Robert J. Marks on Specified Complexity and Meaningful Information (Part II) (https://mindmatters.ai/podcast/ep112)

Robert J. Marks:

Roles are reversed today on Mind Matters News. Dr. Michael Egnor interviews me, Robert J. Marks, on whether a picture of Mount Rushmore contains more information than a picture of Mount Fuji. I get interviewed by a brain surgeon today on Mind Matters News.

Announcer:

Welcome to Mind Matters News, where artificial and natural intelligence meet head-on.

Michael Egnor:

This is Dr. Michael Egnor. Dr. Jeffrey Shallit, who is a mathematician in Toronto, claims that Mount Rushmore doesn't have any more information than Mount Fuji. I'd like to ask my guest today, Dr. Robert Marks, to answer that question.

Robert J. Marks:

In terms of meaningful information, I think it's obvious. Michael, they used to say that it doesn't take a brain surgeon to answer this or it doesn't take a rocket scientist. Well, it turns out you're a brain surgeon and I've done work for NASA and I got an NASA Tech Brief award. I guess that makes me a rocket scientist. So I think for both of us, the answer is obvious, yeah, that Mount Rushmore contains more information than does Mount Fuji. And it's clear from the context that this refers to meaningful information. There's more meaningful information on Mount Rushmore. There's Lincoln and Roosevelt and Washington, and yep, what do we get? Well, if we get Mount Fuji, we just get a big chocolate gum drop. So yeah, there's obviously more information on one than the other.

Michael Egnor:

And we had spoken in a previous podcast about different types of information. Can we say what type of information the additional information on Mount Rushmore is?

Robert J. Marks:

Well, yeah. This is an interesting question. We can ask ourselves, for example ... I'm going to give an explanation and then dovetail into the answer. We can ask ourself the definition of two DVDs, both of which have the same storage capacity. One has the movie Braveheart. one has just random noise in it. And both of them take out the same amount of bytes. Can we say that the DVD of Mel Gibson's Braveheart has more information than the noise? Yes, absolutely if you talk about meaningful information. And as we talked about before, it depends on your definition of information.

Robert J. Marks:

Certainly in the case of Shannon information or possibly Kolmogorov information, yeah, they're the same, but neither of one of those measures meaning. And so one has to go to specify complexity, the mathematics of specified complexity, specifically algorithmic-specified complexity. And I'll give a little pitch here in case people want to read more about it. It's in chapter seven of the book that I co-authored with William Dembski and Winston Ewert called Introduction to Evolutionary Informatics. And the cool part about the book is that it references a lot more nerdy papers that have been published in archival prestigious journals and conferences. You can read it there at a layperson's level, or you can dig deeper and go into the papers.

Robert J. Marks:

I believe that Dr. Shallit was thinking about Shannon information in the sense that a DVD of Braveheart would contain the same information as a DVD of just random noise. If you took a picture of Mount Rushmore and you took a picture of Mount Fuji and you stored them on your camera, both of them might have the same file size, if you will. And in that sense, they are identical.

Robert J. Marks:

One of the problems that we talked about before is people throw around the idea of information without really defining it. I hope that by defining it that we've made this clear, and I think clearly in the context of the statement about Mount Rushmore containing more information than Mount Fuji, that we're referring to meaningful information. I think that that's implicitly obvious.

Michael Egnor:

It's kind of interesting that Dr. Shallit was making the criticism, that he was saying that it wasn't clear that Mount Rushmore had more information than Mount Fuji, using his blog where he types letters and words that other people read. And there's no question that his blog contains more information than either a blank screen or just a screen with random typing. Even the very effort that he makes to deny that Mount Rushmore has more information about Mount Fuji is itself an example of something that has more information than something analogous to Mount Fuji.

Robert J. Marks:

That's a fascinating observation. Him making a statement is actually a self-refuting argument.

Michael Egnor:

And if these guys didn't have self-refuting arguments, they wouldn't have any arguments at all, because everything is self-refuting. You referred to specified complexity. And what is that?

