The 2020 AI Dirty Dozen Hyped Stories: Countdown by Bradley Center Brain Trust Members

https://mindmatters.ai/podcast/**ep155**/

Announcer:

You're in for a real blast from the past this week on Mind Matters news as we revisit the top 12 most over-hyped stories in artificial intelligence from 2020, with the members of our very own Bradley Center Brain Trust. Now here's your host Robert J. Marks.

Robert J. Marks:

Greetings. There are many forces that shape the AI news we read. One is a materialistic ideology that unavoidably leads to the conclusion we are meat puppets, and this conclusion says that AI will someday duplicate us. There are many other reasons for all of the hyped AI stories we see today. Media's everywhere and competition is fierce, articles with provocative headlines and content or clickbait for the browsing consumer. So, we're going to count down today, the AI dirty dozen, the top AI hyped stories for 2020. And we are joined by two members of the Bradley Center Brain Trust. And this is the first time I think that they've heard me call them the brain trust and I hope you like the title. First, we have Jonathan Bartlett. He is the director of the Blyth Institute. The Blyth Institute focuses on the interplay between mathematics, philosophy, engineering and science. And Jonathan is the author of several textbooks and edited volumes, which have been used by universities as diverse as Princeton and DeVry. And he is a senior fellow of the Bradley center. Welcome, Jonathan.

Jonathan Bartlett:

Well, thanks for having me on, it's an honor to be here.

Robert J. Marks:

It's great. We're going to have fun. And the other member of our brain trust is Dr. Eric Holloway, who works for the National Institutes of Health and is a current captain of the United States air force. He has served in both the United States and Afghanistan. He is also a senior fellow of the Bradley Center. Welcome, Eric.

Eric Holloway:

Thank you very much. It's awesome to be here.

Robert J. Marks:

Okay, great. We are going to start with number 12. This is the dirty dozen AI stories of 2020. This is number 12. It's an article from the MIT technology review and the headline is, The Way we Train AI is Fundamentally Flawed and the subtitle is The Process Used to Build Most of the Machine Learning Models we Use Today Can't Tell if They Work in the Real World or not and That's a Problem. Eric, what do you make of this?

Eric Holloway:

Yeah, this is actually a really insightful article and it's true. The problem they identify cross cuts every major machine learning technique out there, because what we call AI is actually more properly called machine learning. And essentially it's just curve fitting. You have a bunch of data points and you find the best curve that fits those data points. Although it's a little bit more complex than just a curve, but it's the same idea. But if you keep that idea in your head, it's also easy to see why there's a problem. So, let's start with a really simple example. Let's say you have a 2D graph and you have a single data point on that graph. That makes sense?

Robert J. Marks:

Yep.

Eric Holloway:

Okay. Now I asked you to fit the best line you can, to that single data point. There are so many different lines that can fit that one data point, and they're all very different lines from each other. And that in a nutshell is the problem with modern AI. Even though we have millions or billions or even trillions of data points, the models themselves that are being trained on these data points are still like this line being trained on a single dot the models themselves are so, so, so complex that even with billions or trillions of data points, that models are still very under specified. And like with the dot example, you can have both a line sloping up and line sloping down, which both perfectly fit that dot, but on other data have very different predictions.

Eric Holloway:

And so that's the problem with modern AI, even these fancy techniques like deep learning, the deep learning model, you'll have many different models that fit the same vast data sets. And these models will have very different predictions on new data that is not contained in the dataset. So, when you hit the real world and you're not just in a lab anymore, the data you're going to be analyzing is going to be very different. And so that's why all these models start to fall on their faces once we go into the real world. And the basic problem is that they're under specified. There's just too many different models that are very different from each other that can fit the same data.

Robert J. Marks:

Artificial neural networks are supposed to be bonafide. Once you train them, you're supposed to subject them to a bunch of test data or validation data, just to make sure that they work and they do what they're supposed to. This is data, which was not used in the training. And it seems that here, they're saying that things not used in the training are for some reason that the data to which the neural network is subjected is not equivalent to the original training data. I mean, this is old news, but I don't think they do a lot of cross validation and deep learning, like deep convolutional neural networks. Is that true?

Eric Holloway:

Well, they'll do validation as they're training it, but the problem there is that their test set, which is supposed to be independent of the training itself actually starts seeping into the training. So, they still don't achieve independence there.

Robert J. Marks:

Yes. Yes. That's an old story in financial neural networks is that we ... I was friends with Jack Marshall who was a professor of financial engineering and he had all of these people come in and say, I have trained a neural network to forecast the market. And they used this idea of training the neural network, and then they tested the neural network. But what they did is they tested the neural network and they say, well this test didn't work very good. I'm going to change my neural network. So, they changed the neural network and then they tested it on the same data. And now he says, well, this works a little bit better in that they went back and they did it again. And finally came up with a result, not realizing that by the person in the loop and the testing data being applied again and again and again, it became part of the training data. I think that's the point you're making.

Eric Holloway:

Right, exactly.

Robert J. Marks:

Which is really, really interesting. Any thoughts, Jonathan on this?

Jonathan Bartlett:

Well, the other thing is that anytime you have a model, there's things that match the model and things that don't, there are things that are inside what you can expect and things that are not. And one of the problems with a lot of the AI work is that there isn't really a clear definition always of why the data is being chosen and why those specific fields. Sometimes it's just what could be measured. And also there's not a good clarity about what's in bounds and out of bounds. For example, in oil pumping, when they have the motors that pump the oil, they have these curves that they do, and they have equations that model pump performance on these curves.

Jonathan Bartlett:

Well, the models are only valid within specific regions and outside those regions, not only are they not valid, they are completely off the mark. There is no relationship between data and reality once you step a moment outside of those bounds. And that's what I see happening a lot with AI is that something maybe within the bounds on the things that that people think to train for, but then that's not necessarily how they're going to be used in the real world. And once you make that switch, then the models aren't valid anymore.

Robert J. Marks:

Excellent. Yes. George Gilder wrote a piece for the Bradley center called gaming AI, and he comments that AI is restricted to something that he calls ergodic. Basically the data from the past is enough to forecast data of the future. And I think that that is something that's a very simple concept, but there's lots of things which are not ergodic in the sense. You can't forecast the stock market. It isn't ergodic. The data of the past will not allow estimation of the future.

Jonathan Bartlett:

The other part of ergodicity is that you can't actually predict the impact that the AI system itself will have. So, if you think about the stock market, not only can we not necessarily use the past to predict future performance, but let's say that I came up with a tool that for some reason could, well, what the tool can't model is what will be the effect of that tool in the stock market. And so even if we could have in ergodic notion of the stock market, that would fail as soon as we introduced a new AI tool that looked at it differently and started trading differently.

