Robert Marks on Information and AI (Part I) (https://mindmatters.ai/podcast/ep111)

Michael Egnor:

What does information and how does information relate to biology and to evolution? That's the topic 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. And it is my privilege to have, as my guest, Dr. Robert Marks. Dr. Marks is the distinguished professor of electrical engineering and computer engineering. And he is the director of the Bradley Center. Dr. Marks it's a privilege to have you with us today.

Robert J. Marks:

It's good to be with you, Michael. It's interesting to be on the side that's being interviewed. So, this is going to be a new experience for me on Mind Matters.

Michael Egnor:

It'll be a little bit of reciprocity.

Robert J. Marks:

It is. Turnabout is fair play, as they say.

Michael Egnor:

That's right.

Michael Egnor:

So, we'd like to discuss today, information. And I know that information is a topic that you have a strong professional interest, and have a great deal of professional expertise. And, probably, the best way to start is to ask what is information?

Robert J. Marks:

That that's a profound question. It turns out that before talking about information, you really have to define it. And there's some interesting questions that reveal the ambiguity of defining information. For example, if I burn a book to ashes and scatter the ashes around, have I destroyed information? Does it make any difference, if there's another copy of the book? If I take a picture, everybody knows the pictures on your cell phones require so many megabytes of storage, am I creating information? So, these are kind of questions, the answers of which depend on your definition of information.

Robert J. Marks:

On another hand, if I'm given a page that has Japanese text on it and I don't read Japanese, which I don't, does it have less information to me than it does to somebody that has a native reading ability of Japanese? Claude Shannon recognized this. He recognized that there was different definitions, at least in academia, in the mathematical sense, if you will, of information. And he said, "It seems, to me," this is a quote of Claude Shannon. "It seems, to me, that we all define information as we choose. And depending on what field we are working in, we will choose different definitions. My own model of information theory was framed precisely to work with the problem of communication."

Robert J. Marks:

Claude Shannon, of course, was a guy that wrote the 1948 paper, which defined information. It was the first paper to use the word bit, a contraction for binary digit. And he laid the foundation for the communication networks that we use today on our cell phone. He was an extraordinary man. And the work that he did in founding Shannon information, had more profound effects on our lives than say, for example, the works of Einstein. He was just really a brilliant guy that worked for Bell Labs.

Robert J. Marks:

So, to answer your question, what is the definition of information, one has to go to the different models of information that there are. There is the Shannon model of information. And Shannon modeled information as bits, and was interested in how we used information to communicate.

Robert J. Marks:

There's also something called Kolmogorov information and Kolmogorov was a Russian scientist and his information was founded by three people, Kolmogorov, Chaiten, and Solomonoff. Kolmogorov got his name associated with it because of the Matthew principle he who has shall be given to because Kolmogorov at the time was the most famous of these mathematicians. But it has to do with structure and the description length of an object.

Robert J. Marks:

Then, there's the physical definition, a great physicist named Rolf Landauer said all information is physical. Now, I would argue with that. I think that all information is physical, if you constrain yourself to the physics definition. And this is in total contrast to the founder of cybernetics, Norbert Wiener, one of the great names of Norbert Wiener, it's a great name. He's famous for saying information is information. It's neither matter nor energy. We can think of, for example, information written on a book on the printed page. That's information that's etched on matter. But we also know that information can be etched on energy in terms of the cell phone signals that you receive. And so, it's not matter nor energy, but both can be places where you can place and represent information.

Robert J. Marks:

And then, the fourth ... the three so far is Shannon information, Kolmogorov information, physical information, and the fourth is specified complexity, specifically algorithmic specified complexity. The models of information that I have just shared with you don't really measure meaning. The purpose of specified complexity and specifically the mathematics of algorithmic specified complexity is a way to measure the meaning in bits of an object. And so, that's what information is. And, again, it depends on the field that you're talking about. And, typically, when people talk about information, they don't take the time to define information, and that needs to be done.

Michael Egnor:

If you consider the information that is put into a system by intelligent agency, is that one of those particular kinds of information, or could that be any kind of information?

Robert J. Marks:

Well, yes, it does turn out that one of the challenges with naturalistic processes is that they cannot be creative. And, therefore, if you see the act of something which is creative, creation for example, it requires external information to be infused in the process in order to guide whatever it is that's being designed to its final design. For example, in biology information is there everywhere. Shannon information is there. Kolmogorov information is there because the description of the human body would take volumes to write down. Specified complexity, algorithmic specified complexity is also there because the body and what it does has a lot of meaning. So, I think in all cases, the relevance to biology is really significant.

