just beyond the nature of mathematical objects or formal thinking, but we can push it further to show that the human person, or at least some aspect of the human person is immaterial as well. So Selmer, welcome back. Thank you again for joining me. It's great to be able to continue this conversation with you.

Selmer Bringsjord:

Thanks, Pat. It's, needless to say three times over, a pleasure to be here with you.

Pat Flynn:

Yeah, it's always fun because we're recording it all at once, but a listener might listen to it over a couple span of days. All right, so very quickly, just so we can hop back into this efficiently, Selmer, I just want to reiterate again, and this is something that we said several times in the first two parts of the conversation, is that step one or phase one of this argument is itself hugely significant in terms of the grand scheme of things or the metaphysical big picture. Because if we've shown that there really is some immaterial aspect to reality regardless of philosophical anthropology or what we think the human person is, that itself is hugely significant and it's something that should not make any physicalist or materialist comfortable.

Pat Flynn:

So I just continually want to emphasize that because seen a number of people attack this argument at this second stage, but to me, the horse is already out of the barn. If you're a physicalist, you really need to try and shut this argument up a lot sooner. To me, the second step, the, "So you are too." Concerning immateriality, I think it's right. I think it's ultimately correct, but I think it's sort of icing on top of the cake so to speak. So I'd like to just get your general thoughts on all that, Selmer, and then maybe we can begin to explore that second stage in a bit more detail.

Selmer Bringsjord:

No, that's right. As long as the icing is sufficiently delectable, I'm perfectly happy saying that the result of step two, which is of course, that we are non-physical, is indeed icing on the cake. I think that fits. I think we really have three phases because if we agree with Ross and company, not only contemporary company, but those going back historically, as you point out, we would say that thinking is non-physical or immaterial and that's awesome and that is needed here, no question. And we build on that argument. There are people who are willing to say that and not say at all that thinkers are non-physical. And this is something we talked a little bit about at the outset in step one. The people who hold this position, the late Dale Jacquette who was actually a philosopher of mind, someone I have great respect for, was someone in this position.

Selmer Bringsjord:

We had some conversations about this and he wrote a little philosophy of mind book that was really deceptively, it was lucid and pedagogically sound, but it was deceptively sophisticated because it conceded at what would be undergraduate consumption level. You know what, "I'm a property dualist. I think thinking and stuff like that is non-physical." We build on that and then we say next now the objects themselves that you're thinking about are going to be non-physical as stage two. And I think Jacquette and others are reluctant to say that trans-finite numbers, the set of all positive integers, if it's smaller, inference schemata algorithms. I'm not sure they'll all go there after they admit that this mental activity is non-physical. So Descartes would be unhappy. I'm not sure what Ross, I really am not sure. I think Chisholm, who Ross studied under, one of the reasons I went to Brown was to study under Chisholm, following essentially in the footsteps of Ross.

Selmer Bringsjord:

And by the way, Chisholm told me, "Don't do this, don't study that." It was a very harsh start because of sort of market forces that he was incredibly prophetic about. He predicted that AI, COGSI and all these things would connect to philosophy in much more vibrant ways and literally thank God I listened to him. But then the next step is what we're talking about now, which is that we are two phrase, what are we? Well, we have to be non-physical, so thinking's non-physical. The things we're thinking about therefore have to be non-physical for the right things. And then the thing doing the thinking about the things that are non-physical is itself non-physical. That's the chain and you're very right and much more in command of the history here than I am. But you're certainly right to point out that this is sort of a wonderfully luminous thread of thought that goes back a long ways. And philosophy has these things. That's what makes it great in many ways.

Pat Flynn:

And just to give people an idea of the historical roots of this, we've been using this triangularity example for a while to illustrate the general thrust of the argument where if we think about triangularity as such,

it doesn't matter how many triangles you draw on a chalkboard, any particular triangle that you draw is always imperfect, it's indeterminate, it's always open to alternative interpretations. It doesn't lock down the specific precise determinate meaning that we clearly have when we're thinking about triangularity as such. You think of medieval thinkers like Aquinas, he's got a sort of wider metaphysical system that he's operating in, but any material thing is any sort of bounded individuated entity. And his sort of link here as I understand it, is that when we think about triangularity as such, it can't just be any sort of particular physical thing in the brain, a brain process or anything like that because then you just run into the same problem.