Robert J. Marks:

Yeah. Fascinating stuff. Number 11, Transparency And Reproducibility in Artificial Intelligence. Now this is a paper from a very prestigious journal Nature, and it questions some of these things. Eric, what is the hype here? What's the problem?

Eric Holloway:

Well, yeah, the problem is now AI is not merely like a research project, but it's also a product and it's a product of some really big companies like Google. I think Google has said, it's an AI first company. I think Facebook has too. So, it's now a really big part of their brand. And so it's in their interest to inflate AI as much as possible. And we see this a lot with the results they released, like with the AlphaGo. And I think now with their AlphaFold protein folding. They don't actually release anything that people can use to reproduce their results.

Eric Holloway:

They just say, hey, we ran these massive neural networks on these massive datasets with massive amounts of compute. And we got super great accuracy scores, and you can use our model to get our same scores, but we're not going to really tell you how we did it. We might kind of hint at it, but we don't give you enough specifics where you can repeat it. And also it's actually out of reach of pretty much anybody who's not Google, because these computations cost like millions and millions of dollars and use massive computer farms.

Robert J. Marks:

There's an old saying in engineering that in theory, theory and reality are the same, in reality they're not. And I think when you reduce something to practice, that's where the rubber meets the road. That's what's going to be important. On the other hand, Eric, doesn't Google make available to the public, this incredible software platform they call TensorFlow and other AI sort of software that they can use? But you're not talking about that, are you?

Eric Holloway:

No, it's not like the tooling. Well, they don't even release all their tooling. They give us little bits and pieces of it. Enough that other people will start getting addicted to Google, but not enough that we can really do what they do. So, they released TensorFlow, but there's always a difference between the tools that Google releases to the public and what they actually use. But also, what I'm talking about though, is the specific technique. So, TensorFlow is a framework that makes it easier to write these AI algorithms, but the actual algorithms and models themselves. That is the secret sauce that Google is not really releasing.

Robert J. Marks:

I see. So, it works and just trust us.

Eric Holloway:

Right.

Robert J. Marks:

I see. Okay.

Eric Holloway:

There's even a bigger picture issue why AI is not scientific. And that gets back to its fundamental assumption that everything a human mind can do, you can do with the computer. Everyone in the AI field just takes that for granted. They're like, oh yeah, of course.

Robert J. Marks:

You're saying that AI doesn't follow the scientific method.

Eric Holloway:

Yeah. The very premise of the field is unscientific. Science is all about questioning your assumptions and testing them before you accept them as valid. But AI is the complete opposite. They take their assumption and treat it as valid and then do all their research and stuff based on that assumption.

Robert J. Marks:

Boy, that's an interesting viewpoint. And yeah, I would agree with you. That's a ...

Eric Holloway:

Yeah, it's ironically how people like to talk about creationism, where they start off with their theology and try to make the science and data fit it. That's exactly the same with AI. They start out with their assumption and try and make all their science and data fit their assumption.

Robert J. Marks:

Goodness, great observation. Number 10, Will Artificial Intelligence Ever Live Up to its Hype. This is a article from a very prestigious publication, Scientific America. And the subtitle to the, will artificial intelligence ever live up to its hype is, replication problems plague the field of AI and the goal of general intelligence remains as elusive as ever.

Eric Holloway:

Yeah. This actually directly builds on what we were just talking about. So, because we have this training problem where they don't really train their models in the way that fits the real world and they don't really have the constraints well-defined and they don't really follow scientific methods and they're not even scientific fundamentally. So, it's kind of unsurprising that once you hit the real world, then all the hype kind of deflates and the author of this article, he looked at, I think, 40 different startups, AI startups that were originally really hyped. They're going to change the way the world is and everything. And after the fact, once they actually started trying to use their

product in reality, then all of their venture capitalists decided, these companies aren't really living up to the hype. The AI is going to be much less impactful than we originally thought.

Robert J. Marks:

That's fascinating. I think there's always been hype associated with AI. In fact, in 1957, I ran across a New York times article, 1957, that the Navy had come up with artificial intelligence that in the future would be able to walk and talk and reproduce. And this was the hype in 1957, this was back when Bernie Woodrow at Stanford and Frank Rosenblatt, I believe at Cornell were doing rudimentary artificial intelligence. And the hype was there.

Eric Holloway:

Well, even at the very beginning, the field was started by like a Marvin Minsky and Claude Shannon, and some of the luminaries of information theory and they were like, yeah, let's just get like 10 of us really smart people and give us funding for like a month. And we'll give you intelligence that can learn just like a person do all the things, just like a person.

Robert J. Marks:

I've heard of that. Do you know the date that, that happened?

Eric Holloway:

Not the quick. No, I don't have the specific date off my head, but yeah, I'll write an article on that, but yeah, it's pretty funny. They're like, okay, just a month and then we'll have something completely like human intelligence and it'll be done. And here we are like three or four, actually eight decades after that.

Robert J. Marks:

Yes. So, it's not an apparent problem. The other thing in this article, there is the assumption that general artificial intelligence, or I think it's called AGI, artificial general intelligence. It keeps changing names. It used to be hard intelligence, hard artificial intelligence, but there's the assumption that this can be achieved. And I think that both you and I are on the page that there are fundamental challenges in computer science that are going to prohibit this from ever being achieved.

Eric Holloway:

Yeah. And the very fact that we have to differentiate the fields now actually points to the problem originally when Shannon and Minsky were coming up with the field, they're like, oh yeah, it's just a computation. We'll just have a fancy algorithm and that'll do it. And now we're finding all these algorithms we thought were going to be the AI turned out to only actually work in very, very small domains and very restricted data sets. So, that's why now we have to be like, okay, well we have AI that does something, but it's not actually artificial general intelligence, because they're all really narrow domains that they actually work on. And this is actually, there's a fellow. I think he's a fellow of the Bradley Institute. But anyways, he's working with Dembski. A gentleman named Erik Larson and he's actually going to be releasing a book, I think sometime next year, about this fundamental difference between minds and machines. He's quoted in this article here.

Robert J. Marks:

I look forward to Erik's book. Erik's book is going to be published by Harvard University Press. So, he has a very, very prestigious pedigree and that should be released very, very soon. And so we're excited about that. Okay. We are counting down the dirty dozen hyped AI stories of 2020, and we're up to number nine. AI Superstar, An AI Robot is Cast in the Lead of a $70 Million Sci-Fi Film. This was reported both of Mind Matters news. And by the way, we're going to supply links to all of these stories on the podcast notes. So, you'll be able to review them yourself and check the accuracy of our claims and commentary. This was published in June, 2020 in the Hollywood reporter. And apparently we're going to have a robot in the lead role of a Sci-Fi film, a $70 million Sci-Fi film. Eric, what do you make of that?