Robert J. Marks:

One of the interesting things about biology, and this is fascinating and it shows, again, the genius of our creation is that one of the things about information that Shannon showed is that if you had a digital representation, then you could communicate exactly. Continuous or analog computing, or processes it degrades. And if you took a photocopy of a picture of your mother, and took a photocopy of the photocopy, and a photocopy of the photocopy of the photocopy, et cetera in about 10 or 12 generations your photocopy would look nothing like the original copy.

Robert J. Marks:

One of the beautiful things about all of creation is that it uses DNA, which is digital. And we know we can take a digital image of your mother, and send it to your wife, who sends it to her sister, who sends it to her son. And each one of these pictures is exactly the same. There is no degradation. So, there is-

Michael Egnor:

That's fascinating.

Robert J. Marks:

Yeah, there's a beauty in the fact that our reproduction is guided by a digital process. This is the reason, for example, we no longer use VHS tape, which was an analog, very continuous process, why we went to DVDs, and ultimately to digital streaming. And it's the same reason we no longer use cassettes. We went to CDs, and ultimately to a streaming music on Spotify, which is all done digitally. So, it's a wonderful testimony to biological design that we are fundamentally digital in terms of our reproduction.

Michael Egnor:

Wow. Sometimes when it's difficult to define, or understand a concept it's helpful to imagine its absence. What would characterize a system that had minimal information in it?

Robert J. Marks:

Well, this dates back to the 18th century, and a mathematician named Jacob Bernoulli, who came up with the idea of Bernoulli's Principle of Insufficient Reason. And his concept was that if we know nothing, if we have no information, no a priori information about anything, the best that we can do is assume that everything is going to be uniformly distributed. It's a very useful principle. But in terms of absent information, if everything followed Bernoulli's principle, we would just be dust kind of spread around.

Robert J. Marks:

Mathematically Bernoulli's principle is exactly the same as maximum entropy. And we know that maximum entropy is what spreads the gas around the room. So, if there was no organization, if there was no informational structure that's where we would go according to Bernoulli and to thermodynamics.

Michael Egnor:

What fascinates me, I believe, is not only that living things quite obviously contain enormous amounts of information, but even just the ordinary laws of nature manifest information and manifest an intelligent cause that, I think, you can see information in snowflakes and the laws of physics, all sorts of things.

Robert J. Marks:

Yes, exactly. In fact, one of the interesting things is that things with high Shannon or even Kolmogorov information happen all the time. We know, for example, that the generation of a single snowflake, as you mentioned, requires a lot of information in terms of either bits or description length of Kolmogorov. And so, the question is, is that meaningful? We're getting back to the idea of measuring meaningful information. And the idea is no, because we know that it snows all the time. So, these improbable high informatic things happen all the time. I think a more interesting question is to ask, what is the meaning of two identical snowflakes? Then, all of a sudden, you get into the idea of meaning. Two identical snowflakes has a greater meaning than a single observation of a single snowflake.

Michael Egnor:

How does biological information differ from information in non-living things?

Robert J. Marks:

Well, I don't know if it does. I think that the information ... now, if we just talk about the functioning biology, talking about things in your field is a little bit outside of my wheelhouse, about talking about the mind and things of that sort. I do believe that after recent study that the mind is very different from the physical part of the brain. So, there's information that occurs external to the brain. But I think, in terms of just the physical materialistic definition, that most of the information can be used to measure what is in the biological entity.

Robert J. Marks:

We can talk about, however, creativity and where the idea of creativity comes from, which is the creation of information. And that is outside of naturalistic or information processes.

Michael Egnor:

Thomas Aquinas, following Aristotle, defined living things as things that strive for their own perfection. And he felt that that was what distinguished living things from non-living things. A rock doesn't every day, wake up in the morning and try to be a better rock. Whereas living things, to greater or lesser degrees of success, try to make themselves better at what they do. They eat, they rest, they interact with nature. They do things to make themselves even better examples of what they are. And it would seem, to me, that that might relate to the difference between information in non-living in living things. That the information in living things is directed to ends, it's directed to purposes that you don't see in non-living things in the same way.

Robert J. Marks:

Yeah, I would definitely agree with that. It does turn out that in order to do the improvement that you're talking about, there needs to be, if you will, a degree of creativity in doing that.

