Artificial Intelligence: Navigating the Hype, Limitations, and Ethics
https://mindmatters.ai/podcast/ep377

Robert J. Marks:

Greetings and welcome to Mind Matters News. I'm your host, Robert J. Marks. We're talking to Dr. Donald C. Wunsch 2.0 or the second, as the old people would say. He is the Mary Kay Finley, Missouri Distinguished Professor of Electrical and Computer Engineering at Missouri University of Science and Technology, where he is the director of the Comer Institute Center for Artificial Intelligence and Autonomous Systems. He's won a lot of awards. He's published a lot of papers. And when I say I've published a lot of papers, usually I say some of them are good and a lot of Dr. Wunsch's papers are good. So look him up on Google Scholar and you'll be able to see some of his papers. He's also a recipient of many awards. We mentioned he was a fellow of the IEEE. He's also received a neural networks pioneer award from IEEE. There's a sister society called INNS, the International Neural Network Society. And there he has received two awards, the Gabor Award and the Lovelace Award.

Dennis Gabor, by the way, was a Noble laureate and the prize is named after him. And I think he won the Nobel Prize for the invention of the hologram. Lady Lovelace was the first computer programmer, and she wrote computer programs in the 19th century for the mechanical computer that was dreamed up by Babbage. And she had some very fascinating things to say about computers. She said, for example, that computers can never be creative. They can only do what they're told to do. So Don is the recipient of both the Gabor and the Lovelace Awards. So Don, welcome. And it's good to continue talking to you.

Donald Wunsch:

Thank you, Bob. It's good to be here.

Robert J. Marks:

It is good to be here. With me, it's getting more and more common to be anywhere. We've talked about artificial general intelligence. You had a curious take on that. You said that it really isn't a viable name. It isn't something that we should consider. Did I get that right? Kind of summarize that in a little paragraph, what you meant by AGI and what your stance is on AGI. Remember I asked you the question and you rejected the premise.

Donald Wunsch:

So a couple of things about that. One is that one thing that we are pretty good at, I won't say that we humans are great at, but we are maybe the only creatures that are pretty good at it, is dealing with paradox. And so a paradox is certainly around the idea of artificial general intelligence. So I'd like to talk about two folks who have weighed in on this topic, both of whom I admire and who have different opinions about it. So Yann LeCun, winner of the Turing Award, he has come out saying that AGI is unfortunately a term that we have to deal with because it's become so popular, but it's really nonsense because our intelligence is not general. We have all developed intelligence. We, by that, I mean humans and animals have developed intelligence based on the needs of the creature.

And so my first PhD student, Daniel Procrav, and we were roommates for a while. And I remember one time he was joking about what he had found about a dog brain. He had found some article about dog brain and it showed how much of a dog's brain is devoted to smell. And his joke was that a dog is a walking nose. And my spin would be that a dog is a genius of olfaction. Well, I personally know and have interviewed a human genius of olfaction, Walter Freeman, who developed the neural network models of

olfaction or the processing of smell. And it's a very sophisticated capability. But if we humans had as much brain devoted to olfaction as dogs, we could probably smell things a mile away that we cannot even smell in the next room right now. And if we had the sensory capabilities, that adds to it.

Robert J. Marks:

So if I had a dog brain, I would smell good. Is that right?

Donald Wunsch:

You would smell well. You might not smell good, but you would smell well.

Robert J. Marks:

Okay. Yes.

Donald Wunsch:

And then for the opposite perspective, Demis Hassabis, the most recent neural networks winner of the Nobel Prize, there have been four. One of them did not win it for neural networks. Leon Cooper won it for semiconductors, but he was on the board of their National Neural Network Society and one of the early neural network startup company founders. But most recently, Demis Hassabis, John Hopfield, Geoff Hinton all won the Nobel Prize. And I would have to put Hassabis first among those, all the prior recipients of Nobel Prize related to neural networks. And in fact, maybe among most Nobel laureates, because I think we will live longer because of the alpha fold for protein folding predictions.

Robert J. Marks:

Let me ask you, Don, I believe that there are other people other than Geoff Hinton or John Hopfield that were more deserving of the Nobel Prize than they were. Unfortunately, a lot of awards today are given by promoters, people that promote themselves and such. What do you think about that? Do you think there's people more deserving of the Nobel Prize than Geoff Hinton and John Hotfield?

Donald Wunsch:

Well, I would say that it's unavoidable that some great talents get overlooked and the field is full of stories of that. I think that there are ... My best prize that I've ever received is the Pioneer Award, and I personally know people who I believe are more deserving than me that don't have it yet. And so there's always an element of that. But I would say that I've met John Hopfield and know his work well. That was the first paper that I read that brought me to the field. I know the work of Geoff Hinton is phenomenal. He has run rings around me in terms of the impact of his publications. But yes, there are people whose contributions are just absolutely stellar who have not yet been recognized in that manner.

Robert J. Marks:

My choice for the Nobel Prize would probably be Paul Werbos for inventing of backpropagation with his Harvard PhD dissertation. And Bernie Windrow for his breakthrough work in the 1960s at Stanford, I mean, he did incredible things with neural networks back then.