Pat Flynn:

You're not going to have the thought, the understanding. And we are talking about understanding here, conceptual ideas, not perceptual ideas of triangularity as such. So this ultimately leads him is to say at least minimally that because of our intellect, the intellect is immaterial. Now again, this might mean that you have to take on wider broadly Aristotelian assumptions of understandings of matter, which are a little bit different than contemporary understandings. But I think when you link this sort of argument up with a system like that, you can build that chain out, I think pretty rigorously, but you do it in a little bit of a different way. So I just wanted to tease some of the historical roots that I see, but we're going to focus on your kind of contemporary formulation at this point Selmer. So yeah, please take it away and yeah, help us see the chain, build the chain all the way through for us, if you don't mind.

Selmer Bringsjord:

Yeah, sure. Before I do that, if it's okay, maybe I don't know if it's okay. So Aquinas, you mentioned Aquinas and totally on target. And Ross was, well, it's obvious to anyone who has some command of the background here that he was even sight unseen for the literature in question, a devotee and a scholar of Aquinas. But he once said, I remember Ross said once, "You know, the thing about Aquinas is he was ..." And I am paraphrasing, but I'm pretty close to the verbatim, "Aquinas was really smart. He was so smart that he would make inferential jumps that worked, but skipped over a lot of the details. "

Pat Flynn:

He's not that the only one that said that of Aquinas. Right? Yeah.

Selmer Bringsjord:

Yeah. Extremely charitable. So I mean, the irony here is really interesting. It's like, "Well, professor Ross, I thought we're in the business of working out those inferences." "Oh yeah, yeah, of course we are." "To see whether it's valid or not." Well, yeah, yeah, of course we are. And that relates to the argument we're talking about here. So he really did Aquinas and more than liked, he was a great scholar of Aquinas.

Pat Flynn:

Yeah, and totally. And I love that comment. I would just remind people that, again, Aquinas was operating with a very broad, very thick and rich philosophical system. There's also just sort of a lot of in the background for Aquinas that is just not in the background for us as well. And it's important to understand that and spell that out if you're going to get him or really any historical thinker, that's just sort of the necessary project that you always have to try to engage in. But that's a conversation for a different day. So yeah, back to you and your development of it.

Selmer Bringsjord:

Yeah, yeah, yeah. So the line of reasoning to get to the point which is the focus here in part three, which is going to be the final step. There are two in the paper, but we've opened our minds actually historically to at least three and more. But the final step is we are non-physical things. So the question is how do we get to the point where we want to consider how to make good on that final step? So in step one, to recap, we considered what could be going on, put barbarically here, what could be going on when someone truly understands a fairly meaningful algorithm that's active in the world, and we went with quick sort. Quick sort is the algorithm from Tony Haar that is still used today in its variants, takes an arbitrary set of numbers that are all jumbled up, integers, and then out comes a wonderfully arranged list.

Selmer Bringsjord:

Of course you have to decide what kind of ordering you want, but starting perhaps with the smallest up to the largest, we said we could do the same thing for numbers for, sorry, letters in the alphabet. English alphabet you throw in a bunch of, and it's more pragmatic perhaps to talk about last names, which we always have to sort. Jumbled up last names come in, sort them alphabetically and quick sort will do that. And this was really an amazing accomplishment that this algorithm was as efficient as it was. So how do you make sure someone understands the algorithm? And what we've seen is if they just, as you pointed out in connection with triangularity, same kind of deal. If they start saying, "Well, I got it because look, let me show you how I can work through the algorithm as they wrote it down on this piece of paper or on this dynamic web page with this little app going here to show you the process. So don't I understand it?"