Eric Holloway:

First of all, I asked, why is this news? We've had animatronics in movies since star wars or the Muppet movies. So, this is just a fancier Muppet movie.

Robert J. Marks:

That's right.

Eric Holloway:

Yeah, they're just fancy puppets. It's like Sesame Street, but with a bit more electronics. And it's really funny reading these articles, because there's a whole lot of anthropomorphic [inaudible 00:18:06] going on with these AIs. They go through great pains to make it sound like the AI is learning something. They're practicing all their lines. And they're trying really hard. They're trying to make it sound like a real person, when all they're just doing is some engineer, underpaid engineer in the back, is running the algorithm a couple of times on new data sets.

Robert J. Marks:

It's kind of like the wizard of Oz behind the curtain, right?

Eric Holloway:

It's really people underneath.

Robert J. Marks:

If you have visited Walt Disney world or Disneyland, you go to the hall of presidents. And it was big news when Disneyland opened, I think in the late 1950s that they had these robots and these robots would come out and they would be dressed up like the president and they would deliver a speech and the mouth would be synchronized to the words. And everybody was really astonished. Now, clearly we've come a long way from that. But this idea of animatronic, how did you say it? Animatronics.

Eric Holloway:

Yeah, animatronics.

Robert J. Marks:

Animatronics has been around for a long time. In fact, I read somewhere that Disney patented his hall of presidents, his technology back a long time ago so that nobody else would duplicate it. Also, I think that maybe this has very little or nothing to do with artificial intelligence. I don't know. They say that they trained this robot. So, maybe there's a training algorithm associated with it. But we see a lot of what I refer to as seductive optics in these presentations and these movies. And as some of these hyped versions where in you come out with a robot.

Robert J. Marks:

One of them that was recent, recent maybe being about a year ago, was Sophia. And this was supposed to be a really, really exciting thing. And people looked at it and said, oh my gosh, artificial intelligence. But the robots Sophia was nothing more than an animatronic robot that synchronized their mouth movements and their facial expressions in order to communicate. And the optics, which was to have it inside a human form, had nothing to do with artificial intelligence. The container of artificial intelligence often has little to do with the driving artificial intelligence itself.

Eric Holloway:

Yeah. And who's to say whether there's actually a person controlling the bot behind the scenes? This is actually really, really old as 1700s or so they had supposed chess playing robots, but really there's a big table where the robot sat and underneath the table, there's just a person hiding in the table, moving the robot's arms. So, this is really old. It's just, yeah.

Robert J. Marks:

I remember that.

Eric Holloway:

[crosstalk 00:20:47] it's more the same.

Robert J. Marks:

And isn't that, in general, true in a much, in a sense to today's artificial intelligence, all of the intelligence we see as due to the computer programmer asking the artificial intelligence to do something.

Eric Holloway:

Yeah, exactly.

Robert J. Marks:

And we might be surprised at the output. I mean, you talk about AlphaGo making the incredible move when it beat Lee Sedol, the Go champion, but that was a surprising move, but there was really no creativity there, because my goodness that software was trained to play Go, and that's what it was doing. And it was doing it a lot better than humans, just like calculators calculate a lot faster than we do. And cars go a lot faster than we can run. So, it's surprising. It's kind of cool, but it certainly is not creative. The creativity came from the computer programmer.

Eric Holloway:

Right. And also if you look at these AlphaGo type things and you look at what they actually do, they even have to achieve that result. They go through like billions or trillions or pentillions of calculations and checks to even arrive at those results. Vastly, vastly more than any human, even over the history of humans playing Go ever do. And so once you look at it more as just like brute force, trying every possibility, it doesn't really seem so impressive anymore.

Robert J. Marks:

Wow. Fascinating stuff. We have already gone through number 12 through number nine, we're on number eight. And the question here, Is AI Really Better than Physicians at Diagnosis? We're told AI is going to replace lawyers and doctors and accountants and all sorts of people. So, let's look at a case of the physicians. This was a piece written on Mind Matters news. And Eric, what do you think, do you think that AI will ever be better than physicians at diagnosis?

Eric Holloway:

Well, I don't know if they ultimately will or will not, but right now they definitely are not. And this gets back to something that John brought up last session about just how unscientific AI sites is. This particular author, he took a look at 10 years worth of studies for deep learning algorithms on medical problems. And only two of them actually relied on randomized trials, while 81 were non-randomized. This means basically people can just pick and choose the type of data that makes their algorithm work well. So, really their results don't really tell us anything about how well their stuff works in the real world.

Robert J. Marks:

Yeah. That's fascinating. Gary Smith, who is also one of the fellows of the Bradley Center, he talks about the idea that when you publish a paper based on statistics, you got a problem. That 90% of his papers that are published on statistics are wrong. They're not wrong, but faulty, I think is the word he used. And indeed, that's the case when you have incomplete or unstructured data that you're trying to train with. Is this your point?

Eric Holloway:

Yeah, basically they didn't follow very strong protocols. And so they can just make a thing that works really well in the lab. And we have no idea if it's going to work in the real life and they tend to not work very well in real life.

Robert J. Marks:

Yeah. Gary points to a case of what did he call it? He called it the Texas sharpshooters fallacy. And the idea is, is that if you have a barn door and you paint a bunch of targets on it and you shoot at the barn door with an arrow, you're going to get close to a bullseye if you have a thousand targets up there. So, there was this one case about pancreatic cancer and they began to look at correlations with pancreatic cancer and well they thought it was caused maybe by smoking. No, it wasn't caused by smoking.

Robert J. Marks:

What about, I don't know. What about chewing tobacco? No. Chewing tobacco. Drinking tea. No. How about smoking cigars or pipes? No, it didn't correlate. What about drinking coffee? Oh my gosh, there was an incredible correlation there. So, they published this in the new England Journal of Medicine and coffee futures fell and people stopped drinking coffee. And in fact, in the end it turned out that it was totally just a coincidental correlation. And subsequent studies showed that the correlation was just coincidental. In fact, another study said, if you drank a lot of coffee, your chances of contracting pancreatic cancer were improved. So, it's just crazy. And I think that, that's one of the problems that we have, but you hold out promise for the future may be, huh?

Eric Holloway:

Yeah. I would say and probably if you restrict the domain enough, you're going to be able to pull out some stuff. But the other problem too, is how they tend to build these systems. They get a dataset from some doctors and then they just go off for a bunch of years and try to make some algorithm that scores highly. What they really need to be doing is working much more closely hand in hand with the doctors and trying to optimize particular parts of their workflow with these algorithms instead of just trying to replace them.

Robert J. Marks:

Okay. So, current and any claims that AI is better than physicians is probably incorrect.