Robert J. Marks:

Well, it's built on Kolmogorov ... I'm going to get, get a little bit in the weeds here. But Kolmogorov complexity is based on the shortest description length you can have of an object. The reason I really like Kolmogorov information theory is that it is the link to the physical idea of information. We know what mass is. We know what energy is, but what is information? What's a physical link to information? And I think that the description length is a good example.

Robert J. Marks:

To illustrate, imagine that we have a three-dimensional printer and we want to write a program. All three-dimensional printers need programs in order to operate. We're going to write one program prints a bowling ball in three dimensions. Then we're going to write another program which generates a detailed bust of Abraham Lincoln down to the detail of the wart on his right cheek and his upper shaved lip. I'm shaving my upper lip right now with my beard and everybody says I look Amish and I point to the fact that Abraham Lincoln shaved his upper lip, and so we would have to get Abraham Lincoln's shaved up upper lip.

Robert J. Marks:

And the question is, if we had the two programs, if we had the bowling ball and we had Abraham Lincoln, which program is going to be the longest? It's obviously the one of Abraham Lincoln because with a bowling ball, you say, "Print a sphere and put three holes in it." But with Lincoln, you would have to specify his lips and the beard and the mole and his eyebrows and everything else and it would be a much longer program. Therefore, Lincoln, the bust of Lincoln has more complexity than the bowling ball. And this is what Kolmogorov complexity measures in terms of information. The longer the program is, the more Kolmogorov information it has.

Robert J. Marks:

Now, the interesting part is that if you wrote a program to do a three-dimensional bust of Lincoln, I wrote a program to do a three-dimensional bust of Lincoln, one of our programs would be longer than the other one. Which one is the proper description length? Well, Kolmogorov complexity asks the question, well, there must be a shortest program somewhere that generates the bust of Lincoln. Whatever the length of that shortest program is is the Kolmogorov complexity of Lincoln. This is Kolmogorov complexity, which is a component of specified complexity.

Michael Egnor:

If you have three different systems, you have a bowling ball, you have a bust of Lincoln, and you have the atoms that would make up a bowling ball or a bust of Lincoln reduced to individual atoms and just distributed throughout the universe, just dust, which has the most information?

Robert J. Marks:

Well, again, yes, we are talking about Kolmogorov information, which is description length, description length in terms of the computer program that we're required to duplicate the object. It is that computer program that I'm arguing will be longer for a 3D printer to print a bust of Lincoln than a bowling ball. Now, there is the open question that how specified do you want to get as far as duplication. You want to get down to the atom? I would say probably not. You're just interested in the surface of the bust of Lincoln and the surface of the bowling ball.

Michael Egnor:

But wouldn't something with maximal entropy have more of that kind of information than a bust of Lincoln?

Robert J. Marks:

Yes, and this is where the rub comes in. Let's have a bust of Lincoln versus say a rock that we picked out in the driveway. Now, this rock might have a bunch of dimples and indentations, and it might have the same complexity that the bust of Lincoln does, but that's only one part of specified complexity.

Robert J. Marks:

Complexity is one part of it. The other one is specification. Why do we recognize that there's more meaning in a bust of Lincoln than there is in a rock that you pick out of your driveway, even though the computer programs that generate them are of the same length? It is that in terms of context, in terms of experience, the bust of Lincoln is more meaningful and this can be folded into the idea of Kolmogorov complexity in order to come up with this idea of algorithmic-specified complexity.

Robert J. Marks:

The idea of algorithmic-specified complexity uses the idea of conditional Kolmogorov complexity. And the idea here is that you bring into the interpretation of the meaning your experience. And so this would answer the question that I answered on the first podcast, or I guess I asked it, about the text in Japanese that I can't read versus the text in Japanese that a Japanese person could read that was fluent in Japanese. Well, for them, it would have more information because they had the context to interpret it.

Robert J. Marks:

Specified complexity has two components. It does have to have the complexity and then it does have to have the specification. Those two things combined give the overall measure of algorithmic-specified complexity, which measures the meaning of an object.

Michael Egnor:

From the standpoint of information theory, how does a bust of Lincoln or a statue of Lincoln differ from Lincoln?