Selmer Bringsjord:

Well, not so fast. Not so fast. You understand it when you understand the background deep concept behind all these physical incarnations, that's when understanding arrives. So we ask the question, okay, are we really sure we can't have understanding by virtue of the fact that the would-be understander stands in some relation to a particular artifact or embodiment? And we came to conclusion, no, there's no way. There's no way there's understanding there. So our explanation for there being understanding, which as you point out already packs a bit of a wallop for physicalists and materialists or not for them, but against their worldview is that no, there must be this thing out there, quick sort, this type if you will, this abstract type. So now we have to get down to brass tacks and figure out how we get from the result of step one, which is that to our nature as non-physical things, and by the way, I'm not saying we don't have physical bodies. This isn't some kind of absurd position that the physical world is unreal, or at least that we're not in part in command of bodies.

Selmer Bringsjord:

We have hands, we have brains, we can use them, but they aren't us. When I say we are non-physical, the result of the next step, what I'm saying is we as persons, we as the kind of amazing entity that can reason, deliberate, maybe we're not happy when we close our eyes, we think about the kind of thing we are ethically, maybe we're not happy with it. This kind of thing is a person. With Naveen, we're saying this is a non-physical thing. So we get the next step done by turning to by any metric, another titan in a different quarter of philosophy and that's John Searle. That's where this changes very much.

Pat Flynn:

Yeah. Real quick before we go to Searle, and I love that you incorporated Searle's thought into this, very influential and brilliant thinker in his own, just to maybe help people understand what we mean by

understanding, I'm going to use the triangularity example again. I'm just using boxes and boxes of triangles in this conversation. This, again, tracks all the way back to the classic distinction between perceptual and conceptual ideas or even the difference between imagination and understanding. So we don't image triangularity, we image particular triangles, but we understand triangularity and there's certain things that we understand that we can't even image. Think of a shape with a trillion sides, a tarragon. Right? We understand that, but try and form an image of it. Yeah, good luck. I can't do that.

Pat Flynn:

So not everything that we understand can be imaged. Right? And there's clearly the space of reasons is richer and wider than the space of physical causes, if you will. So this again is a distinction that's very ancient that this argument I think is leaning into, even if it's not always made explicit, but I think triangularity helps here again, because something that we very much understand, again, the sort of denial of understanding it as we argued before, is ultimately incoherent and unsustainable. But don't mistake that for thinking that we can image something. So I just wanted to put that clarification out there before we move further just may help people.

Selmer Bringsjord:

No, that's great for many reasons. Returning to triangularity however many times it takes is wonderful. And I've already said that, we looked at some of the reasons why. But in addition to what you just said, a sort of warning to people who would explore the territory we're in, I think is extremely wise that we can't assume we can image or draw or such. The thing that we understand, everyone actually already knows this, but few people perhaps take account of it at least with sufficient focus and clear mindedness. What I love about part of Leibniz's Modus Operandi, the great example here that fits what we're talking about is the calculus. So certainly filled with very interesting logical mathematical objects that if you take them away from the human race, they aren't going to have a mission to Mars that works. Right?

Selmer Bringsjord:

This is for motion for things that move and change. Also true of economics, where you're looking at phenomena that don't center around physical things moving, you take these objects away from the thinking of the agents doing the stuff, you're dead, including the computational artifacts that process the information. But the amazing thing about Leibniz's stuff and the calculus is his notation has not only survived, it is the staple. I could walk into any, here I am, I'm on campus at RIPI. If I go find a calculus class, I can walk into that class and in the worst case scenario, look at the textbook sitting on someone's desk and say, "You mind if I open this up and show you how Leibniz imaged and drew some deep concepts that are in the background of the calculus?" Now, in his case, what's great and connects to what you said is that the notation is suggestive at a conceptual level.

Selmer Bringsjord:

It's not that the image is the thing. So Leibniz was great because just crazy publication plan in his case here. Fortunately we're talking about people who were wise enough and organized enough and smart enough to get a book together. Leibniz would just write this stuff down on pieces of paper, wonderful drawings and stuff and notation and throw it into a cabinet, and that was the end of it. So thankfully we have his correspondence and people are still excavating what's in the cabinet, but you're very right about that, what you say about understanding in the sense that specifically it can't be counted upon this thing that you understand to be imageable and drawable.