Eric Holloway:

Right.

Robert J. Marks:

Okay. Number seven, AI can implement video games just by watching. This was for an article called Learning to Simulate Dynamic Environments with GameGAN.

Eric Holloway:

Yeah. This one is kind of fun to look at, but you won't really be selling these video games that makes millions of dollars. So, it's able to learn some kind of feedback matrix based on looking at the game screen and the players' input. And so you get something that looks a little bit like Pac-Man or a little bit like that game Doom, but it doesn't stay coherent for very long, like walls will appear and disappear. And a ghost will pop up and disappear. So, it's not super coherent, but because you already kind of know what's going on with Pac-Man, you can kind of squint your eyes and say, yeah, that's a Pac-Man game.

Robert J. Marks:

Oh. So, in other words, they train some artificial intelligence with a number of games and this artificial intelligence creates a game. Is that the idea?

Eric Holloway:

Right, right, right. Yeah. And it's not creating a new game. It's basically just reproducing what it already learned.

Robert J. Marks:

Very interesting.

Eric Holloway:

So, they train it on a whole bunch of screens of Pac-Man and player input. And it just learns how to map the input to different screen frames and kind of finds the gradient between those. So, what they can do with that is they can randomize it and come up with random variance of Pac-Man, but still it remains Pac-Man in general, just a much weirder kind of Pac-Man.

Robert J. Marks:

Now things are much more sophisticated today, but I had a friend, Russ Eberhart who trained a neural network to compose music and he used only four or five songs. And when you listen to the synthesized song, you can hear the refrains very clearly from the original songs. The AI is much better than that today, but it sounds like something similar is happening then. Now they call this GAN. GAN stands for generative adversarial network. Tell me what a GAN is.

Eric Holloway:

Yeah. I believe that's what it is. So, what it does is first it learns like a basic model from its input. And then it generates a new variant of that input from what it learned. And then it learns from that again. And so it's kind of a feedback cycle of it learns a little bit. And then adds that to its data source and then tries to learn from that again. I'm not quite sure how that works so well, because it seems like you'd end up perpetuating errors you learned all over the place.

Robert J. Marks:

Well, you do. I think GAN is the source of the fake faces that we see recorded. Yeah. And that's really interesting. In fact, our editor and our director of Mind Matters, Austin Egbert just published a paper on GAN where he applied GAN to radar sort of data for extrapolating data when it was a little bit sparse.

Eric Holloway:

Oh, very nice.

Robert J. Marks:

Yeah. So, it's still at the formative stages, but GAN has some interesting things, but in terms of this particular application where they tried to extrapolate games, that didn't work too hot, huh?

Eric Holloway:

This is what I see with pretty much all of the convolutional neural network type results. And like the GPT results, which I think we'll talk about a little bit later. If it's generating text, if you look at just a few words or a few sentences, or maybe at the paragraph level and you squint your eyes, you can kind of get something that makes sense out of that. But once you start stepping out and getting the bigger picture, it falls apart, because the neural network is really good at learning these very closely related relationships, but it doesn't really have a concept of the overall structure of anything. And so that's why you see these video games too. Like Pac-Man, he'll

move around and within like four or five squares, you see pretty much the same maze, but once you leave an area and come back to the area, then starts misremembering what it came up with before. It's kind of like a bad dream of Pac-Man.

Robert J. Marks:

Fascinating stuff. We're counting down the dirty dozen hyped AI stories of 2020. And we're on number six. And that's what doctor captain mentioned. Eric mentioned GPT-3. Those are four alphanumeric letters that rhyme GPT-3. And there was a headline that says there's a subreddit populated entirely by AI personifications of other subreddits. Well, first of all, what's a Reddit for those of us who are not socially media savvy?

Jonathan Bartlett:

So, Reddit is just, it's a site where people post links and comment on them. So, it's just, Reddit is grouped into categories and subjects, subject matters, and you can go and find what's interesting to you and people post links and articles. And then you make lots of comments, but it kind of develops a social gathering type of feel. And so basically there was some posters who were posting within some of these subreddits, the subcategories, and it took a while before anyone noticed that these were actually bots that were posting.

Robert J. Marks:

Well, it's interesting that people were able to notice. Now GPT-3 stands for, I looked it up, generative pre-trained transformer 3. And some of the headlines of GPT-3 were just kind of scary. Wasn't this, when it came out, the developer said this might be too dangerous to release, because of all the fake headlines that it would generate.

Jonathan Bartlett:

They've made lots of different claims about GPT-3. And it is indeed. I mean, it's impressive as a demo. I mean, it really does do some impressive text generation. In fact, I think someone actually built a cogeneration system based off of it. So, you could kind of describe in plain words what you wanted the code to do, and it would actually generate a functional code to do what you asked it to do. So, it's actually got quite a bit of a kind of wow sizzle to it, but it turns out that it's not, once you try to get it to do anything serious, it kind of loses its luster.

Robert J. Marks:

Yeah. GPT-3 was trained with billions and billions of articles, including all of Wikipedia and a bunch more. And I think one of the big claims from GPT-2 to GPT-3 was this great, massive increase in the amount of training data that it did. And you could just take a few words and prompt it and boom, it generates a paragraph corresponding to those words. And a review on wired said, GPT-3 was provoking chills across Silicon Valley. But like you said, it was one of these real quick sort of things where you didn't into too much of depth. And I think it was you in your article that you wrote for Mind Matters news said, it's very impressive if you don't look too closely, is that right?

Jonathan Bartlett:

Exactly. It's one of those things where, when people see some of these results, I think people start expecting things that they really shouldn't be expecting from these sorts of systems. For example, one thing that was really impressive is that this is a text processing engine, but it turns out that it can do math. It can do basic arithmetic, but it turns out that once you get past three digits, it doesn't do basic arithmetic at all.

Robert J. Marks:

Oh really?

Jonathan Bartlett:

Yeah. So, if you asked what's the number before a 100, it would tell you it's 99, if you ask it what the number before a 100,000 is, it would say 99,009, which is not the number before a 100,000.

Robert J. Marks:

I see, okay.

Jonathan Bartlett:

Anyway. So, it's one of those things where, because ... I can just imagine somebody, some mid-level manager playing with this and giving it lots of simple arithmetic things. And just assuming that this thing, since it did all of the examples he threw at it actually could do arithmetic. And if he then said, hey, use this as our engine for this. And we expect people throw arithmetic at it, then as soon as they get into four digit numbers, it starts breaking. This is the sort of thing that if you take these systems too seriously, then they can wind up causing damage in the end when you expect them to be more than they really are.

Robert J. Marks:

Yeah. That's what I understand also. GPT-3 was able to write like short paragraphs that were just astonishing in their coherence. But if you ask them to write a chapter, all of a sudden that coherence was lost. Is that a fair statement?

