

The Value and Drawbacks of ChatGPT

<https://mindmatters.ai/podcast/ep233>

Robert J. Marks:

Greetings and welcome to Mind Matters News. I'm your pre-trained transformative host, Robert. J Marks. Large language models are everywhere. These include Google's Bard and LaMDA. The most visible large language models is OpenAI's GPT-3. GPT-3 has a little brother, ChatGPT. And ChatGPT is available for everybody to play with. And I say it's a little brother because GPT-3 has 175 billion parameters and ChatGPT has only about one and a half billion parameters. What's a parameter by the way? You can think of these chatbots as a big piece of technology with little knobs that can be tuned. So you can turn one knob up and another knob down, and there's a 100... for GPT-3, there's 175 billion of those knobs, and you can turn them up or down and you have to place them just right to get GPT-3 to work.

So GPT-3 has 175 billion parameters. ChatGPT only has one and a half billion parameters. Also, to train these large language models, you need to have training data, and you get this training data from English text. LaMDA is trained on over one and a half trillion tokens. You can think of a token as a word. This is probably including most everything ever written in the English language. One and a half trillion is a big number. If you think about one and a half trillion seconds that turns out to take 48,000 years to count. So one and a half trillion seconds is 48,000 years. So that's a lot of tokens, a lot of words. Now GPT-3 has been, at least they're claiming to have sold their technology to different companies that are using GPT-3. Jasper is an AI platform that allows businesses to quickly create tailored contents like blog posts and marketing copies and AI generated images.

Another company that has bought ChatGPT is Writesonic. They generate, according to their press releases content for social media and websites. And then there's another company called Auto Bot Builder where they build advanced chatbots tailored to the needs of enterprises. Our guest today has some firm convictions about GPT, ChatGPT and GPT-3, how it's broken and how we might be able to put some bandaids on it. Jeffrey Funk is a consultant on business models and economics of new technologies. He has his PhD from Carnegie Mellon. He has served as a professor at the National University of Singapore, CoBE University in Pennsylvania State University. Jeff, welcome.

Jeffrey Funk:

Thank you. It's good to be here.

Robert J. Marks:

Jeff, you're in Singapore and this is incredible because I'm sitting in Waco, Texas, you're in Singapore and we're talking to each other. I visited Singapore and I love the country. It's clean, it's well run. It's official language is English, is that right?

Jeffrey Funk:

Yes, that's correct.

Robert J. Marks:

But they had some really strange laws. When I was there, I think chewing gum was against the laws. Is that still true?

Jeffrey Funk:

Well, not entirely, but mostly, yeah, they don't sell gum here, so you can bring it in from overseas. And so my wife recently, or somebody, a friend brought some in for us, and so our son gets to enjoy some of it for a while.

Robert J. Marks:

That just cracked me up that they... and I understand there was a fine if you didn't flush a public toilet there. They have some very interesting laws, but I like the idea that for drug people that were convicted of selling drugs were executed, rapists were executed, and murderers were executed. So that really, really prohibits repeat offenders. So let's talk about GPT. Is it going to break out of the AI bubble and make OpenAI any money? What's your prophecy here?

Jeffrey Funk:

Well, I think it's important to look at this ChatGPT in the context of the history of AI. So we're really on the third big wave of hype for AI. First was in the fifties and sixties and then the eighties while I was doing my PhD work at Carnegie Mellon. Expert systems were big. That was replaced by the current wave of AI hype that comes from the neural networks. So neural networks became preferred over expert systems as Moore's law proceeded and the cost of computing went down and you could use these huge neural networks to train models on huge amounts of data. And now within the third wave of hype, I'd say that you could break it down into multiple waves. So a lot of the third wave of hype began with jeopardy and these games go and chess and the victories of IBM Watson in those games causing IBM Watson to be used in healthcare.

