AI Smash Hits 2020 Part II (https://mindmatters.ai/podcast/ep117)

Robert J. Marks:

We continue our countdown of the top 10 AI smash hits of the last year or so, 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, for our countdown we are joined by two members of the Bradley Center Brain Trust. Jonathan Bartlett is Director of the Blythe Institute and a Senior Fellow of the Bradley Center. Welcome Jonathan.

Jonathan Bartlett:

Thanks for having me on Bob.

Robert J. Marks:

And Dr. Eric Holloway works for the NIH and is a current captain in the United States Air Force. Dr. Holiday is also a Senior Fellow of the Bradley Center. Eric, welcome again.

Eric Holloway:

Thank you and that will be my presidential slogan. It will be, holidays for all from Dr. Captain Holloway.

Robert J. Marks:

Dr. Hollaways for all, Holidays.

Eric Holloway:

Yeah. Holidays for all.

Robert J. Marks:

Oh, your name does rhyme with holidays, right? So you take a Holloway holiday, I guess, right?

Eric Holloway:

Yeah, yeah exactly.

Robert J. Marks:

Okay, excellent. Well, it's good to have you back. We're counting down the top 10 list. We've gone through number 10 to number six, and we're going to finish up the list today. Just as a reminder, we consider only AI accomplishments and not great promises of the AI future. If you look at most of the top 10 AI lists that are on the web, they all have hedges about their adaptation in the future. As you know,

we are surrounded by great AI that we take for granted. It's everywhere and we're numbed by familiarity but if you think about it, we have Alexa, we have Google maps, Uber, GPS, even cell phones.

Eric Holloway:

My rice cooker.

Robert J. Marks:

Your rice cooker? Oh, that's right, it has fuzzy logic in it. That's right.

Eric Holloway:

Yeah, my laundry machines or a thermostat. Yeah, thermostat, the best AI have.

Robert J. Marks:

What are some other things? Face recognition, like when you sign onto your computer and as we talked about last time. No, I don't think we talked about last time, but as voice to text transcription, that comes to mind. And these are all things which we have right now that can be labeled as artificial intelligence, but we take them for granted. The question is what amazing AI will we have in the future?

Eric Holloway:

Yeah, one of the biggest ones we all use is Google. They've done great stuff using machine learning in huge corpora of texts.

Robert J. Marks:

You're right. And we take it for granted. We take it for granted that we can get 20,000 responses to a query in just a few milliseconds. It's just amazing.

Robert J. Marks:

Number five deals with application of artificial intelligence in entertainment. Walt Disney has a great history of applying technology to entertainment. They had automatons in the 1950s in the Hall of Presidents, most Disneyland, and later on at Disney World, where the Presidents came out. They did gestures, their mouth moved in accordance to what they were talking about. Disney had some patents on it and Disney again is a leader in application of technology to entertainment. So number five on our list is deep faking entertainment. We all know about deep fakes where we can generate images, which really don't exist. John, what are deep fakes in AI and what is Disney doing that's going to wow us?

Jonathan Bartlett:

Disney has already done some of the deep fake research. I remember going and seeing some of their... They had some animated stuff at Disney World many years ago when I went there, where it was kind of animated in real time. The deep fakes are really interesting because we've had an explosion in the popularity of some of the deep fake things where you have the Reface app that takes your face and puts it on some movie stars in clips and special movies. And that's kind of really grown in the popular imagination.

Robert J. Marks:

Is that available for free?

Jonathan Bartlett:

I don't know if it costs money. I'm pretty sure it's free though. I've never used it myself. I've just seen... My Twitter feed is full of people putting their face on various things.

Robert J. Marks:

It also shows you have a good self-image, right? Okay. So go ahead.

Jonathan Bartlett:

So, deep fakes ... People are worried about the potential for using them for evil and that's definitely a worry because you could ... I've seen deep fakes of people making Obama or Trump say all sorts of awful things. And if you weren't aware of the technology, you might think that there really were videos like these that existed. They do make it hard for people to recognize truth from reality.

Jonathan Bartlett:

But there's also a lot of practical applications people can do as well. It speeds up animation. You can think of animation as a giant deep fake project. And so the ability to do real-time deep fakes helps people do some filmography and some special things. But another really interesting advance in deep fakes is for compression. You basically ... What some people have figured out is that you can, you can basically deep fake yourself. And basically with the deep fake technology, it requires fewer bits once you have a baseline image to translate the changes in your face and whatnot over the wire than it does to transmit actual video.