Jonathan Bartlett:

Yeah. Somebody did a, I forgot if it was a, I think they did a series of blogs with GPT-3 and they actually were really good, but it turns out that they said that they did it unedited, but really what they did is that it's unedited in the sense that they didn't actually modify words, but it is edited in the sense that if it said something nonsensical, they would try something else.

Robert J. Marks:

Now, you talked about, there's a subreddit populated entirely by AI personification of other subreddits. That's the title of the article. But you mentioned that somebody noticed that this was a product of GPT, unless somebody volunteered, but how did they know that this was generated by GPT? Is it something that can be recognized by people? Do you know?

Jonathan Bartlett:

I don't remember exactly how they wound up figuring it out. But I mean, at the end of the day, usually AI is wind up saying something that's completely nonsensical. One of the things that GPT-3 does, somebody was poking at it a bit. And if you asked it basic questions about the United States, it could tell you who is the president of the United States in different times, but you could also ask it who is the president of the United States in 1600? And it would give you an answer and not recognizing that the United States didn't exist in 1600 and you could ask it how many eyes does a blade of grass have. And it would give you an answer of blade of grass has one or two eyes.

Robert J. Marks:

I saw that. Yes. I saw that article. Yes. So, yes.

Jonathan Bartlett:

And so that's usually how you wind up sussing these out is they'll start up talking something that sounds logical, but winds up being more or less nonsense.

Robert J. Marks:

Oh, this is one of the limitations of AI. Isn't it? Is it can only think inside the box, it can only interpolate on its training data and extrapolation outside of the box has to require creativity and artificial intelligence doesn't have that creativity. So, that's the reason that could be fooled so easily, or at least that's one of the indicators that it's not as wise as it seems. Okay. Number five, Lack of Sleep Could be A Problem for Artificial Intelligence, as we continue our countdown. Now this is from Scientific American. Lack of Sleep Could be A Problem for Artificial Intelligence. Does artificial intelligence need to sleep, Eric?

Eric Holloway:

Yeah. I looked into this a bit. It's a little bit hard to figure out what they mean exactly by sleep. And it seems to be one of those cases where they're trying really hard to make an analogy behind some kind of obscure mathematical thing to do and everyday life just to make AI sound more human like. My best guess is, well, what they say they do is they train these networks and then they have to subject the networks to waves of noise that, in their opinion, resemble something about the brainwaves during sleep. And then apparently the networks are able to learn more effectively. What I suspect might actually be what they're doing is they're just adding random perturbation to the weights after some training, which is a standard technique. And they just happened to like one particular way of adding noise to the network.

Robert J. Marks:

That's what struck me too. There's a method in training neural networks called simulated annealing, where in, you do basically add noise into the training process to make it much more effective. And there's other things such as weight saturation avoidance, where all of the weights, all of the interconnects are so big that they kind of saturate each of the neurons. And so you have to back them off a little bit. So, you have to halt your training in order to back these things off. But these are problems which have been known for, I don't know, 30 or 40 years, these are techniques which have been, which people have practiced for a heck of a long time. And this is an example of what I referred to as seductive semantics. It's like you said, Eric, that they are

trying to make this thing sound more human. And they do that by trying to relate it to human attributes when the relationship really isn't there, is it?

Eric Holloway:

Right, right.

Robert J. Marks:

So, it's very frustrating. Okay. We're down to the final four. This sounds like a basketball tournament. The final four. Number four of the hype list is Elon Musk is Claiming Self-driving Cars Will be Here Next Year Again. And this was an article, which was, I believe written by Jonathan from Mind Matters news. And I think that self-driving cars have made some advances, but this is clearly clearly hype. Isn't it, Jonathan? Tell us what's going on.

Jonathan Bartlett:

Rob, have you ever seen the movie Groundhog day with Bill Murray?

Robert J. Marks:

Yes, I have. He wakes up to the same environment every day again and again and again, yes.

Jonathan Bartlett:

Yes. And so this is kind of what we have. Elon Musk has been claiming that he's going to have self-driving cars next year since 2016. Now I have to say, part of me loves Elon Musk and part of me can't stand the guy. And I appreciate his humor. I appreciate the fact that he kind of is more approachable than a lot of the other tech billionaires, but there's also this kind of hucksters salesmanship that just really drives me the wrong way. And so he's been saying, he's actually been selling self-driving cars since 2016. People are literally paying him thousands of dollars for this feature that doesn't exist. And he says, oh yeah, it'll be here next year. Next year, I promise. And he says that your car will actually be worth more, most people when they buy a car and they drive it off the lot it's worth less as soon as you drive it off the lot.

Jonathan Bartlett:

And he says, oh, our cars are going to be worth more, because you're going to be able to make money with them by simply sending, when you go to sleep, you can send them out to earn money for you by being a robo taxi. So, you don't have to be there. And he makes claims like this and he makes them every year. And yet he's also, it's not surprising that he's making them right now, because last year he did it right before a $2 billion capital raise for his company. And now he's doing it right before a $5 billion, that's billion with a B, a $5 billion capital raise.

Jonathan Bartlett:

And so he keeps on. In 2016, he said that you're going to be able to summon your car from across the United States. And it would be able to come and get you on its own, finding charging stations on the way. And it wouldn't even need a driver. And said the only thing that could stop that was if we didn't get regulatory approval. Anyway, he keeps on saying that it's going to be

next year, next year, next year, he's saying it again. And anyway, I just wish the media would stop falling for it.

Eric Holloway:

His company supposedly now worth more than Apple.

Jonathan Bartlett:

Oh yeah. So, basically he's got this company. So, Tesla motors is it's a tiny percentage of the car market, but it's basically worth more than the rest of it combined in terms of the value of the stock.

Eric Holloway:

Yeah. Who said you can't make money with science fiction?

Jonathan Bartlett:

Exactly.

Robert J. Marks:

Okay. Well George Gilder, who is one of the co-founders of the discovery Institute in an interview said that Elon Musk is a tremendous entrepreneur, but kind of retarded thinker, which I thought was an interesting statement. And in a conversation that I just had with Gregory Chaitin, that's going to be a podcast which comes on in a while. He said his heroes in life were Stephen Wolfram and Elon Musk. He really thinks highly of Elon Musk and his innovations. And clearly he's done some stuff, but he's also a salesman, isn't he?

Jonathan Bartlett:

Yeah, indeed, he is.

Eric Holloway:

Well, on that note, Elon Musk's innovation, just looked like somebody read a couple of Sci-Fi books and decided to try and sell the ideas to the government.