Donald Wunsch:

Both of them are close friends. Yeah. So I would put them up with anybody who's ever walked the face with the earth in the field. And another one that I would list is Steve Grossberg.

Robert J. Marks:

Steve Grossberg, yes.

Donald Wunsch:

They are giants, and all three of them have been a tremendous influence on my trajectory.

Robert J. Marks:

You mentioned you got an award and you said you didn't deserving of it. I have a good-

Donald Wunsch:

No. No. I didn't say that. I said there were people even more deserving.

Robert J. Marks:

Oh, you said they were more deserving.

Donald Wunsch:

Yeah.

Robert J. Marks:

So I got a line for you. I received an award one time and I, of course, wanted to present myself as a humble sort of person. And I says, I don't deserve this award, but I have lower back pain problems and I don't deserve that either. That's a good way to accept awards in the way that I think.

Donald Wunsch:

So I'll contrast with this with the field of fuzzy logic. In the field of fuzzy logic, it's very easy to point to one person. And Lotfi Zadeh-

Robert J. Marks:

Lotfi Zadeh. Yeah.

Donald Wunsch:

Was the father of fuzzy logic. And so anybody that works in the field agrees. They might not agree on anything else, but they agree Lotfi Zadeh was the father of Fuzzy Logic. And I knew him too. You can see my interview of him together with Walter Freeman, who I mentioned earlier. Those are available on YouTube. And these are giants and we're very lucky to live in a time when it's as if technological progress was compressed. I still think that we're in the infancy of AI and even in the infancy of computers. I used to say we're in the stone ages of computers. We may have advanced to the bronze ages of computers. There's a lot left to do. But the field of computing has moved so fast that a lot of the giants have overlapped. I'm about to turn 65 in two months. So in the short span of time that I've been in the field 40 years, there have been so many of the giants that have overlapped with that period of time. People like Claude Shannon who wrote perhaps the most influential master's thesis of all time saying that these ... Yeah.

Robert J. Marks:

He did digital logic. He worked for Bell Systems. Didn't it revolutionize the way the Bell systems hooked up their connections?

Donald Wunsch:

Yeah. Well, his doctorate was one of two people that revolutionized information theory, but his master's was to say, "Hey, these switching circuits that you use for telephones, you could do Boolean algebra with them too, so you could use them for computations." And so the idea to use electronic switching circuits instead of the Babbage mechanical devices for computing, so you could call him the founder of electronic computing. So that's a pretty good master's thesis. Most people would trade their doctoral dissertation for that.

Robert J. Marks:

He was also a member of the Dartmouth Conference in 1965 where people like Marvin Minsky and John McCarthy and Claude Shannon and Ray Solomonoff got together. And that's where the word artificial intelligence was created. I think it was John McCarthy that did that.

Donald Wunsch:

Yeah, that's another one, by the way. Minsky came to the University of Washington and gave a talk while I was a graduate student and I attended that talk. So there are a lot of good examples of people who have contributed to the field that are still with us. And so Eric Kandel, Principles of Neural Science, he came to the university and gave a talk. And I attended that talk. He gave that talk to the medical school and that giant tome on principles of neuroscience. He later got a Nobel for that. So anyway, the names that we have mentioned so far, not Babbage, of course, not Lovelace, but many of the names that we mentioned were alive during this period and many still are alive who have made just towering contributions. It would be like knowing Einstein and Maxwell and Newton and Kepler. So it's just a marvelous period to be alive in this field and we can all be very grateful for that.

And so we do need to realize in neural networks, there are many people who have made ... And many ... I mean, still relatively small handful compared to the thousands or even millions of AI enthusiasts, and there will be many more who come because of the many great minds who've been attracted to the field. But a lot of the founding contributions have been made during our lifetimes, and many of those people are still with us. So it's really just an amazing moment in history.

Robert J. Marks:

Yeah. I remember, let's see, who won the ... Gosh, Roger Penrose won the Nobel Prize for his work in Black Holes, but unfortunately Stephen Hawking had passed away. He should have got the Nobel Prize in tandem, but you have to be alive to get the Nobel Prize. So yeah, I think that a lot of things are happening. I do maintain though that this exponential explosion in artificial intelligence that people talk about Moore's law and things of that sort, exponential increases are never sustainable. And I think that eventually the curve has to level out. When it levels out and how it levels out, I don't know what that is.

But look, I want to talk about your paper. You wrote a paper called Artificial General Intelligence is Nowhere Near, Artificial Specific Stupidity is Already Here and Policy Matters. And one of the things you do, you emphasize autonomy as a critical, and if you will, a missing component of what we're calling AGI. Why is autonomy such a hard problem compared to the pattern recognition or large language models?

Donald Wunsch:

Well, and in fact, autonomy is something that the field is attacking. There's a lot of research targeting agentic AI and tools being even rolled out, some on a trial basis and some that you actually can use. But I do think that this is a major limitation. And so in my paper, I say that insect intelligence is ahead of the field of artificial intelligence in terms of these autonomous capabilities. Because if you released a fly into this room and I were trying to kill it, it would be a pretty fair match. Despite my much larger brain, the flies' agility would serve it pretty well. And if I keep the door closed eventually after a large amount of effort, I would succeed, I might break some things in my office in the process. So it's amazing how much an insect can achieve with such a small brain, and we're just not that good at dealing with that.