Michael Egnor:

Sure.

Robert J. Marks:

In other words, the entity has to have creativity. This is one of the things that we argue a lot about in artificial intelligence. Will artificial intelligence ever be creative? And I maintain, I feel it's almost like a mantra now, but artificial intelligence will never be creative. It'll never understand. And, currently, it has no common sense. So, these are materialistic attributes that are beyond the capability of artificial intelligence, which certainly can be applied to the human mind.

Michael Egnor:

Sure. I mean, I've always thought of artificial intelligence as just representation of human intelligence. And, in a sense, that the term artificial intelligence is an oxymoron that if it's artificial it's not intelligence. And so, intelligence must be human. And all the intelligence that's in computers, and computer programs, and machines is all human intelligence that is representative of those devices.

Robert J. Marks:

It is. In fact, what you mentioned is exactly the tests that Selmer Bringsjord, he is a professor at Rensselaer used in his test for whether or not artificial intelligence would be creative. His test was that does the computer program do something, which is outside the explanation, or the intent of the programmer? And there has been, thus far, no artificial intelligence that has done this.

Robert J. Marks:

Now, we have surprising results from artificial intelligence. I would say the AlphaGo that beat Lee Sedol in the Asian game of Go, the most difficult board game of all at one, at one point it did this surprising move and everybody went, "Oh, oh that's incredible. This was creative." No, that was not creativity. AlphaGo was trained to play Go. And that's exactly what AlphaGo did. It was playing Go better than human beings.

Robert J. Marks:

But we see this all the time. Computers play chess better than I can. I think was Emo Philips he says, "Yes, computers can beat me at the game of chess, but I have yet to lose a match in kickboxing." So, there will still be things that humans are better at. I always love that joke.

Michael Egnor:

So, what specifically is the relevance of information to biology, and why is information so interesting when it's applied to the study of living things?

Robert J. Marks:

Well, I think one of the things is that it's beyond the explanation of materialism. We hear a lot of thoughts today about Ray Kurzweil's singularity. This is where AI writes better AI that writes better AI. And pretty soon the AI, it is claimed, is going to have the intelligence equivalent to that of a human being. Well, of course, if we have AI writing better AI that is beyond the original intent of the computer programmer, then we have creativity, don't we? And creativity is beyond the ability of artificial intelligence. So, Ray Kurzweil's idea isn't going to work.

Robert J. Marks:

There's also the thought that someday we will have something which is called artificial general intelligence, which will be intelligence at the human level, including things like creativity, understanding and sentience. And this will be achieved someday. And no, this all requires creativity in terms of the computer program. And that will never happen.

Robert J. Marks:

I have a colleague, Roman Yampolskiy, he is at the University of Louisville. And all this talk about artificial general intelligence, he posted a post on social media on April Fool's day in 2016. Let me read it to you. Roman is quite a character. He said, "Google just announced major layoffs of programmers. Future software development and updates will be done mostly via recursive self-improvement by evolving deep neural networks." In other words, he said that this AGI had been achieved and this allowed everybody to be laid off at Google because they weren't needed anymore. And it was a great joke. But one of the interesting things he said in social media, he got a bunch of likes but also, interestingly, a few requests from journalists came in, and they said they would like to talk to them about this wonderful developing story. So, to non-experts the joke was not obvious. But I think that anybody that does artificial intelligence, it was obviously an incredible joke. And I think it's humor that illustrates this idea that no, we're never going to have this sort of computer generating software that generates better and better AI, isn't going to happen.

Michael Egnor:

Right. There's a computer scientist in California named Judea Pearl, who I find to be a fascinating guy. And he's, over the past several decades, really pioneered in the field of causal analysis, which is the ability to ascribe quantitatively, using mathematics, causal inferences in processes, which historically has been very hard to do. Statistical analysis were able to infer correlation, but they weren't able to infer causation. And he is trying to figure out how to allow machines to make causal inference, because to have machines that are worth using that have any kind of, what appears to be, autonomy the machine has to be able to infer cause. But, of course, all that inferred cause is human intelligence that's put into the machine. Very interesting aspect, I think, of information theory is the ability to infer a causal relationships as opposed to simple correlations.

Robert J. Marks:

Yeah, that's very interesting. Of course, the people that do word mining, and natural language processing are big into things such as correlation. And correlation is easy to do. Talking to my friend, Gary Smith, who was a professor of economics at Pomona College, he said, this is very dangerous. And it's going to be very dangerous in the area of data mining. He said because we're going to come up with inferences that have nothing to do with each other.