Robert J. Marks:

Okay. Okay. So, you're really high on Elon Musk. I can tell you. The challenge with self-driving cars, as I've learned is that there's five different levels and there's kind of this mushy fuzziness when he talks about self-driving cars. And he's assuming the top level, isn't he? Self-driving cars that will literally replace the human being in all sorts of environments. And I think that there's a lot of doubt that level five will ever be achieved. In the lower levels, I think I learned from you, Jonathan, in one of your posts at Mind Matters news, that we're actually driving self-driving cars right now, according to the definitions of self-driving cars. Is that right?

Jonathan Bartlett:

Yeah. So, basically self-driving in the lowest levels. It just means that the car is doing some driving feature without you. And so if you think about cruise control, although most car

companies don't use the term self driving to refer to cruise control, that actually technically fits the definition of level one self-driving. Most of the time, if a car company refers to their cruise control as being self-driving, they're usually referring to adaptive cruise control, which also looks at the cars in front of them to see how fast they're going, but really any sort of cruise control technically fits the definition of level one self-driving.

Eric Holloway:

So, if my car were really badly out of alignment and did right-hand turns all by itself, would that be self-driving?

Jonathan Bartlett:

Indeed, it may.

Robert J. Marks:

Okay. Yeah, my car has a mind of its own. My daughter has that thing on her car, which I drove where you put the cruise control. And if somebody pulls in front of you, it automatically adjust your speed to have three or four car lengths. And you can choose how many car lengths there is. And I love it. I don't have to wear out my thumbs and pushing all those buttons and slowing down and speeding up. It does it automatically. So, I like that, but that's at a lower level. And Musk in talking about these things and driving across country is kind of assuming the level five. Isn't that right?

Jonathan Bartlett:

Yes. So, level five means that basically you don't need a steering wheel and you can go anywhere. So, any place that I would normally want to go with my car, there's no limits. I just tell the car where I want to go, and then I can go sleep in the back seat and it will take care of everything.

Robert J. Marks:

My heritage is in West Virginia, in West Virginia there are dirt roads, which are notched out of the mountains. So, imagine a bunch of mountains and you put a little cut in the mountains and those are the roads and they're dirt roads, and they're single lanes and you are driving along and you meet a logging truck coming at you, and you have to scooch over right to the edge of the road, where you're just about ready to fall over the cliff. And that logging truck just sneaks by you. I don't think self-driving cars at level five are ever going to achieve the skill of driving on a West Virginia road. I can't conceive of it.

Jonathan Bartlett:

Yeah. So, the thing that makes me doubt the ability for level five, that specific instance is a good one, but just in general, a lot of our city streets, the way that we drive, the way that the streets are set up, they're geared towards social navigation. That is we understand what the car next to us is doing. We have a kind of, we can look at someone and wave them through. Sometimes you get other hand signals that are not as happy. And so there's a lot of social navigation. Actually sometimes in really congested traffic. People will actually invent a lane. I've seen that happen before in traffic where in really crowded streets, if the street is wide enough cars will just

sometimes decide to, hey, let's add an extra lane to the street and they'll crowd together into a new lane. And so there's all these social aspects to driving that I don't think that you're going to be able to code a computer to understand all of these different social aspects.

Eric Holloway:

Unless absolutely everybody else has their own smart car.

Robert J. Marks:

That's right.

Eric Holloway:

Yeah. So, now they're going to enforce everybody to have their own smart car just so that one smart car will work.

Robert J. Marks:

Do you think that, that's going to happen? Bill Dembski wrote a thing where he said that one of our choices is to either make the self-driving cars smart enough to appropriately navigate, or we are going to have to change the environment and all of the rules to adapt the artificial intelligence to us. And the question is which one we do.

Jonathan Bartlett:

I think changing the environment is the one that's more likely to happen. And that's level four. So, level four is basically where you say within these defined parameters, the car will drive itself. With the additional stipulation that if the car ever goes outside those parameters, it has a safe enough way to get out of the way out of the traffic so that you don't have to immediately assume control. So, if you can imagine, let's say you might have a level four that can navigate neighborhoods. So, it's going at a low enough speed that if it ever encounters a situation that it doesn't know how to handle it could simply pull over and stop and wait however long you needed to go and for you to go and assert control over the vehicle. But level four, the car is doing all the operations.

Jonathan Bartlett:

You can sleep in the back, but there's only a limited segment of road or a limited set of environments where it works. And that's kind of where we're most self-driving that's been successful has gone, is they've done high resolution maps of areas. They've determined that in certain locations, there's not going to be a lot of unexpected things happening. They have streets that are easy to navigate. They're at low enough speeds that they're not going to hurt anybody. The roads are isolated enough that you're not going to worry about pedestrians accidentally coming across suddenly. And so by mapping it out and having enough knowledge of the environment, they can make a car for that environment. And that's generally what they've been doing when they're successful.

Robert J. Marks:

Jonathan, we just did a paper with a student of mine, Sam Haug. And it was about the idea that the more complicated an AI system is the more contingencies that you have. And many of these are unexpected contingencies. So, if you have a broad AI system, you're going to have all sorts of things, which the AI is not programmed to respond to and it's unavoidable. And it requires a heck of a lot of tests. So, this idea of the environment fooling you is very real. There's going to be lots of situations, lots of scenarios that are unexpected. So, we have that Elon Musk is claiming self-driving cars will be here next year again, reliving Groundhog day as Jonathan said, but we are making some advances in self-driving cars, but maybe not at the level five. Number three, can AI really know when it shouldn't be trusted. The title of the article from Science Alert is, Artificial Intelligence is Now Smart Enough to Know When it Can't be Trusted. Eric, what's going on here.

Eric Holloway:

Well, first of all, I'd like to note that the title does not say that AI can know when it should be trusted. So, you could just have an AI that says never trust me. It's always going to be right.

Robert J. Marks:

That goes back to fundamental detection theory, right? You have a 100% detection, but you have a very high percentage of false alarms too. Huh?

Eric Holloway:

Yeah. Now, as to what they actually did, they added some kind of confidence level to their results. So, if it's really low confidence, then you know you can't trust it, but the converse does not apply. They can't say that when it has high confidence that you can trust it. There's a very solid well-proven theorem called Gödel's second incompleteness theorem. And it says for any system that can reliably tell you that things are true or false, it can not tell you that it itself is reliable. So, if they ever did create an AI system that can tell you, oh, you can trust what I say. Then at that point you can precisely cannot trust it.

Robert J. Marks:

It reminds me of the credence paradox. He says, everything I say is a lie. So, that's where you're getting to. Is that right?