And so that comes back to my criticism of the term artificial general intelligence to say that intelligence that has been successful, both artificial and natural, has become specialized to a task at hand. And the counterpoint would be Demis Hassabis has explicitly stated artificial general intelligence as a goal, but he has wisely taken task after task after task to rise up a chain to move towards those objectives. So he started out with games and then made general game engines that could learn a set of different games, but that's still a very constrained environment. And then the problem for which he won the Nobel is a great example of the success of ... So I think it's not a bad thing to have a goal that is reaching for well beyond our grass, reaching for the stars of quixotic effort, but I think it helps to have the humility to say this is a quixotic effort and not to proclaim victory prematurely. There are many examples in history of a premature claim of victory that backfire.

So I don't think that any of the claims that we have already achieved AGI are anywhere near true. And also, I don't think the projections that we will achieve AGI by 2030 are true either. There are such simple failure modes of these systems that ... We have failure modes too, but still there are failure modes of these systems that humans would simply say that's not intelligent at all. And yet there are things that these systems that can do far better than humans, but when my dad brought home a calculator in the '70s and I had fun calculating factorials, well, it could calculate factorials far faster than I could. You just put in a number like 99 factorial and it gives you, in scientific notation, how big that number is. There are all sorts of things that computers can do better than us, but there have long been things ...

So the John Henry song of the driving the railroad spikes in a machine that could do it, but even having a railroad in the first place. There are things that machines can do that humans can't do, but humans created these machines. And the same is true with computers. There are things that ... You can design a computer to do something that you can't do, but there's still a lot of things that you can do that the computer could never do. And so if the building catches on fire, if the building catches on fire, you leave the computer behind and you run out of the building. Too bad for the computer.

Robert J. Marks:

Yeah. So speaking of things that computers can't do, you talk about Moravec's paradox. What is that? And why do AI systems still struggle with this problem, Moravec's paradox?

Donald Wunsch:

Well, so I'm quoting from an article by LeCun and I cite it. And as I recall, it's the idea that you do something and when you succeed, it's instantly, okay, that's not AI. That's a nice thing you did, but that's not AI. And I might be getting that wrong, so I'll follow the citation. But basically that is a recurring theme. So for example, when I entered college, they had just gotten chess machines that could play at the master level, but there were little hacks you could get to every now and then beat one. For example, they underestimated the value of pushing upon towards a queen. So they could do great tactical levels and be well ahead of you. But if you could slip a pawn down close to the sixth or seventh rank, you might be able to force them to trade a rook for that pawn because you're threatening to queen it and

you might be able to get some advantage. I'm nowhere near master level, so it was hard for me to exploit that very often, but still I noticed that these little weaknesses around the edges of such systems.

But still to play at that level was considered a goal of AI. And then fast-forward to '97 and IBM trots out at great expense, Deep Blue and beats the human world champion, Gary Kasparov. Well, at the time I was already very interested in the Game of Go, and I literally thought it would take a century to get there to beat a human champion in the Game of Go. And I would encourage people to watch the documentary Alpha Go. And there's another one that came out just a few weeks ago and the name escapes me. I'll put it on my website. But both of them involved the work that was done by DeepMind, led by Demis Hassabas, another major technical contributor was David Silver. He collaborated with the Turing Award winner, Richard Sutton, and they developed a computer that would use reinforcement learning to self-improve. It didn't take anywhere near as long as I thought. So about a decade ago, they beat one of the top human players, and then subsequently they beat the then reigning human world champion in the Game of Go.

And that is so impressive because it has a lot of challenges that chess doesn't have. So the number of board states is much larger. The non-local interactions are sometimes even harder to see, and the evaluation of a board state is much more challenging. So it's NP hard just to evaluate a board state, and it's peace-based complete to attack the problem of the Game of Go. And this is to say, if every chess computer, if you had billions of the Deep Blue systems all reprogrammed to attack the Game of Go, you would not solve it by brute force. But what they did still had a lot in common with brute force techniques, but the reinforcement learning tools that were applied were very powerful.

Robert J. Marks:

Incredible.

Donald Wunsch:

And very well-designed. And it was a major breakthrough, major breakthrough.

Robert J. Marks:

Emo Phillips had a great story. He says, "Yeah, a computer could beat him at chess, but it was no match for him in kickboxing." Talking about swarm intelligence, backtracking a little bit. I met a guy at Stanford. They called him the fly man. His total job was to explain the way that the fly works. And one of the big questions that he asked is, "How does a fly land upside down on the ceiling?" And he was working on that. I don't know if they solved that, how flies can land upside down on the ceiling. You talked about things that insects can do that are really remarkable.

Donald Wunsch:

Right. That's an example that I like to trot out as well. And so you might recall that I had Robert Pinter on my dissertation committee. He had a joint appointment between electrical engineering and zoology, and he was an expert on insect vision and insect modeling. And so we have an expert on that topic here, and he knew about Pinter's work actually when he came here. So anyway, that's a fascinating problem. You might remember Karen Haines from the IEEE Neural Network Council. She was coordinating student volunteers.