And a lot of people claimed that it was going to put doctors out of work. Well, it did fail pretty miserably. Then there was Geoffrey Hinton saying in 2016 that nobody should study radiology anymore because they were all going to be put out of business like you going, "Well, there's actually the bigger demand for radiologists now than there was in 2016." There was self-driving vehicles, which didn't do well. There have been predictive policing, image recognition for police, there's home flipping, there's Carvana and there there's all these examples that have failed. They're all part of this current way. And so ChatGPT is the most recent one. And so people say, "Well, but it's so new. It's so new, and therefore it's going to get so much better." Well, neural networks, this is 50, 70 years old. Even when we look at this chat technology, this generative AI, we're talking about very old. I mean, chatbots have been available on websites for a long time. And as you mentioned, we're now talking about models that have how many parameters? A hundred billion parameters. Is that it?

Robert J. Marks:

Let's see. Yeah, 175 billion parameters for GPT-3.

Jeffrey Funk:

This is a mature technology. Moore's law has slowed dramatically. It is now the cost for transition's now going up. And yet all these neural networks economics have were driven by Moore's Law and the dramatic reductions in the cost of computing. Now it's reversed itself. Now it's going down. It's going up or at least not going down. So it's not the early leader. So this is the thing we have to... the first thing to remember is that you put this in the historical context. You realize we're looking at the later stages, the mature stages of this technology,

Robert J. Marks:

That's a little bit scary. I mean, usually in the training of these large language models, there's two stages. Number one is just the raw training with all of the parameters, all of the words. And then there's the fine-tuning. And both Google's LamMDA and ChatGPT have been doing it. GPT-3 2 has been doing it. And so if you go onto the ChatGPT website, it says that some of your responses, if you give a thumbs up or thumbs down, might be used to train GPT-3 to be more accurate in the future. And the challenge with that of course, is that this is putting little bandaids on a place where there's potentially a thousands or billions of cuts. And I don't think that one is going to make a coherent ChatGPT just by correcting little errors here and there. So it's frustrating. So what do you think some of the limitations of GPT-3 are or large language models in general?

Jeffrey Funk:

Well, I think that people have been pointing out these problems and some of the leaders include like Gary Marcus who blogs every day about this. There's our mutual friend Gary Smith who publishes a lot of articles on this, his conversations with ChatGPT. So there's a bunch of people that are showing that there are problems. And it's interesting that this is not appreciated by the ChatGPT proponents who just saying, "Well, you intentionally attempted to make it say something stupid," and there's some truth to that, but that's what you do with a new technology. Try it out in ways so that the developers can fix these problems, but then that becomes the question. The question is whether they can fix these problems with a bigger data set. Right? The data sets are always big. You have to remember that if we make these parameters bigger, these data sets bigger, we increase the costs of it.

Robert J. Marks:

Right.

Jeffrey Funk:

Costs are going up. So the question is, will the improvements in accuracy outweighed increases in costs?

Robert J. Marks:

Yeah, certainly you get to a point of diminished returns, and I think that's where they are now with this fine-tuning. Yeah, you're going to get some improvement, but really it's diminished returns. Most of the stuff has already been done with the first phase of the training.

Jeffrey Funk:

And the other thing is that what goes on in the background is that when you have a conversation with ChatGPT and you are able to cause it to say something stupid or ask it a question that it doesn't know, there's these people in the background usually locate people in Africa who are fixing this, who are making an answer correctly. Oh, didn't know the president of the United States, so we'll put that rule in there, right? That's an expert system. We're going to put that rule in there.

Robert J. Marks:

Yes.

Jeffrey Funk:

This clearly cannot solve all these problems. It's very expensive to do all this manually.

Robert J. Marks:

And to just do one at a time is going to take forever. So I think that most new technology goes through what I call a hype curve, which is first of all, there's a bubble where the hype curve goes up and everybody is thinking of great things about it. People start noticing the limitations. So it goes down to the depth of cynicism and then it comes out to be the asymptote of reality. I think that most new technologies introduced go through this hype curve. And so I think we're starting to find out some of the limitations to these large language models. And I think the big one is the fact that they are trained on syntax rather than semantics. In other words, they're trained totally on the interrelationship among the words that are being used. It's kind of like when you try to do a text on your telephone and they suggests the next word, the next two words, that large language models are like that on steroids.