Eric Holloway:

Yeah. Also, so, let's put back down to a more practical level. Let's say it does have some kind of confidence level and can say it's fairly non-confident about some results than others. You still may not even want to trust that lack of confidence level. There's another theorem called Rice's theorem, which says any non-trivial property of a program is impossible to program itself. So, you can't have a program that can always say that, hey, my confidence level is reliable. So, if they can precisely set it up in a constrained environment, then you can probably get some kind of confidence out of it. But it definitely, they cannot do anything like the headline claims, which is a artificial intelligence that is smart enough to know when it can't be trusted as weighed in general to be something you can actually do with computers.

Jonathan Bartlett:

What's even worse is if you read the first paragraph of that article, the first paragraph of it just goes to the total science fiction land. It says, how might the Terminator have played out if Skynet had decided it probably wasn't responsible enough to hold the keys to the entire US nuclear arsenal? As it turns out, scientists may have just saved us from such a future AI led apocalypse by creating neural networks that know when they're untrustworthy.

Robert J. Marks:

Oh, good grief. Okay. Yeah. One of the big problems with the AI hype is the confusion of science fiction with science fact, and people need to be more cognizant of that. We're counting down the dirty dozen hyped AI stories of 2020 with Eric Holloway and Jonathan Bartlett. We're on number two. And this one just kind of makes me mad. Number two, Sam Altman's Leap of Faith. Eric, what is going on here with Sam Altman, who is he? And what's his leap of faith, which is totally incorrect, I believe?

Eric Holloway:

I would actually say Sam Altman is totally correct. He's actually taken the AI kind of trend to its logical conclusion, because if AI is truly as great as it should be, like we can actually reproduce human intelligence and then it could feed into itself and then take off forever. Then the crazy claims he's making here are actually correct. So, I would say it's not Sam Altman that's crazy. He's just the logical conclusion of a crazy movement. And he says stuff like, I'm only going to focus on creating AI, because once you get AI, it's going to embed absolutely everything else. He calls it the light cone of the future. And then he makes these funny venture capitalist sells, like, instead of saying, hey, we're only getting you to give you a certain percentage of the profit. He says, well, once you get a hundred times return on what you invest in us, then we're going to have to give the rest to charity. Like he's over promising in kind of a, trying to undersell his over promising. Pretty hilarious.

Robert J. Marks:

Now this guy is no slouch. He is the, what is he? The president of OpenAI or something like that.

Eric Holloway:

Yeah, well, no, he has a fantastic history as a great venture capitalist. He came up with some company called Loopt when he was just in his early 20s that he sold for millions. And then he took control of YCombinator, which is one of the most successful venture capitalist firms in Silicon Valley, which has a pretty nice lean startup approach, or at least they used to. And then he took that approach and even made it better. So, he has a great background. And so that's why I say, he's not crazy. It's the movement that he's kind of heading up. It itself is crazy. And he's just taking it to its logical conclusion.

Robert J. Marks:

A friend of the Bradley Center, Roman Yampolskiy on April fools, put out a tweet on social media. And he said, "This is incredible. Google fires all of their programmers, because they have developed a super AI that will write all of the programs of the future." And if you just think about that, it's just really ridiculous, yet he got a lot of thumbs up and he was even contacted by

people in the media that said, we want to talk more about this. And he said, look, it was a joke. It was simply a joke.

Eric Holloway:

Yeah. And if you, so let's look at this from a narrow perspective. If I told you, hey, I have this neat little black box and you can plug anything you want into it. And this little black box will power it forever. It just creates an energy out of nothing. No one would take me seriously, but what Sam Altman is claiming is exactly equivalent of that. But in information theory, instead of with energy. And actually that if he was right about information theory, then you could probably actually turn that into a source of infinite energy too. So, they've essentially the perpetual motion machine for computer science.

Robert J. Marks:

That's really, that's very interesting. And of course, this idea of AI writing better AI that writes better AI assumes that AI is creative. We don't have time to get into the so-called Lovelace test, which is a measure of whether AI is creative or not. But according to the Lovelace test, artificial intelligence has yet to be creative.

Eric Holloway:

Yeah. And in fact, well, the things we were just talking about, like the open GAN generating games and GPT generating texts, at least GPT, actually that's part of the Sam Altman's company and all of his AI advances, even though they're pretty remarkable in themselves, they illustrate exactly this. The only things they're doing is regurgitating all of their training data, just a more finer grain interpolation between data points. But it's all just reproducing what somebody else wrote. There is zero creativity and these AIs that have come out.

Robert J. Marks:

Wow, it's really an embracement of materialism and determinism, isn't it?

Eric Holloway:

Yeah. Yeah. And the ironic thing is that the more they'd buy into materialism, the less they actually create.

Robert J. Marks:

Right. And I think that our stance is well-grounded in computer science and why people don't recognize this, I don't know. There's lots of people that believe AI will never be creative. This includes the recent Nobel Laureate, Roger Penrose and his book, an Emperor's New Mind and Satya Nadella, who is the CEO of Microsoft said basically the same thing. He said, "In the future we're going to do a lot of things with artificial intelligence, but creativity is always going to belong to the programmer." So, there's lots of people that understand the limitations of AI. Yet, there is still this, I don't know, theology out there that we're going to reach this idea of a singularity. No, it isn't going to happen. It isn't going to happen.

Eric Holloway:

Yeah. And I would say it is actually close to kind of a religious belief, because I had this conversation with other people and I'm like, well, I'm skeptical that the mind can be reproduced with the computer, but then they'll say, well, I mean, we evolved and all the things that evolved ultimately came from just a physical laws and atoms bumping into each other. So, at least in theory, we should be able to create AI. So, it is a logical deduction from a certain frame of reference.

Robert J. Marks:

Yeah. It's unfortunate. I still maintain AI will never be creative. It'll never be sentient, it will never experience Qualia. It'll never understand what it's doing. It'll add the number seven and three, but it doesn't understand what the number seven or three are. So, that's the limitations of AI, which is unfortunate and apparently not recognized by other people. Okay. Number one, this is the number one of the dirty dozen hyped AI stories of 2020. And the number one has to do with Elon Musk again, by the way, number two, the Sam Altman, this OpenAI, that was an Elon Musk venture, right?

Eric Holloway:

Yeah. He collaborated with Elon Musk. Elon Musk is claiming AI is the biggest existential risk the human race faces. So, he wants to make friendly AI.

Robert J. Marks:

Yeah. And the interesting thing, he never talks about what the second most existential risk for humanity is. And I would actually put thermonuclear weapons as a more of an existential risk.

Eric Holloway:

Also think about what they're claiming to be creating here. They're claiming to be creating a entity that is all powerful, all knowing. And since it's friendly, it's also all loving. So, what would be another name for such an entity?

Robert J. Marks:

Oh, God.

Eric Holloway:

So, they're basically trying to create their own God.