Robert J. Marks:

She ended up in Australia, didn't she?

Donald Wunsch:

Right. But not before she did a PhD with Tom Caudell. She wrote her dissertation in the ... So I think it was Castleson Advisor Masters and she wound up working with Tom Caudell on her PhD and she did work on fly vision. And so what flies can do is really amazing. And the scenicuanan is the ability to land and hold on upside down landing on the ceiling. And yeah, I have yet to see an AI-driven system that can do that. That's not the same as if you made a drone copter that could reverse the force of the propellers and force itself to be on the ceiling. The fly can land on the ceiling and hold on. And so that's a remarkable control problem. There's all sorts of stuff like that.

So I think we need to have a lot of respect for natural intelligence. It goes a lot further and a lot deeper. And also it's more parsimonious than anything that's being done in AI. So having systems that can do so much with so little. So people are talking about tens of megawatts, hundreds of megawatts, even gigawatts for up and coming data centers, but we make do with 20 watts. And so prior to LEDs ... This was the amount of wattage of a nightlight. So dim bulb is no longer an insult to one's intelligence. It's actually a great compliment. What we can do with our dim bulbs is really remarkable.

Robert J. Marks:

It is. Yeah. It is incredible. Let me ask you another question about AGI. And you made this point in your paper, and I've been thinking about it, and I agree with you. I think that 90% of the articles that are posted on the internet from these news services and stuff belong more in the National Enquirer than they do as a news story. And you raised the question, is AGI more of a marketing term than an engineering target? I'm not even sure who coined the term AGI. And what do you think? What do you think?

Donald Wunsch:

Yeah. I cited a book by the title of AGI. And so I think it's the first citation in my paper. I browsed that book. It's good as why I cited it. I speak about it favorably in the paper. So I think that's the best source for looking at the various definitions, the history of it. I also quote from the OpenAI definition of AGI. And some of this actually, it's not just marketing, there's also some legal ramifications. OpenAI recently, I think renegotiated ... That might not be the right word, but they have clarified the understanding of their agreement with Microsoft because they had some terminology in there about when AGI was achieved, that it would change the nature of their agreement. And I think that they realized that that term was too malifius to yield a legally enforceable agreement. So I believe that they've renegotiated those terms.

I didn't put one of the news stories on my website. Maybe if I find one that I like, I might dig into that and put up a good one. But you were talking about a lot of the things should belong in the National Enquirer. I have an AI newsfeed on my website.

Robert J. Marks:

Yes. By the way, I looked at it Dom and it's excellent. You really keep up on things. So I would recommend it to our listeners.

Donald Wunsch:

Yeah. I update a few times a week. It's personally curated and I don't put stuff in it if it's real hype, even if I like it. So for example, I'm a user of the paid version of ChatGPT, one of the higher tiers. I'm a user of the paid version, the Google Gemini Pro. And some of the other ones I use too. I use Microsoft Copilot. And so these tools, they're flawed, but they're useful. So in all of them, I've found shortcomings that I

would not tolerate from a human assistant, but they will make promises. They'll say," I will do something, "and then they give you something and they say, "I've done it." And it turns out that there's just a hollow shell and they didn't do anything that you described within there, particularly in the space of coding and they're improving. So I don't think that things will forever be that way, but I do think that if you want to be a user of AI, particularly LLMs, you need to be very aware of its limitations.

And that's one thing that I think earlier in the field where people were not talking about such general tools, then there was a very clear sense that I'm designing an AI system to do X and outside of X, it's useless, but within X, it can do what I want it to do, and then I can diagnose if it needs to be improved in this way or that way. And that's a good way for an engineer to think about AI. And it's very tempting to think of an LLM as an Oracle. It does have a broader range of user interface. And so you could ask it about all sorts of things, but it takes more responsibility on the part of the human to decide whether what you're getting back is actionable, whether it's useful, whether it's true, whether it's been hallucinated or not.

So you need to calibrate carefully. Often I'll be doing something right in the center of the strike zone of the AI. I'll be asking it for looking up things about a particular scientific topic, for example, and it'll give me back something and I'll say, "No, you need to try again. You need to work harder on this. You're giving me something very shallow and there's stuff out there that I know exists that is a better answer to this. Go back, but I'll give it some hints. I'll give it some direction of look for this." AI has been and will continue to be a powerful tool in the hands of humans. And the more that people realize that and use it accordingly, the more successful that they'll be in the application of AI. And there's a lot of work that's high consequence where you really need to get it right or you're going to lose money or maybe lives. And so it's really important. If there's one thing that engineers need to know about AI, it's that the engineering mindset has always been to work hard to verify that the system works as intended and to realize and declare the limitations of any system that you feel.

Robert J. Marks:

Ronald Reagan used to say, "Trust but verify." I think with AI, it's backwards, isn't it? It should be verify, then trust.

Donald Wunsch:

Yeah. Maybe verify and then trust, but still verify.

Robert J. Marks:

But still verify. The other thing is that when you get responses from these different platforms, they seem so sure of themselves. Another Emo Phillips joke, and I think this is really true of these large language models and these AI platforms is that computers are not really intelligent. They just think they are. And certainly they present themselves as they are intelligent.