Robert J. Marks:

Oh. Well, I think they're two different total worlds. For the bust of Lincoln, we're just only interested in the outside, the external surface. We're not really interested in what goes inside. I believe we're constrained with whatever the physics of the 3D printer is if we're printing it on a 3D printer and whatever it fills it in with, it fills it in with. We're only going to rest it in the external. That was the intent of my example.

Michael Egnor:

The thing is that that ... and it's certainly true, but I'd kind of like to take the observer out of it, meaning to say it let's not consider so much what we're interested in, but rather what the actual differences are. How does a statue of Lincoln ... and let's say that you made the statue in such a way that you also had a statue of Lincoln's internal organs. You tried to make it as detailed as you could. How would the most detailed statue that you could imagine differ from Lincoln himself, information theory wise?

Robert J. Marks:

Well, I think that in terms of meaning, Shannon information, Kolmogorov information, physical information where people talk about Landauer said, all information is physical, which is true if you go there, none of those address meaning. The only way to address meaning of which I am aware in information theory is to place it into context, into something which is meaningful. And that context must come through experience.

Robert J. Marks:

In order to read that thing in Japanese, you have to have the experience of having learned Japanese. I can show you a picture and you might say, "Oh, I see a couple of women and a boy there," and I look at the picture and I says, "Michael, that's a picture of my family. This is my son, Jeremiah and Joshua, and my daughter Marilee and my wife Monica," and so that picture is going to have more information than it would for somebody that has never met my family. In that sense, it is all based on context.

Robert J. Marks:

I'm not sure, for example, how an alien would come down, say a bulbous blob sort of alien with no form and look at a bust of Lincoln and think that it had any meaning. It would have to have the context of knowing what humans look like. And if it had no idea what humans look like, it would just sit there and flip its lips and said, "This is just like a moon rock."

Michael Egnor:

As I mentioned in the past, I'm fascinated by the traditional tonistic and scholastic definition of living things. That is that they are things that strive to perfect themselves. And what Thomas Aquinas meant by that is that there are purposes built into nature, final causes, what we call teleology broadly, and those purposes provide goals for things in nature, that things in nature tend to change in the direction of those goals. But what is unique about living things is that they act of their own accord to achieve their goals, whereas non-living things are acted upon, but don't act on their own accord.

Michael Egnor:

An example would be no matter how detailed a statue of Lincoln you make, the statue wouldn't be trying to make itself a better statue of Lincoln, whereas Lincoln tried to make himself a better man every day. What made Lincoln alive was that he was always trying to be a better Lincoln, better in terms of more fully realized or perfectly himself, as we all do. There's always a striving in living things. There's no striving in inanimate things. Statues don't try to become better statutes. They can only be made better by something external, but they don't try it themselves.

Robert J. Marks:

Yeah. This is interesting because this gets back to the idea of algorithms and the idea that everything naturalistically must occur in an algorithmic sort of sense. If you have a bust of Lincoln that you want to print, you're generating a computer program that follows a type of algorithm. What you are describing is this intent to better yourself is non-algorithmic. I would maintain that it is beyond the scope of naturalism, beyond the scope of information theory to capture, at least as I know it right now.

Michael Egnor:

The other connection that I think is absolutely fascinating here, the connection was drawn by the scholastic philosophers, is that there is a remarkable analog to this idea of things in nature, and particularly living things, striving to perfect themselves, striving towards a goal. And that is intentionality, which is a technical philosophical term that refers to the fact that all thoughts are directed at things. Every thought you could have is about something. If you think about it, you can't really think anything that isn't about something. What do we say? "I'm thinking about." Every thought has something it points to.

Michael Egnor:

And the implication there that the scholastic philosophers drew was that the tendency in nature for things to go to ends, to go to goals, and particularly the tendency of living things to perfect themselves, is the kind of thing that has to have a bind behind it. That is that there's no striving unless there's a more profound organizing mind that creates the striving. And so the ancient philosophers connected the idea of intentionality in the human mind to the idea of teleology and final cause in nature. And that was one of the reasons that, for example, was Aquinas' fifth way of demonstrating the existence of God.

Robert J. Marks:

Well, let me ask you a question. We're all familiar with robots, robotic mice finding their selves a path through a maze, for example, in order to get food, for example. And this again is a robot. Wouldn't these robots be attempting to ... I guess robots don't eat food, but they get some sort of reward at the end, a teleological award, of course. But wouldn't these robots be designed to improve themselves by getting better and better?