Pat Flynn:

Yeah, yeah, yeah. Very good. Glad we just got that pointed out real quick. So Selmer, all right, so to swing back around then, you just brought up Searle.

Selmer Bringsjord:

Yeah.

Pat Flynn:

Explain how you incorporate some of his famous thought experiments in the development of your argument.

Selmer Bringsjord:

Yes, sure. Well, there are people who have forgotten in AI today, or maybe they've willfully forgotten that in 1980, Searle caused an absolute earthquake, because of my age. I was there, I experienced it in AI, in artificial intelligence, the field devoted to roughly building computational artifacts. That's where the artificial comes in that have a level of intelligence that at least approaches what we see in the human case. So Searle says, okay, in this 1980 paper, he says, "Okay, we understand a lot as humans. Kind of a big deal for us. That's how the whole educational system works." I think we have this overused overheated term now in AI machine learning, this machine learning, that deep learning, everything is learning. Actually none of it is learning in the sense that has been in use for conservatively, explicitly in use for 2.5 millennia. "So that's a big deal." Searle says, "So let me ask you, so how does that actually work when we look at a machine? It seems to me that what your computers are doing as just moving symbols around, and I don't see how that could possibly ensure understanding." And that's just gentle.

Selmer Bringsjord:

And then he gets much more forceful and says, "I can show you that it doesn't amount to understanding. I will become a computer and I will manipulate symbols." And this is the famous Chinese room argument. And when I, in the room, manipulate symbols according to what he calls a rule book, it will seem to people on the outside, for example, that I actually am a competent understander of Chinese. Because in could come what he calls squiggle squoggles, the rule book says, do such and such with it, send back a new sequence of squiggle squoggles. People on the outside are like, "Wow, this system understands Chinese." And Searle would say, "No, it doesn't. I don't understand any Chinese." Which was of course actually true. And far as I know is still true. He understood no Chinese. So this was the Chinese room argument. People not only forget about it in AI by the way, but forget that it appeared in behavioral and brain sciences, which has a venue type that allowed Searle to respond to all the critics. And he was just, by any metric, I think his replies were really quite devastating.

Pat Flynn:

Yes. Yeah, I think it is another very, very powerful argument that has withstood the test of time and many, many objections. Right?

Selmer Bringsjord:

Oh, absolutely. So all right. We have to ask ourselves about Searle's argument. So we affirm the argument, I've defended the argument repeatedly in print. I've actually defended it against some objections that Searle did not consider with some colleagues. So there's a very much a sort of traditional

scientific methodology going on in the article we're talking about because it's building on the shoulders of folks coming before, and that's part of the essence of science. So all right, let's say as Navina and I both say, and we both firmly believe, Searle has demonstrated that the mere manipulation of symbols incarnated in a machine really is patterns of ultimately at the lowest level patterns of electrical tokens that doesn't amount to understanding. All right? So let's think about where we are in the overall progression. If we have Searle on hand here and this lesson, what could possibly be going on if we understand some logical mathematical concept, if we understand modus ponens which we talked about, modus tollens which is in the chapter, if P then Q, not Q therefore not P, take that away. You eviscerate mathematics.

Selmer Bringsjord:

If we understand quick sort, the algorithm we've talked about, and if we understand triangularity perhaps in the main three geometries, Euclidean, Lobachevsky, and Riemannian. If we understand that, is it that we understand it because we manipulate physical objects quickly and process them? Is that where the understanding is? For the opposition, they don't have many alternatives here. So once we get to the objects are non-physical, but we understand them, we then have to ask the question, okay, well what does the understanding consist in? Tell me in the real world. I'm over here. Triangularity is out there, what's going on here? How does it actually work? And if that person tries to appeal essentially to what now is the dominant response, there's no question about this. The dominant response to this question today in 2024 is, "You moron, look at what's happening. Look at the world, look at it. GPT-4, you ask it a question, you get it back, you ask it a question even about some things of course in mathematics, what are you talking about? You get the answer back, it understands." No, absolutely, and you see in the cracks that have appeared.