Donald Wunsch:

So there's one thing in the article this reminds me of, and that is that there's a section about the threat of AI taking over. And so anyway, I think it's a red herring. So there are many threats of AI that really deserve our attention now. And a huge one is the threat of AI to our privacy because this device ... I won't put the side with the brand on it because I'm not blaming any particular company, but our smartphones have created an amazing ability to spy on us. So I would say that they have reduced the cost of spying on us by a factor of at least 10,000 and maybe 100,000. So they can track our every motion. They can be turned on even when you think they're turned off. You could then use the

microphone and the camera and the GPS on the device. You can do all sorts of things with this if you have either the right authority or the right hacking capability to do that. And if you have all that, installing such software might cost you a few thousand dollars, but once that has been installed, then the marginal cost of using it to spy on somebody is just a mouse click. It's negligible.

So anyway, the ability to do all sorts of things to spy on people has decreased, I would say easily by a factor of 10,000 and maybe much more. But what has still been expensive is to decide which people you want to spy on. And AI has reduced that cost. So in the past, I'd say we could take some cold comfort in the fact of, although it's very cheap to spy on us, it's expensive to decide that we're worth spying on. And furthermore, I'm not a senator or a billionaire or a CEO of a multinational company or somebody even more important than a senator. So my relative obscurity is a shield against the ease of which it would be to spy on me. So yes, if somebody really wanted to know to spy on me, they could spend a tiny bit of marginal costs. There'd be a one-time cost, but then the marginal cost to spy would be very low. But I can take some comfort in that the cost of deciding to spy on me would be very high. Well, AI easily lowers that barrier.

So if you had some parameters that you were trying to decide who to spy on, then AI would dramatically improve the ability to say that people with let's say a certain income level, a certain political persuasion, a certain sexual orientation, a certain religion, a certain geographic location, all sorts of criteria ... A certain level of history and various public records, whatever they might be, a certain purchasing pattern. There are all sorts of things that would make it much easier to spy on people, intrude on their privacy, and that doesn't even begin to open the Pandora's box of social media. And so the things that we voluntarily disclose. So it's gotten to where the risk of AI to privacy is not some hypothetical future risk that we might arrive and we're debating about whether it's coming in three years or five years or 10 years or 30 years or never. These are risks that are already here that are not sufficiently appreciated.

And so I would say in that sense that the concerns about the AGI and the policy implications about what to do about AI have got certain risks backwards and that indeed it behooves certain entities to encourage you to look at this family of risks and ignore this family of risks that I'm pulling off of the screen. But this family of risks are already here and already a danger. And if you say, "Hey, look at these. Actually, look at these. These are the ones already here and already in danger." "Don't look at those ones. No. Those are no problem." No. That's a pretty big issue.

And also regulation creates barriers to entry. So there are entities that can hire armies of lawyers. You think about the trillions of dollars being spent on AI. And so literally of the top 10 companies by market capitalization, all of them basically trillion-dollar companies, nine of those top 10 are AI companies. So how many lawyers can they afford to hire? As many as they want. So to have government regulations, they're needed, but they have the ability to hire lawyers and lobbyists. So they can create barriers to entry by having regulations that make it hard for new entrants, and they can also lobby to not have regulations that they're not very interested in. And so for example, privacy would be a good example. That privacy might interfere with their ability to market to you or their ability to do various things.

Robert J. Marks:

Well, my daughter wants to put in an app and follow me wherever I am so she can note wherever I am, she will know. And I've resisted it. I don't want to surrender that. In terms of privacy, I was doing some work with a nonprofit and this one guy said, "Well, maybe he can make a contribution." And the people at the nonprofit looked up this special service and they showed everything about this guy. He was rich. He was one of these guys that you said that AI should pay attention to. The Bible says that God knows the number of hairs on your head. It didn't go down to that detail, but it was pretty dang-on close.

Donald Wunsch:

I'm giving God a much easier problem with each passing year. You're counterbalancing it. You're giving him a little more work to do here, but ...

Robert J. Marks:

We are. And we have a professor here that got a very nice NSF grant about mining information from publicly available social media. And his task was to do things like identify people that are possibly violent, have violent tendencies, might be school shooters or people that might have a tendency to commit suicide and things of that sort. And all of that could be mined from all of this social media and the person could be identified. So you're exactly right in terms of the loss of privacy. But again, with Amazon and amazon.com and some of these other services, I have surrendered my privacy. I've surrendered my privacy to Google because their service is so doggone good that it's worth that trade off. So that's where we're at, I think.

Donald Wunsch:

Yeah. And by the way, this month I think is the last month that Google Maps timeline. They were keeping that on their servers. And now if you want to keep your timeline, it's pushing it to your device, but you can still keep your timeline. I like my timeline, but it's a privacy trade-off.

Robert J. Marks:

What do you mean by keeping your timeline? I don't know what that means.