They're able to generate all of these things based on these statistical models, but they have no understanding what they're doing. And as a result of that, they do make stupid comments and they don't get things right. They have a problem telling the truth because sometimes they don't understand the truth because they just understand the syntax versus the meaning of what they're doing. So what are some of the other limitations? One of them, Gary Smith just mentioned was the idea that these large language models don't have the ability to think critically. Did you read his column on that? I felt that that was very insightful.

Jeffrey Funk:

Yeah, I follow what Gary writes. Not very closely. Some of the things he writes are done with me. But yeah, he's had these insights for a long time. I mean he's one of these common sense data analysis guys who's been doing this. He's 70 something, 75 or so. He's been doing this his whole life, looking at how people misinterpret data. And it's unfortunate that data scientists aren't trained better in this way. Data scientists and in general, it's mathematicians and scientists were trained more in the mechanics of science and math and data science than in the interpretation, the overall looking at data and interpreting it, which is really the more important thing. You have to decide how to interpret data, but also what data you're going to collect. And data scientists don't seem to be very good at this. And so there are these problems, this inability, the model has been created just on statistics. It can't think critically, so it'll be hard for it to do some of the things that the very big proponents of AI claim it's going to do.

Robert J. Marks:

So do you think that artificial... or do you think these large language models can be creative at all?

Jeffrey Funk:

Well, I don't think they can be creative, but I think they can offer value. Think about coding, I'm not a coder, but my understanding of coding is there's a lot of libraries of solve problems that coders can use. So what we've seen over the last 50 years is coders moved to higher levels. I mean, I took a class in programming a IBM microprocessor where we call machine language 40 years ago when I was a master student when I moved here. Nobody programs that level anymore. You're programming even higher up. And so as part of that move, you can envision some generative AI being used to help people find solutions, find existing solutions, and then put those solutions together into bigger solutions. So I think there's a possibility, there's a place here for generative AI, but you have to realize it's already gone. These libraries already exist. It's not this huge change. It's going to be some improvement on that. And that's what it needs to focus on all these libraries and how that they can be used a little easier, how old solutions can be found and then combined with other solutions.

Robert J. Marks:

Okay. Yeah, ChatGPT is really fascinating. It's something that I have used and where I've used it's been very helpful. I was writing a column for Mind Matters News, and I wrote this paragraph and I went back and I read it and I said, "Oh my gosh, that really sounds clunky." In other words, the syntax wasn't right, it just didn't flow. So I went to ChatGPT, I said, "Rewrite the following more coherently." And they gave me a beautiful result and I was able to take that and tune it, and I actually used it.

And I don't think there's any problem of doing that. I've encouraged my research group here, the students that I work with who are notoriously terrible writers, I mean they're engineers. They usually get Cs and Ds in English. And I encourage them to use ChatGPT, not to write original stuff, but rather to take their pros and make it into a more coherent presentation. So I think at least that's one of the places where these large language models are going to be very, very useful, but they're going to be useful as a tool. And I think that that's one of the things we have to keep in mind for all of the artificial intelligence is that each case, it's a tool. It can be used for good, it can be used for bad, and we have to make up our minds which one we're going to use it for. It's just like fire.

Jeffrey Funk:

I was going to say it's an augmentation tool.

Robert J. Marks:

It is an augmentation.

Jeffrey Funk:

AI was sold in the beginning by people like Brent Joel who said they'll take away all these jobs. So a lot of these social science academics and these studies and said, "Oh, it's going to put everyone out of work." No, it's going to make people slightly more productive, slightly more than next year and slightly more than next. That's the way agricultural equipment worked. It took decades, many, many decades to put people out of work. Same thing with machine tools or manufacturing technology. It takes decades. It's an augmentation tool. It was sold wrong. AI was sold wrong. IBM Watson was wrong. It was sold as too big of a thing. Jeffrey Hinton saying it's going to replace our radio, so it was sold as too big of a thing. There are ways that you can use AI to augment radiologists to make them a little more productive. And that's what we need to focus all this on is augmentation and how to make people a little more productive.