Selmer Bringsjord:

Whenever the reasoning that's required to supposedly demonstrate understanding has got to be accurate and precise and formal. These things fall apart because they don't have in their data things of the sort that we're talking about. This is an empirical fact. Right? GPT-4 does not have the formal Euclidean definition of a triangle. It literally doesn't have that. It's got pieces of data that don't relate to the logic of what Euclid came up with and was dealing with. But nonetheless, all I'm saying here is for this article, what we have on the scene right now only makes the whole thing stronger.

Selmer Bringsjord:

There is no understanding on the part of these systems because of ultimately what Searle pointed out. Now, I tend to think Searle may not agree. I tend to think with a number of others in terms of the history of this argument and philosophy. And I have a TED X talk on this. So Leibniz said this before Searle. Right? Leibniz said, "Look, you know what? You think thinking-"

Pat Flynn:

His famous mill argument. Right? Yeah.

Selmer Bringsjord:

Yeah, exactly. Exactly. It's the same idea when you go into a mill, you see all kinds of physical stuff swimming around. A mill might not compete in complexity with what we've seen in a so-called deep neural network. It's the same thing. What you're going to see you go into it is little physical things being moved around extremely fast and that's it. But that doesn't give you understanding. So we can't be the

kind of thing like that since we do understand the objects, the concepts we're talking about. And that's it. That's not it. I mean, there's more to it and you've encouraged listeners to go to the article and I wholeheartedly concur, but that's it. So if you don't want to go to step two, your best bet is to knock down Searle. Okay, good luck to you. Because what's happening again in modern AI, contemporary AI is just, I don't know where Searle is. I don't know what he's doing. He's not perfect.

Selmer Bringsjord:

Chisholm, by the way. We talked about Chisholm. Chisholm used to quote Searle all the time because he'd say things like, "We need to have a level-headed. Let's just get a commonsensical, level-headed, take on the issue." And he rarely would quote people. He was fine with Aristotle. Okay. But Searle, in my recollection by far, got in more than anybody else. So I think Searle, at least in thought here, should be sitting around whatever, smoking a cigar, relaxing, and then looking at this crazy stuff that we see from modern so-called generative AI and thinking, "Wow, this is awesome." I just read this article over here where this person says, "I asked it a dirt simple question, and out of the blue it starts hallucinating and comes up with a cockamamie response." And Searle's sitting there thinking, "Yeah, because it doesn't understand anything about what you're asking."

Pat Flynn:

Yeah, really, really great points, Selmer, and I love that you brought Searle's thought into this because again, I think it's very powerful and it's a good way to sort of bridge that gap or finish the chain that we were talking about. But it does really connect, again, I think very well with the ancient thought. Aquinas's thought that there's two ways a form can be contained in reality. One is in a material way, a particularized way. We've talked about why that is a problem if we're going to try to make sense of formal thinking through that. But it also says that form could be contained in a way that form contains form, in an immaterial way. And that's what he says the human intellect is. Right?

Selmer Bringsjord:

Absolutely.

Pat Flynn:

And I think that Searle's arguments can offer some contemporary reinforcement for that. And if you have a sort of wider, again, I would say broadly neo-Aristotelian metaphysic, great. It all links up really nicely actually. It's like a really great systematic account where we can account for a huge range of phenomena and data with I think a very powerful philosophical system that has some really smart, impressive people behind it. And that stands up really well even today. So in terms of theory comparison, I think if you incline in that direction, you're in very good company. And you can use the thought of thinkers that might not themselves identify with these philosophical schools of thought to actually lend very significant support for them being Searle. Last I checked was not a theist or a dualist of any sort, but he offers some of the best arguments against reductive physicalism out there. So that's great. Right? The hostile witness of sorts. Right? Yeah.

