AI Dirty Dozen 2020 Part I (https://mindmatters.ai/podcast/ep113)

Robert J. Marks:

We count down the top dozen hyped AI stories of 2020 today on Mind Matters News.

Announcer:

Welcome to Mind Matters News, where artificial and natural intelligence meet head on. 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, medias everywhere and competition is fierce, articles with provocative headlines and content are clickbait for the browsing consumer. So we're going to count down today the AI dirty dozen, the top AI hype stories for 2020.

Robert J. Marks:

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. 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 in 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 will 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 a 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.

Eric Holloway:

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. Does that make sense?

Robert J. Marks:

Yep.

Eric Holloway:

Okay. Now, I ask 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.

Eric Holloway:

The models themselves are so, so, so complex that even with billions or trillions of data points, the 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. 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.

Eric Holloway:

And these models will have very different predictions on new data that is not contained in the dataset. 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.

Robert J. Marks:

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 networks."

Robert J. Marks:

So they changed the neural network and then they tested it on the same data. And then he says, "Well, this works a little bit better." 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:

The other thing is that anytime you have a model, there are 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.

Jonathan Bartlett:

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. 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 like no relationship between data and reality once you step a moment outside of those bounds.

Jonathan Bartlett:

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'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. Even if we could have an 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. This is a paper from a very prestigious journal, Nature, and it questioned 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 like it's an AI first company. I think Facebook has too. It's now a really big part of their brand. 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 the 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 of millions of dollars and use massive computer farms.

Robert J. Marks:

There's an old saying in engineering that "In theory, theory and practice are the same. In reality, they are 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 a little bits and pieces of it enough that other people will start like getting addicted to Google, but not enough that we can really do what they do. Yeah, 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, yeah, what I'm talking about though is like the specific technique.

Eric Holloway:

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. 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.

Eric Holloway:

It's ironically like 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 an article from a very prestigious publication, Scientific American, 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, what do you think?

Eric Holloway:

This actually directly builds on what we were just talking about. 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, 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 like really hyped.

Eric Holloway:

They're going to change the way the world is and everything. And after the fact, once they've actually started trying to use their product in reality, then all of their venture capitalists decided, "Yeah, 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:

Yeah, 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. 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 Marvin Minsky and Claude Shannon and some of the luminaries of information theory. And they were like, "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 that quite. No, I don't have the specific date off of my head. 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... Actually eight decades after that.

Robert J. Marks:

Yes. Yes. 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.

Eric Holloway:

Right.

Robert J. Marks:

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. 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 turn out to only actually work in very, very small domains and very restricted data sets.

Eric Holloway:

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 at 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. 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 scifi film. This was reported both at Mind Matters News.

Robert J. Marks:

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 scifi film, a $70 million scifi film. Eric, what do you make of that?

Eric Holloway:

Well, first of all, I asked, why is this news? We've had animatronics in movies since like Star Wars or know The Muppet movies. This is just a fancier Muppet movies.

Robert J. Marks:

That's right.

Eric Holloway:

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 anthropomorphization 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 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:

Yeah, 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. 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 its Hall of Presidents' 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 referred to as seductive optics in these presentations and these movies and some of these hype versions, wherein 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 robot Sophia was nothing more than the animatronic robot that synchronized their mouth movements and their facial expressions in order to communicate. The optics, which was to have it inside a human form, had nothing to do with artificial intelligence.

Robert J. Marks:

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 not like a person controlling the bot behind the scenes? This is actually really, really old, 1700s or so, they had a supposed to 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. This is really old.

Robert J. Marks:

I remember that.

Eric Holloway:

The more things change, it's more the same.

Robert J. Marks:

Isn't that, in general, true in a sense to today's artificial intelligence? All of the intelligence we see is due to the computer programmer asking the artificial intelligence to do something. And we might be surprised at the output. I mean, you talked 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 that software was trained to play Go and that's what it was doing.

Robert J. Marks:

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. It's surprising. It's kind of cool, but certainly it's 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 to even have to achieve that result, they go through like billions or trillions or quintillions 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. 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. Well, we have counted down from 12 to nine of the dirty dozen hype list for 2020 with Bradley Center Brain Trust members Eric Holloway and Jonathan Bartlett. Next time we meet, we'll go through the next four, and we will eventually get all of the top 12 hype stories from 2020 on Mind Matters News. And so 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 Egbert. The opinions expressed on this program are solely those of the speakers.