Donald Wunsch:

If you look in Google Maps, there will be a menu item on the left-hand menu dropdown for your timeline, and it'll show everything that Google Maps has detected you doing over the time that you've authorized it. And so by the way, I do share mine with my wife and son. So basically, if something were to happen to me, they could find out where I am, or at least where my phone is. They might not find out where I am, but they can find out where my phone is. And actually that helped one time. So one time I didn't have the same ... This is a blue phone phone jacket, but I had a black phone jacket and I was in Washington DC and I had a cart with a little black thing where the child can sit in. And in DC, they charge you for the bags. So I had my own bag. So I checked out and then I put stuff in my bags and when I packed up all my bags, I didn't spot the black phone on the black child seat and I left it in the cart.

And somebody came along and instead of turning the phone in, they took the phone. And then I didn't realize this, but I mentioned it to my wife after I did realize it, but I had to run off to a meeting. And my wife called the phone and the person said, "Well, I want a reward for finding the phone." And so she could see where the person was and he told her where he was. And this is remarkable. You know Hong, she's fearless. She went off in downtown DC to go meet this person and she was on her way there and I guess she ran into some people that she told the story to on her way there and they said, "Oh no, don't do that. That doesn't sound safe." And they called the cops and the cops came and a cop went with her to meet the person. And he changed his mind about asking for the reward. And she said, "No, I said I'm going to pay you the reward, I'll pay you the reward and take the phone." And it turned out he was one of these vendors downtown on the National Mall and he gave her and the cop a free bottle of water. And so everybody was happy, but yeah, Hong is fearless. When I found that out, I was horrified, but I got my phone back.

But anyway, she was able to do that because I had enabled her to track my phone. The guy had told her where he was, but she was able to also see where he was. So it was interesting.

Robert J. Marks:

Oh my goodness.

Donald Wunsch:

I was off at DARPA or something. I came back from my meeting and she told me this whole story.

Robert J. Marks:

That's really astonished. So you've talked about some of the things that we should be careful about with artificial intelligence. On the page that you curate where you actually go through news articles, you said that you censor some of the articles.

Donald Wunsch:

No. I don't censor any of them. I either put the whole article or I don't include the article.

Robert J. Marks:

Okay.

Donald Wunsch:

I curate them.

Robert J. Marks:

You're selective. You curate. Okay. That's a better word.

Donald Wunsch:

Yeah, I curate them. Yeah.

Robert J. Marks:

So what are some of the key things that you look at in an article about artificial intelligence that let you know that it's clickbait or junk news or not worth posting?

Donald Wunsch:

Well, also there's a lot of worthwhile articles that I don't put because I'm interested in an engineering approach to AI. I'm interested in where the field is moving. I'm interested in its limitations and also its strengths. So for example, I will have more stuff related to the energy footprint of AI than maybe some other person with similar desire to run a feed. In fact, one of my friends, Daniel Tauritz of Auburn University, he used to be here at Missouri S&T and he's got an AI newsfeed and he does a separate cybersecurity newsfeed and they're really excellent. Daniel Tauritz, T-A-U-R-I-T-Z, is a really good person to look up, just remarkably prolific about choosing these articles. And he chooses different ones than I do, but very good ones too. He doesn't put any garbage up there.

So for example, I put up a New York Times article about the people who are suing OpenAI for these sycophantic chatbots that have encouraged people to commit murder and suicide and stuff like that. I've put up things that relate to policy. So some of the policies that are currently changing the field that are being put out by the federal government, some of those I have put stories out about those. Or sometimes links directly from pronouncements. There was an executive order on April 23rd about AI for

K through 12 that I thought was important. So I just put that. There were articles about it, but I just put the executive order out there saying that the agencies need to do more to educate all our students about AI. And so basically things that are in the sphere that I think my audience would be interested in. So there are all sorts of things that are ...

Oh, like AI for games is huge. That book that I mentioned about artificial intelligence, that was written by Julian Togelius, and he is also one of the people who's done a lot with AI for games. And so I'm very interested in AI for games, particularly in the early phase of my career. I was working on AI in the Game of Go in the late '90s and for about a decade beyond that. And some of the things that are done with AI for games are still fascinating. The work that DeepMind did on StarCraft immediately after beating the Champion in Go, that work was also quite fascinating. But I would say I don't post as much on entertainment AI and AI for games as I would've done if I had started the site a decade ago. I wouldn't quite say it's a solved problem, but I would just say that it's a relatively mature part of the AI field and I'm still very interested in it, but more as a customer than as a topic that I think is changing the world.

But I do think that the tools that they developed in terms of AI games are extremely powerful. And so the LLMs are built on these reinforcement learning tools. The reason that they're so widespread is because these techniques really work. And so if you have designed the right cost function, if you define the problem correctly, if you define the problem in a focused enough manner that you can predict what the AI is likely to be able to be beneficial for you, then those tools are extremely powerful. So for example, I was once asked about future of AI and I was saying that I thought that I personally would be more likely to trust say an AI hedge fund manager or a mutual fund manager than I would be to trust an AI baby center.

But I would suggest that one of the risks of AI is that maybe the mass majority of the population are making the opposite determination and that I don't think is good. So I would say that a lot of people are de facto allowing AI babysitters because there are AI algorithms that are driving the social media feed, there are AI driving the gaming and making those more addictive. And so people are essentially subject to AI babysitters and their performance might not yield the desired result. Will not yield the desired result, I'll say it a little stronger. Whereas AI might indeed be able to be designed to do better than a lot of professionals in money management.