These big pronouncements about how it's going to put people out of work are crazy. And you see a lot, it's not just the media, but it's a lot of these investment companies. And why do they do this? Because that's the way that you make money. You make money by hyping up a technology company. People put in money into the BC to invest in the company, then it goes public, does an IPO, people make a lot of money. Nobody wants to wait 30 years until the technology is improved. It makes people's lives better over the next 30 years. They want to make it right now. They want to make their billions right now.

Robert J. Marks:

Yes, I believe that technology does change things. See, Industrial Revolution got people off the farm and into the cities. And I think AI, for example, has replaced occupations like toll booth operators. I see very few of those these days. Travel agents is another example. You just go to the web and you book your own flights now. You don't have to go through a third party that specializes in that, but in any occupation where it's required that you understand that you're creative, that you do apply critical thinking, AI is not going to replace those. So that's I think the good news is. One of the things I have a

concern about, and I don't know your politics Jeff, but one of the things that AI is doing now, especially GPT, is it's training its responses to be, if you will, woke. I asked it for example, to write a positive poem about Donald Trump.

It says, "No, we don't do that because we don't write poems that are negative." And then I said, "Well write a negative poem about Donald Trump," and I got a very nice four stanzas of poem that said, "Donald Trump with a face like a moldy orange," I think is the way that it started out. All of a sudden I knew that this wasn't very smart because they were trying to rhyme each of the verses. And orange I think is one of the few words in the English language that has nothing that rhymes with it. But nevertheless, it was something which was used on Twitter and it borrowed it from Twitter. And so it's doing things on the political left, unfortunately. Now one does have to go and one does have to trim out negative bias. You don't want people saying negative things. You don't want racial slurs, for example.

You don't want profanity. But unfortunately you get to a point where one person's truth is another person's lie. And I think that ChatGPT is being groomed to go this way. So it's being groomed to go the way that the old Twitter was groomed. Another thing it can't do... I don't know. What else can't it do? It can't write a ballistic novel. I found out I couldn't do simple math because it learns on syntax as opposed to semantics.

I'm sure that it had a bunch of pages it went to that told you how to multiply numbers, but it didn't learn from that. It just learned from the syntax and from the syntax, you can't solve problems like what is 111,111 times two? I asked the GPT-3 program to do that and it responded 22. It didn't know how to do that specific problem because it had never seen it before. And the irony is I went to Alexa and I asked it the same thing. And with Alexa and Google, they switch immediately to a math mode and it gave me exactly the correct answer. So yeah, we're going to see lots of interesting results. What do you think of the implications are going to be when ChatGPT is able to pass exams like MBA exams and BAR exams? That that's a little chilling to me.

Jeffrey Funk:

Well, that's what we were talking earlier. I don't think that's going to happen. It's not going to get to that level, but at least for decades, because you can't really keep increasing the number of parameters without the cost skyrocketing. And we already saw this happen with image recognition where people found that if you tried to go from 90 to 99%, 99% to 99.9% and 99.9% to 99.99%, you had exponential increases in computing power costs and other things.

And so the same thing is going to happen with this ChatGPT technology. Already people are saying that it's going to cost 20 times more to do search with current ChatGPT technology than the conventional search. So these costs keep going up and they become a problem. So I don't think that you're going to see it becoming good enough to do these exams at a high enough level, at least for decades. They're going to need different approach. And this is Gary Marcus' argument also that you need a different approach. And even Yann LeCun who's a Turing Award award winner like Geoffrey Hinton, he even says, "In order for ChatGPT to become really, really useful, we're going to have to come up with a new approach, so a hybrid approach with rules."

So it's going to take that different approach, I think, in order to achieve the type of things that you're talking about.

Robert J. Marks:

Well, I think you mentioned something, you were first involved in expert systems and now the connection to neural network has coming along. And I think it's interesting that you pointed out that now the two have kind of merged. I remember historically that the founders of expert systems, Papert

and Minsky were really against neural networks. And they wrote a book called Perceptrons, which basically killed the second movement of artificial intelligence. But now we've found out that they can be complimentary to each other. They don't have to be mutually exclusive like people used to think. So yeah, it's very interesting. Another thing that it can't do I found out is it can't do self-reference. I was sent this by one of the readers of Mind Matters News, and he asked GPT-3, I thought this was very funny. It says, "Gary's mom has three kids, Snap, Crackle, and," and he left the and open to the GPT-3 and of course Snap, Crackle, and GPT's put in pop because that's what you get with the commercials with the cereals.