Michael Egnor:

Sure, but the design is externally imposed on them. There's no inherent tendency. That is that if you take a robot and you just plop it down in the middle of a desert and watch it for a while, if it does do things, it won't do them for long. That is that robots are conglomerates of silicone and copper steel and things like that. And you can leave silicone, copper and steel out in your backyard or on your desktop for as long as you want and it will never do anything that even remotely resembles a robot. The only way that silicone, copper and steel become robots is if a human being assembles them and programs and makes them do it.

Robert J. Marks:

Their intentionality is external.

Michael Egnor:

Yes. Yes. Yes. And that external teleology or external intentionality is what characterizes inanimate objects. They can't make themselves better in any way unless some intelligent agent comes along and pushes them and makes them do it, whereas an intelligent agent can make itself better without being pushed, or I should say not even intelligent, a living thing makes itself better. Bacteria make themselves better. And they're not intelligent in the sense that we think of intelligence, but bacteria make themselves better. But grains of sand that are the same size as a bacterium don't make themselves better.

Robert J. Marks:

What this reminds me of in computer science and artificial intelligence is the bias that is placed into artificial intelligence. There's some people that that would hope that artificial intelligence, for example, could filter out hate news. No, it's not going to be able to filter out hate news without having a bias from the programmer of what is hate news. And it is, I think, a firmly established fact through computer science theory, like the no free lunch theorem, that you cannot build a computer science without an intention, without a bias and that computer programs without bias are like ice cubes without cold. You just can't have them. we would expect intentionality in computer programs and artificial intelligence to always be programmed in by the computer programs. That's a good point.

Michael Egnor:

What terrifies me about artificial intelligence, and I don't think one can overstate this danger, is that artificial intelligence has two properties that make it particularly deadly for human civilization. One is concealment. That is that even though every purpose, every single purpose in artificial intelligence is human and all comes from humans, it's concealed. We don't really understand it. We don't understand Google's algorithms. There may even be a situation where Google doesn't understand Google's algorithms, but all of it comes from the people run Google. The concealment is very dangerous. We don't know what these programs are doing to our culture. And it may be that no one knows, but they are doing things.

Michael Egnor:

And the second problem, which is an enormous problem, is one that Rene Girard, who's a French philosopher, wrote about extensively and that is the concept of mimetic contagion. Girard felt that-

Robert J. Marks:

Say those words again. I didn't get it. Mimetic? I feel like I'm doing the interview now.

Michael Egnor:

Oh, that's okay.

Robert J. Marks:

Mimetic contagion. What does that mean?

Michael Egnor:

Mimetic contagion ... Girard was a literary theorist who was also a philosopher and I think one of the most brilliant men of the past couple of centuries. He was a brilliant man. And he felt that what made us human was our desire to imitate, that we are imitating animals and no other animal imitates anywhere near the way we do. And we imitate particularly other humans' desires. That is that for example, advertisers know this, that if they show a popular quarterback drinking a certain brand of soda, other people will want to go out and buy that same soda. But that's kind of an odd thing. Why would we imitate what that guy wants?

Robert J. Marks:

Because the guy looks happy and we want to be happy.

Michael Egnor:

Right. Right. Girard developed this remarkable system of sociology and anthropology based on this idea that humans are veteran imitators and they imitate desires. And he said that one of the most dangerous things that happens in human culture is what you call mimetic contagion. And what that is is it's a contagion of imitation. I imitate you. You imitate me. My neighbor imitates me, who imitates you and then you imitate my neighbor. And then this whole thing just becomes an explosion of imitation, which also can lead to jealousy, to violence. For example, if I imitate your love for your wife, that's a real problem. If I imitate your desire for coffee, that's not such a big deal, but when you start imitating other people's possessions, other people's significant others, then you have war.

Michael Egnor:

And so one of the problems with artificial intelligence is that it allows us to imitate others without even knowing what we're doing and it allows it to happen at the speed of light and simultaneously all over the world. That is that I can imitate a guy in China at exactly the same moment that everybody else in the world imitates the same guy and it takes zero seconds to do it. And that's never happened before. Humanity has never had that kind of interconnectedness. And that mimetic contagion, according to Girard, is lethal to mankind. I mean, we will destroy ourselves.