I've worked in this field since the time I was your student in the late '80s and early '90s. I've published in this field as early as 1991 and some publications beyond, but I've never used it on my own money. I thought that I could do better. But I do think that it would not be that hard to design a tool using AI that would be the average, say, professional money manager, professional mutual fund manager, professional hedge fund manager. And there are other people who have reached that conclusion, and some of them have already done this and made a lot of money. So you can look up Bill Simon. There's a book about him called The Man Who Beat the Market. You can look up David Shaw, David Shaw & Associates. These are people who have been interested in this and turned them into huge enterprises.

Robert J. Marks:

Don, I knew the first professor of financial engineering in the United States. I forgot his name, but he was an interesting guy. And he was approached all the time by people that said, "I came up with a neural network that can forecast the market." He didn't even have to look at the software. What he did to these people that approached him, he asked the question, "What kind of car do you drive?" Because if they had been successful in doing this, they would've applied it and been very, very rich. So about the forecasting, it's very ... I think it was Niels Bohr. There's an old Danish saying that Niels Bohr, the quantum mechanic guy said, he said that forecasting is very dangerous, especially if it's about the future. We talked about some of the things that we use, some of the tools that are used to talk about fake

news. I think one of them is that people that make totally speculative statements about what the future of AI is going to be.

I also look for seductive optics and seductive semantics where in ... Well, for example, anthromoporophizing. They make a robot look like a human being or a dog or something like that. So all of a sudden there's association with artificial intelligence with something which actually exists and exists in the sense that it's a human being. My wife, she used to just love our dog and I said, "Don't love the dog. It's a dog. It's not a person." I think it's the same thing with AI. We look at it and sometimes say it's a human being. It isn't. I told my wife it's a dog. I think that people that talk about the future of AI is, "Hey, it's a computer."

The other thing that I look at for fake news is the source. I always take with a grain of salt the words of people like Sam Altman or Elon Musk that are promoting what they are doing. They're promoting their company, and of course they're going to come out and they're going to be positive. So a lot of these things can be used to filter what the good AI news is and what the bad AI news is.

Donald Wunsch:

By the way, you mentioned computational intelligence and financial engineering. And I believe you and I were both at the conference in about '94 in New York City.

Robert J. Marks:

Yes.

Donald Wunsch:

And I think we shared the hotel room in fact at the Marriott. And anyway, I got in a conversation with a fellow who was doing modeling of ... I think he was using stochastic differential equations and he was modeling deviations from the Black Scholes model or something like that. And I had a fascinating conversation. And as you might recall, before I joined your lab, I got a master's in applied math and I think that that was my edge when coming to talk to faculty members. And I came to you and you were my first choice and thankfully you took me as your student and I found out that you could write down the math as fast as I could read it. So we're no slouches in math.

But I had a talk with a guy and he was doing this mathematical modeling using stochastic differential equations and other tools, and he was an economist. And so one thing that I'm warning people about hubris in the field of AI is that like the business school students, the English majors, people like that, they are not what maybe people in their 60s and older envisioned those fields as being when they went into STEM disciplines and saw these other people in non STEM fields. Some of them can run rings around us. And so this guy, when I looked at his papers, I had had this conversation with him and I was fascinated. I looked at his papers and I said, "He's doing good work, but it's an awful lot of work to study his work. I would love to collaborate with this guy." He was very friendly and had a great conversation for a better part of an hour and just fascinated by what he was doing. But I realized in 1994, I'm not in a position to keep up with this guy. His name was Philip Dybvig. He got a Nobel Prize a couple of years ago at Washington University in St. Louis. So I had correctly assessed the chops that he brought to the table. But that's one exciting thing that was already true back in the '80s and '90s and remains true.

The people that get attracted to this field are just stunning, what they bring to the table. And it does not matter what discipline they come from. I have learned things from our English faculty that changed the way that I approached the field of AI. We have psychologists and sociologists just doing amazing things. I love to attend the lectures of our top historians, and some of them are just doing fantastic stuff. And so the field has always been multidisciplinary and it's getting even more so. So the people who are good at

things like prompt engineering are very different than the people who are good at the mathematical analysis. And then the mathematical tools are sometimes very different from the ones that we focused on. So even PDEs and field equations and stuff like that, they're important, but there are other things that are important too. And so that's one thing that I love about the field is the diversity of disciplines that goes into the field is just staggering. I feel like a kid in the candy store, and so there's just so many different things to look at and read about and learn about and different people to talk to.And that's one of the most rewarding things about being in this discipline is that you can work on and learn about almost anything. If there's something you're curious, you can work on it, you can learn about it.

So your former student, Ewan Sanjudge, I got his master's from you, he's very interested in applying LLMs to the Bible and there's so many different fascinating questions about that topic and he's very excited about it. In philosophy, of course, I was just revisiting the work of the late Daniel Dennett who just passed this year, I believe, and he did some marvelous work in a philosophy of AI. He collaborated with Douglas Hofstadter who got his doctorate in physics, whose dad was a Nobel Laureate in physics. He decided to try something different. After he got his doctorate in physics, he went into a computer science department as a new assistant professor, decided to write a book. His first book got a Pulitzer Prize. So yeah, Douglas Hofstadter, Gödel, Escher, Bach: An Eternal Golden Braid.