Robert J. Marks:

There's a great website called Spurious Correlations and they just show how, in big data, you can get correlations between totally unrelated causes. They have a graph, I'm trying to remember one of the exact ones, but they had a graph, something as people who died strangled in their bedsheets versus the cheese consumption of Wisconsin. And the curves lined right up. I mean, they were totally correlated, but clearly there was no causation there. An establish of causation, thus far, is going to have to require human intervention, yes. So yeah, I agree with you.

Michael Egnor:

It's kind of interesting because, right, the inference to causation seems to be a leap beyond ordinary algorithmic statistical analysis of something. And it's a leap, I think, that only human beings can make. And I think it's also a very big profit science, especially in medical research that correlation is certainly not, in itself, conclusive evidence of causation. And there's all kinds of errors made in medical research because it's assumed that correlates imply cause, which they don't necessarily.

Robert J. Marks:

And this is interesting because most of the journals that accept papers based on statistics are only into the correlations. They require an r value of such and such, meaning that the data corresponds to a high degree of correlation. But the problem is, of course, is we do have these spurious correlations. And there is evidence of repeated correlations that are obviously true. If you smoke cigarettes, it's going to kill you, that's obvious. But for first time sort of situations, you're not really sure. And this has led to the conclusion of one researcher that up to 90% of the papers published in the literature that are based on statistics are flawed.

Michael Egnor:

I mean, I would tend to disagree. I think it's much higher than 90%. I think, that's a real underestimate of the number of papers that are published.

Robert J. Marks:

Okay well, Michael, I was ready to argue with you with it, and it turned okay.

Michael Egnor:

Yeah, 90% is very conservative, yeah.

Robert J. Marks:

Which is really amazing. And that's the reason that today, coffee is good for you, coffee is bad for you. You hear all of these fleeting studies that are reported in the news as gospel and it can have a terrible effect. It can make you paranoid. And so, I try to ignore all of these any more since Gary Smith pointed this out to me, yeah, it's really terrible.

Michael Egnor:

And I actually think that this notion of Pearl's that obviously that causation is a different inference from correlation and that we need to develop a way of quantifying it, and of applying it to machine learning is absolutely fascinating. And I really think that it's a uniquely human thing to understand causes. Machines can crunch numbers, and do things that don't provide insight. But to get insight into what causes what, I think, requires a human being.

Robert J. Marks:

I think, I agree with you. I can see, for example, in natural word processing that one could look in an event which precedes another event. And if you saw that a number of times, you can infer that maybe one caused the other. But I wonder what the dangers are of doing something like that?

Michael Egnor:

Well, a good example, or the smoking example that you used is that smokers frequently will have yellowing of their fingertips because they're holding cigarettes. And they also, are predisposed to get cancer. And there's no question then, that yellowing of your fingertips correlates with having lung cancer, but it doesn't mean that yellow fingertips give you lung cancer. Or that they cause lung cancer. You have to get the causal arrows right. Yellow fingertips and lung cancer are both caused by a common factor, which is smoking, but they don't cause each other.

Robert J. Marks:

I remember a similar story about the ice cream consumption and murder rate in New York City. That the ice cream rate would would increase, and then the murder rate would increase, not realizing that they were both just related to the rise in temperature. And after people got tired of eating ice cream, they went out and killed each other, I guess, is the bottom line.

Michael Egnor:

We actually, when I was a resident in neurosurgery in Miami, I was there during the drug wars. And so, we would get gunshot wounds to the head coming into the ER constantly, except they would always stop when it rained. Miami has quite a bit of rain, so when we would have an hour or two of rain, the ER would just go completely quiet. Nobody would come in. And then the sun would come out, and then people would shoot each other again. And it was fascinating. But people wouldn't shoot each other during rain.

Robert J. Marks:

So, the causation would be the good weather causes people to kill each other.

Michael Egnor:

Yeah. Or that there's something about being wet that protects you from gunshot wounds, which is kind of interesting. Yeah so, the causation thing, is a very subtle question.

Michael Egnor:

Well, thank you, Bob, for that information. We've been talking with Dr. Robert Marks about information theory and, particularly, its relevance to biology. And thank you for joining us. And we will see you again. Thank you.

Announcer:

This has been Mind Matters News. Explore more at mindmatters.ai that's mind matters.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.