Michael Egnor:

And so I see the concealment of meaning in AI. That is that every fragment of meaning in AI is from human beings, none of it is from machine, but we don't see it that way. We don't even understand it very well. And it can happen like kerosene with a match. It could happen at incredible velocity and incredible ferocity. And these are incredibly dangerous things that we're dealing with. Frankly, I think that some of our political crisis in this country right now is because of that, is because of the bias inherent in our information that we're getting and enormous potential for invitation for imitation, for mimetic contagion.

Robert J. Marks:

I just watched a Netflix documentary called The Social Dilemma, which talked about the impact of social media and Google and all of the mining, the data mining that is done by these big networks that correspond to the concealed information that you talked about. And you're right. It's chilling. One of the things that they mentioned is that there's only two industries that refer to their customers as users, and that is social media and drug dealing.

Michael Egnor:

Yeah, sure.

Robert J. Marks:

And that was ... yeah. And so the impact that was made really makes me want to quit social media altogether, but I tell you it's addicting.

Michael Egnor:

Oh, absolutely. Absolutely.

Robert J. Marks:

So one has to do partial withdrawal. Maybe I need to go into a 12 step program or something.

Michael Egnor:

The problem is that they know it's addicting and I think probably one of the reasons that it's addicting is that they've made it addictive.

Robert J. Marks:

Oh yes, absolutely.

Michael Egnor:

And we don't even understand it. And frankly, they may not even fully understand it. That is that it's incredibly dangerous stuff, incredibly dangerous. It also has potential for good, but wow, the danger that we're facing, I don't think we comprehend what this means.

Robert J. Marks:

So concealment. The thing that bothers me about AI mostly is that it's unintended consequences. Certainly in social media, there's unintended consequences. But if you look at things like self-driving cars that kill pedestrians, and there's lots of pushback against the military using autonomous AI because it might not do something that it's supposed to do, and I think these are real concerns and they boil down to if you are going to develop AI, you better make sure that AI does what you intend it to do. That's another one I would add to the chilling aspects. I do like your idea of concealment also. I think that that is frightening.

Michael Egnor:

What concerns me a great deal is first of all, the widespread belief among people who engineer AI that AI has the potential to become conscious or to have its own intentions.

Robert J. Marks:

And it's surprisingly widespread.

Michael Egnor:

Oh yeah. Yeah. Frankly, they all believe it, or practically. And it's a collective insanity. I mean, nobody in their right mind actually thinks that a machine can think. The belief that a machine could think is along the lines of thinking that your television set is trying to communicate with you. No. The people who made the television program are communicating with you through the television set, but the television set isn't trying to do anything. It's just a piece of metal.

Michael Egnor:

And these AI and engineers are smart enough to know that, but they don't seem to, and two things scare me. Number one, that the people who are designing AI aren't smart enough to figure that out. And number two, that maybe they have figured that out and they're using it in ways that they're not being honest about, and both of those concepts are terrifying.

Robert J. Marks:

Yes. I do think that some of these testimonies about control of the masses before Congress are going to historically be revealed to be similar to the testimony of tobacco executives about the effects of cigarettes.

Michael Egnor:

Absolutely.

Robert J. Marks:

They know what they're doing and it's going to come out somewhere.

Michael Egnor:

Right. And I think that the primary motives have been to monetize it. Obviously, they want to make money. And frankly, I think that will always be the motive. I think they're just trying to be trillionaires instead of just billionaires. But the thing is that there are certain cultural and social structures that can be built that make it more lucrative. And that's very concerning. That is that there are certain cultural contexts having us feel that constantly acquiring new things is what will make us happy, as opposed, for example, to praying to God, but praying to God doesn't make them money, but buying new cars does so they push the buying the car. It's pretty scary stuff.

Michael Egnor:

We've been talking with Dr. Robert J. Marks from Mind Matters. It's been fascinating and a privilege. Thank you, Dr. Marks, and we hope you'll be listening to us again soon. Thank you.

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.