Robert J. Marks:

Oh, yes. Yes. Great book.

Donald Wunsch:

Has written many great books since that.

Robert J. Marks:

Great book. Yes.

Donald Wunsch:

He's written many great books since then, so one of the giants of AI. So there are all sorts of people making all sorts of contributions. It's not all about the citations, it's not all about the awards, and we've just scratched the surface. So even if you know the subset of the subset of the subset that you want to focus on, you can easily spend three, four, five decades to delve into that facet of it, and there's a lot of work left to do. So if somebody wants to see AI, AGI, let's say, I don't think that it's coming, but if somebody wanted to see AI that is across a broad spectrum of capabilities, including physical and mental capabilities, including being able to interface with the real world in a diverse way, if they want to see that in their lifetime, I recommend to do research on a longer lifetime. I think that will give you higher odds of seeing these things in your lifetime than doubling down on just the AI aspect of it. So the embodied AI is already huge. Like I talked about, if the building caught on fire, I'd leave my computer here and get out of there. To have embodied AI that can do something about the broader range of things, the broad range of things that we encounter on a day-to-day basis, it's just not that near.

But if you carve out a real specific problem that we're interested in that's worth a lot of money, then yeah, AI might solve it way faster than you think it can. You might think that this is a century, you might find out next year, bam, they've got it solved. But as soon as you make it more general, it's a whole different ballgame. And so I am a big fan, as I already said, make clear, of both Yann LeCun and Demis Hassabis. Hassabis deserves that Nobel Prize for AlphaFold. That's a very specific intelligence. It's not a general intelligence. But reaching for that star of general intelligence helped him develop the tools that enabled him to solve one of the most important problems that we faced since the beginning of this

century. And I think that the experts thought that that would take a century.It used to be that if you could solve one protein folding problem, you could get a seven-figure grant from NIH immediately, and they solved like a quarter million protein folding problems when they published their first result. And they're way beyond that now. There are literally millions of people using those tools, and so the impact is just tremendous, absolutely tremendous.

So focused research on AI, I would say there's almost no limit to what it's going to do, but as soon as you start making these grandiose claims ... I'm glad that people are excited about it and trying to do that. I just think that that is not coming soon.

Robert J. Marks:

Yeah, I agree.

Donald Wunsch:

I'd love to be proven wrong, but I think that attempts to prove that wrong, and that's why I say near the end of the article, the challenge for IEEE members is that engineers have always had the attitude, the best way to predict the future is to invent it. And so people who disagree with what I'm saying in the article, I say that's encouraged. And the best way to express that disagreement is not with some grandiose claims, but I'm saying this from here in Missouri, the show me states. So the results can speak for themselves. And right now, the results are not telling a story that this is anywhere near. The results are fantastic. I have stock in some of these companies for a good reason. I believe that they're producing results.

Robert J. Marks:

A lot of people think that some of these companies like OpenAI and Grok and Nvidia are bubbles.

Donald Wunsch:

Well, I wish I had stock in OpenAI or Grok. Private equity is a whole nother thing, but I don't think that we're at an AI bubble. I think that such a thing could happen, but I personally don't think we're at one. But I could easily be wrong about it, but the reason that I think that is that there's a lo.

T you can do with brute force, a lot. And the problems that people are working on are so incredibly valuable. I think that there will continue to be investment in expanding the capabilities of AI because they're ... So what is it worth to have a better drug target that can conquer cancer or Alzheimer's? What is it worth to stop paying people eight figure salaries to run a mutual fund if you could do that with an algorithm that you own? What is it worth to be able to lower the cost of getting something into space by three orders of magnitude?

What is it worth to have a better solution for energy availability and more? There are things that are worth not trillions, but 10s, hundreds of trillions or more that are problems that AI might be able to reduce the cost of and improve the efficacy of by orders of magnitude. So even though these top nine companies that I mentioned are all trillion-dollar companies, nine trillion is a small number compared to the potential ... Now, actually those nine companies add up to about $25 trillion in market cap. $25 trillion is a small number.

Robert J. Marks:

It would make a big dent in the national debt, I'm sure. Hey, we've been talking to Dr. Donald Wanch about his paper, Artificial General Intelligence is Nowhere Near, Artificial Special Stupidity is Already Here and Policy Implications.

Donald Wunsch:

Artificial Specific Stupidity.

Robert J. Marks:

What's that?

Donald Wunsch:

It's Artificial Specific Stupidity.

Robert J. Marks:

What did I say?

Donald Wunsch:

You said artificial special stupidity.

Robert J. Marks:

I know some people with special stupidity. Okay. So this is-

Donald Wunsch:

I have been guilty of being one of those people from time to time.

Robert J. Marks:

That's true.

Announcer:

That's it for this week. We'll be back soon with more from Dr. Donald Wunsch. Until then, be of good cheer.

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.