Selmer Bringsjord:

It is great. And he was, during my time of consideration of the Chinese room argument at Brown when I was a graduate student, he was, Searle that is, was a mysterion really. He's like, "Oh yeah, it's biology. I'm not going to tell you that I'm a dualist or anything like that, but it's biological." In other words, it's

mysterious. Okay. That's not really something that has a lot of explanatory power. So I'm not a thoroughgoing Searlean, but I would say, I'd add

Pat Flynn:

He may be a property dualist. Sorry, I should qualify that.

Selmer Bringsjord:

But yeah. Yeah. And this is all really, no, that's very good. And this is really all about, and I mentioned Jacqueline earlier, Dale Jacqueline who was very much a property dualist, but he would not go to the much more important question for who we are, what can become of us, where do we come from, and what is the meaning of life is, "Look, if there's some exotic mental properties that are non-physical, I don't think we need to get too worked up about that either way, frankly." But the empirical support that Searle is receiving from the just total unreliability of symbol-manipulating systems of gargantuan size is to me also fascinating. So you're right about the additional warrants produced scientifically here by the theory having explanatory power, but you also have this empirical confirmation.

Pat Flynn:

It's funny you mentioned the empirical confirmation. I really do think the philosophical argument settled the matter so much that even if you thought that these machines had understanding, that wouldn't make me a physicalist. I would just say, "Okay, great, a sufficient level of complexity has been reached that we have a new induction of a form." Aristotle could totally go with that, right? It doesn't prove physicalism. It just means we discovered something really fascinating and then we can just run the same arguments for these machines. Right?

Selmer Bringsjord:

Yes, absolutely. But we both know-

Pat Flynn:

That isn't the case empirically. Right? Yeah.

Selmer Bringsjord:

Well, we both know that it's not how it will go in the mind of diehard non-objective proponents.

Pat Flynn:

But that's how it should go, I think. Right? So I would not rest this argument on those empirical claims. I don't think that ultimately decides the matter. I think you're right on how to assess the actual data of the matter. But if it were the other way around, I wouldn't suddenly swing against these arguments. I would just keep driving them through to the conclusion that... We've induced a new form here by creating something so complex that, and that's how the system works, that you can reach a certain degree of material complexity that disposes toward the induction of a new form if you follow in Aristotle. And that can happen. Right? I don't really expect it to happen.

Selmer Bringsjord:

It's going to happen. I think it's going to happen behaviorally in the minds of a large number of humans. They won't know that they're being Aquinas-ish by doing it, but they're going to invoke this new form and they may even worship it. It's going to happen.

Pat Flynn:

Oh, that's an interesting thought. Now, all that, of course is a conversation for another day, but I think it just may be an interesting thing to put in people's minds. Of course, I agree with everything that you've said so far somewhere, and I think you've done a really wonderful job in this article. And again, I have to encourage people to read it. There's a lot more detail in there, and you address other objections that are certainly worthy of consideration. And that volume, once again, to remind people of the title is called Mathematical, Objects, Are Non-Physical. And you are too. And you can find that in the volume Minding the Brain. I'm sure links will be provided in all the usual spots if you haven't already secured a copy of that. So Selmer, before we go, just any final thoughts you have. And I would love to also just further hear what you're currently working on and how people can just keep up with you and future projects and all that good stuff.

Selmer Bringsjord:

Yeah, thank you. This has been sort of really positively wonderful, and the connection that you've made to things outside of what we discussed in the article is really cogent, deep, and so forth. So I mean, I really mean that. I've done a few things like this, and my experience with Bob Marks was likewise. So thank you so much. As far as today, well, right now I'm puttering around with my academic home page and about to bring up something outside of my university, but it's easy to find Selmer Bringsjord, there's only one as far as I know in the world. And if you see anybody with the name Bringsjord, other than someone in my family, they probably just took it from the spot in Norway, which is still findable in Apple Maps as Bringsjord.