Robert J. Marks:

Yep. And in fact, that's the topic of the book by John Lennox, 2084, where he talks about artificial intelligence and some of the hyperbolic claims, which are made about the future of artificial intelligence. Okay. Number one has to do with Musk again, mind games, Elon Musk wants to connect your brain to a computer this year. And he says, it's going to be awesome with his so-called Neuralink mind chip which he is preparing to launch. I've read some about this. He's implanted it in some people, hasn't he? This is a report from the US Son. Jonathan, what's going on here?

Jonathan Bartlett:

So, Neuralink is basically what he's done is he's created these ultra thin wires and kind of a robotic sewing machine that can insert these super thin wires into a brain. So, Musk's ideas that he basically believes that everything interesting that happens in your brain is basically electronic signals. And therefore, if he can get electrodes in there, anything that's wrong with your brain, if you can get enough electrodes pumping data fast enough that he could fix whatever's wrong with your brain by simply offloading it to some sort of an external processor. And so that's the idea of Neuralink that. So, he can basically make a Jack that connects an external computer into your brain and take over functions.

Robert J. Marks:

That's really strange. I think I'm already connected to a computer, but I don't have to have a chip on my brain. I use my fingers on my keyboard that that links me to all of the knowledge in the world. When are you going to get your implant?

Jonathan Bartlett:

I don't think I'm really trusting anybody with that anytime soon. That seems a little invasive. Although some people are clamoring for it. They're like, yeah, I want to enhance myself. And part of me wonders what kind of, if there are self-esteem issues that are circling around that.

Robert J. Marks:

That's interesting. Do you know if he's had any success at all in this Neuralink transplant?

Jonathan Bartlett:

As far as I'm aware, they haven't done anything with humans yet. I could be wrong on that, but generally they've been doing rats or mice. Anyway, this all kind of traces back to, I wrote an article about halfway through the year, both level five self-driving and Neuralink, both have an interesting connection with them. And that is this myth about the mind that the mind is just basically a computer processor. And this is kind of what you and Eric were talking about a moment ago, this myth about the mind that all it is, is just extended computation. And so for Musk, anything about the mind that's wrong, he can fix because for him, everything about the mind is signals. And so all he has to do is get something attached to your brain that's processing signals fast enough, and he can fix it.

Jonathan Bartlett:

Now that's a presumption. It's actually a huge presumption. I imagine he's got to know that, that's a big leap of faith, but he's pushing it as if he knows that, that's the answer. And that's the thing that's frustrating is that the things, the claims that he makes for this are just outlandish, because he goes into things that we actually don't even know what the causes are. And he claims that Neuralink will be the solution. And so if, to say that a device that is not even been tried out is the cure for something for which we don't know the cause, that seems a little over-hyped to me.

Eric Holloway:

He should rename his company, 42.

Robert J. Marks:
Is that the Hitchhiker's Guide?

Eric Holloway:
It's the is the answer to life universe and everything. 42.

Jonathan Bartlett:
Yes.

Robert J. Marks:
Is that the Hitchhiker's Guide To The Galaxy?

Eric Holloway:
Yeah, Hitchhiker's Guide To The Galaxy [inaudible 01:05:09].

Robert J. Marks:
Oh, okay. Oh, that's kind of funny. Intelligent design. I think there's three reasons that we can have this complexity that we observe. One is an intelligent creator. The second one, which has purported is panspermia that all of this complexity was planted here on earth by some aliens. Elon Musk actually put forward a third hypothesis of intelligent design, which is that we are all simulations. We are all computer simulations. We live in a big Sim world. And I wonder how his Neuralink ties in with his theory that we are all simulations, any ideas?

Eric Holloway:
Well, I think he's actually ... I don't think it's another company, but he's funding individuals who are trying to find bugs in reality like the old movie the Matrix.

Robert J. Marks:
Wait, wait, bugs in reality.

Eric Holloway:
Yeah. Well, that's the conclusion. If you think we live in a computer simulation and then presumably it's written in some kind of code, and if the coder is not perfect, then there's going to be bugs in our simulation. So, he's trying to find bugs in reality, kind of like the glitch in the matrix.

Robert J. Marks:
Or kind of like the Truman show when that big thing falls out of the sky. My goodness.

Jonathan Bartlett:
Now there's kind of a faulty logic that goes to why a lot of people think we live in a simulation and I'll give you the logic then I'll tell you what the problem is with it. And that is that if you imagine that we could simulate a universe. Well, so, let's say that there's only one actual

universe, but then we figure out how we can simulate a universe. Well, as soon as we can simulate a universe, if we successfully simulate that universe, that means that in that universe that we're simulating, there are going to be creatures who figure out how to simulate a universe. And as soon as that happens, we're going to have more simulated universes then we have actual universes. And therefore your chances of winding up in the simulated universe are actually much higher than your chance of existing in the actual universe.

Jonathan Bartlett:

And so that's the logic that's oftentimes used. So, the problem with that is that it always takes more stuff to simulate something then the thing that you're simulating. So, for example, I can make a model of atoms moving around, but it actually requires entire computers, which are all made of trillions of atoms to make that simulation. And so you actually kind of wind up with a space problem that you actually, you can't simulate as much as you have reality. And so even if you could make a perfect simulation of reality, it would have to be a smaller reality than what you're simulating it.

Eric Holloway:

What if it was a bunch of nested lossy simulations?

Jonathan Bartlett:

That's possible, but then you'd have to ... You'd wind up being really lossy really fast.

Robert J. Marks:

I'm sitting here trying to getting back to the Neuralink, trying to understand what the Neuralink would do to me. Currently I can only keep a couple of things in my brain. Like if I multiply two three digit numbers, I have to write them down and I can't do the whole thing. I do it kind of one step at a time. Going through all of the multiplication processes that the little algorithm that we use to multiply two three digit numbers. And so the brain only has this capacity of keeping kind of one thing in the forefront of your mind at one time, I'm trying to understand how Neuralink would improve that. I'm not sure, maybe it can, maybe there is something that can be done, but do any of you have any thoughts on that?

Eric Holloway:

It would make identity theft really interesting.

Robert J. Marks:

How is that?

Eric Holloway:

Well, let's say it worked and I mean, hackers are hacking all our bank accounts, then next they'll be hacking our brains and taking over our actual bodies.

Robert J. Marks:

Oh my goodness. Do these Neuralinks, do they have any wireless connections? Do you know? I hope not. I hope not. Hey, we've been working our way through the dirty dozen hyped list with Bradley Center Brain Trust members, Eric Holloway, and Jonathan Bartlett. We are not going to be totally negative. We are on a subsequent podcast going to go through the top 10 smash hits of artificial intelligence for 2020. There's lots of exciting stuff happening in artificial intelligence. So, until then 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 Egbert. The opinions expressed on this program are solely those of the speakers.