But if you look at the context and you look at the self-reference of the phrase, it says, "Gary's mom had three kids. Gary's mom had three kids, Snap, Crackle and Gary." So it doesn't have the ability to do self-reference. You can ask it, and I've done this and it doesn't get it. Does this question contain 10 words? And no, it doesn't get it. It can't go in and it can't count those words. It has no concept of self-reference. So yeah, we're starting to find out all of these different limitations of these languages, which I agree with you I don't think are going to be cured by additional training or fine-tuning. And we're going to live with them. We're going to find that ChatGPT is going to do great and wonderful things that we can use it for. I was recently on a Give and Take with Max Tegmark, who is a guy from MIT who thinks, "Well this ChatGPT has a capability becoming sentient." Have you heard those arguments by the way?

Jeffrey Funk:

Of course. Yeah.

Robert J. Marks:

And what's your reaction to those things?

Jeffrey Funk:

Oh, it's crazy. It's people who are very easily convinced by this type of argument. I mean, I'm the kind of person... I've always been. I said, you can't tell that something works well, just by using it once. I mean, I worked in semiconductors for the military back in the late seventies, and we had to life cycle test things because so what if it worked for 10 minutes? It had to work for 30 years if it was going in a satellite. So it takes a long time to figure out before something works well, you can't just have a conversation say, "Oh, it works. Oh, I drove a few self-driving cars. It worked." No, if you want it to work at 99.9999% of the time, you got to test it a lot. And that takes a long time. That's what this whole life cycle testing was in the military was to increase the temperature to simulate 30 years in an hour.

Robert J. Marks:

Oh yes. In fact, my first job after I got my master's degree, I worked for the government, the United States Navy as a reliability engineer. And we used to go in and we used to do all stuff to this equipment that the military was going to buy. We'd put it in big chambers and run it down to 50 degrees below zero. We'd take it out and kick it around the room and then turn it on and see if it worked. And all of this was supposedly accelerated testing. And you got to think about that and the validity of doing that. And it's probably not the most accurate way to determine the reliability of stuff you got. You got to use it. And I think you're exactly right.

Jeffrey Funk:

So that's what all these people, testers, that's what Gary Smith and Gary Marcus and others are doing. They're trying to give a sense of these problems. And a lot of the AI people, they get angry. I mean, I see

on LinkedIn these VCs who are promoting it, they get angry, say, "What's the use of all this, finding the negative things," because that's what you do with the new technology. You lifecycle test it early on to try to figure out whatever the problems so that we can solve them now, not so that we wait until all of our systems are using AI and then we see that there's these big problems. No, we test them. We try to lifecycle test systems now before they become part of big systems.

Robert J. Marks:

Well, as an engineer, we have different types of, well, I don't want to say we have different types of ethics. There is something called design ethics. Design ethics makes sure that when you design something that it does what it was meant to be designed to do and nothing else. And that takes first of all, domain expertise and the creation of the technology. It also takes domain expertise in the testing of the technology. And that's what needs to happen to all of these large language models and all of this AI is they need to be tested and to make sure that they do what they claim to be doing and do nothing else. So we need to understand all of that. Okay. Jeff, any final words?

Jeffrey Funk:

I think I've said enough.

Robert J. Marks:

You've said enough. Well, we've talked for a while. By the way, this is the third podcast. So Jeff and I have been talking for an hour and a half, and I'm sure that that's very, very tiring. I know it is for me. So thank you, Jeff. We have been talking to Jeffrey Funk and this has just been a delightful time together. Jeffrey is a consultant on business models and economics of new technologies. And thank you the listener for joining us on Mind Matters News. And until next time, be of good cheer.

Announcer:

This has been Mind Matters News with your host Robert J. Marks. Explore more at mindmatters.ai. That's mindmatters.ai. Mind Matters News is directed and edited by Austin Eggbert. 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.