Selmer Bringsjord:

In terms of current stuff. Wow. My wife says that I should make some changes now that I'm 65, and I should work less and work on less projects. And all I have to do is mention one name that I actually, as far as I know, I've hardly ever mentioned in print, because I don't want to jinx myself. And that's Thoralf Skolem, a Norwegian logician. I think when he hit 65, he considered that some kind of graduation from high school, and now he gets some real work done. So I kind of feel the same way. And maybe tomorrow a meteorite will take me out, but I have a lot of projects going on, and if I had to pick some that are connected to what we've discussed, maybe just pick one or maybe two. Well, I had the great fortune to discuss one really deep logico-mathematical conundrum with John McCarthy, one of the founders of AI, modern AI in 1956, which was the continuum hypothesis.

Selmer Bringsjord:

And that had a huge impact on me. We actually had two conversations, and he was effusive, and he didn't like my stance on AI, but he liked the logic part, he loved that which is why he is willing to talk to me. I'm working on the status of some of the very objects that are alluded to in this conversation or conversations and the article in connection with machines, whether machines can convince us that they do have understanding, despite the fact that I know, and as you point out, they won't have understanding. So I'm doing that with a few really sophisticated things, including things that relate directly to the continuum hypothesis. So that's kind of abstruse. Another thing I mentioned two things.

Another thing is McCarthy did write a paper called Free Will Even For Robots, and it's short. He thought that AIs could have free will.

Selmer Bringsjord:

And I wrote a critique of that paper, and I had a chance to briefly discuss this with him. It turns out John McCarthy wrote, but never published, a fuller treatment of his position that AIs could have free will. It was a very deflationary sense. You can find it on his, he's deceased. Find it on his website, which is still maintained. So I'm working with help from people in my lab, you got to have a roboticist for this to try to make, in some sense, free will correlate to human level, free will somehow semi real in a robot. So it's not like I don't think AI is ... Chisholm told me, "Selmer now that you were settled on working in philosophy as it relates to AI." And that was the idea right after I graduated as an undergraduate.

Selmer Bringsjord:

So he said, "Don't do only formallogical. These logicians, if they're pure logicians, they don't get a job. You got to connect this to something." And he made me go back and forth. So when I came back, I think the third time I came up with computation and AI, because I knew a thing or two about it, he just said, "Yes, do that." And the next thing he said was, "How much have you had? How much have you studied?" And I said, "Well, I've only had two programming classes." He was like, "Well, I think you're going to have to rectify that." So then I started studying computer science at Brown. I took tons of compsci classes and so forth. So I want to see how close the actual AI engineering for a robot, a real physical robot, can get to implementing McCarthy's account of free will.

Selmer Bringsjord:

So now, warning, McCarthy was really, really, really smart and really well-read. So he did not think that his concept of free will was all the way there to what you and I enjoy. And I accept that a hundred percent. But there's a lot of interesting work to be done here, and I'm working on that, and I hope that it leads to some cool explorations. I hope that I can follow through on Chisholm's advice, which was basically, "Selmer, if you're going to start spouting off about AI, which is what philosophers do, you might want to start studying AI itself."

Pat Flynn:

Wow, very, very cool. And I look forward to keeping up with that, and I'm sure many of the listeners will as well. So really, really appreciate it. Selmer. This conversation has been an absolute delight. One more time. I want to remind people to, if they have not already, to grab a copy of *Minding The Brain*. Again, the links will be in the usual places. And finally, if you're enjoying this episode, if you've enjoyed this series, please consider leaving a review of the podcast. Whether you're listening on my show, the host of *Philosophy for the People*, or *Mind Matters*. It's always a joy to team up with the team over at *Mind Matters*, and they always throw such interesting guests my way, such brilliant people to talk with, and it's a pleasure for me. So Selmer, thank you again for this great conversation and want to encourage people to not only grab the book, but leave a review wherever you listen, wherever you listen to podcasts, that helps the show grow, helps spread the word, and hopefully we can do this again soon. So thank you all. Thank you, Selmer.

Selmer Bringsjord:

Thank you, Pat. Thank you very much.

Announcer:

This has been Mind Matters News. 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.