Robert J. Marks:

Okay. Let's talk about compression. I usually explain compression as the idea that it's like transporting dehydrated food. You take the water out, so it's cheaper to ship. And then at the end you put the water back on at the destination. That's what compression is motivated by. Is that right?

Jonathan Bartlett:

That's a really good analogy for it. Yeah. And the problem with compression in general is that there's no general way to compress things. There's no generalized algorithm that will compress any stream of bits. But the nice thing is that usually what we want to transmit is not any stream of bits. It's usually very specialized streams of bits.

Robert J. Marks:

What about zip files or PNG images? They use a common compression algorithm, don't they?

Jonathan Bartlett:

Exactly. So the compressions that we generally use is because the bit streams that we have in our files are not just any bit streams, they usually follow patterns. So for example, I can zip up my text file and make it really, really small because I'm using texts, which is only a subset of the bits available for what I'm doing and and then I'm writing them in words, which will make it more regular. And I'm putting those words into some of which are really common sentences, which make it compressible. So each of these levels of expectation allows you to compress your signal to some degree. And so basically deep fakes do, is they separate out at a really deep level. The bits that are background and the bits that are needed for the foreground. And honestly your mind actually does a deep fake as well.

Robert J. Marks:

Oh, how that?

Jonathan Bartlett:

Our connection between our eyes and our brains are not as high bandwidth as you might imagine. And so basically when you look straight ahead, the optics are focused on what's straight in front of you, but your mind is putting together a lot of what's around. You're actually seeing more than you can actually see because your mind is basically faking some of it for you. So anyway, so that's what deep fakes do is they take a small amount of data and they separates out different pieces of it and can replace the parts of it that more or less matter.

Robert J. Marks:

Eric, how is Disney using deep fakes in entertainment?

Eric Holloway:

Well, Disney is using deep fakes and entertainment as a way to capitalize on not having to hire lots of really expensive actors. So you can have a few expensive actors, they do their thing, and then you copy their body movements and face. And now you can just hire a bunch of cheap actors and stick the expensive actors faces on them. Or you can go in other directions like you can stick cartoon characters on them and you can make animation a lot simpler for cartoon characters because now you can use human bodies to do your animation for you and then just throw a cartoon suit on them virtually. So there's a lot of possibilities here.

Robert J. Marks:

That's really interesting. I wonder if we'll ever have a deep fake superstar. And I say that it might sound funny, but if you think of, for example, the brand of Betty Crocker. Betty Crocker, who sold things like cake mixes and things of that sort, was a brand that was that of a lady and a picture of a lady that was totally made up, but the brand itself became worth millions and millions of dollars.

Eric Holloway:

Wow.

Robert J. Marks:

I wonder if we're going to have deep fake superstars. That would be really interesting.

Eric Holloway:

There's actually a music band like that. The music band had ... It consists of four cartoon characters and they've made a couple of hits that have gotten to number one, but the actual human behind it is a single guy who just makes it all up on his own.

Robert J. Marks:

Really?

Eric Holloway:

Yeah. Called the Gorillas.

Robert J. Marks:

Gorillas. Okay, well that's good to know. So let's all be watching out for this application of deep fakes in Disney entertainment. Fascinating. We are counting down the top 10 smash AI hits of 2020, and we're up to number three, paralyzed man moves in mind-reading exoskeleton. This is exciting. This is where AI is helping the handicap. Eric, tell us what happens here.

Eric Holloway:

Yeah. Now this is a really practical and really useful application of AI. There's pretty much no other way to do what they've done except with the use of AI and something like Elon Musk's Neuralink. It's not ... I don't think it's as invasive, but this might actually work better with the Neuralink. But basically they stick a number of probes into this man's brain. And so they can read the brainwaves and then they have a machine learning system that can learn what his brainwaves correlate to in terms of body movements. And then they take the machine learning model and they use that to control an exoskeleton. So this is a man who's completely paralyzed. Who'd have no way of moving otherwise, but they can hook his brain up to this machine learning system, which moves the exoskeleton. And now he can actually walk around and move his arms.

Robert J. Marks:

Does the exoskeleton learn?

Eric Holloway:

I don't believe the exoskeleton itself learns. I mean, it's not in production yet. They will have to do a lot more. So I think it's pretty limited they have to do a lot of training time with him offline. I think he had to play a video game a whole lot. And then once the AI gets some kind of idea what his brainwaves mean, then he can have some kind of rudimentary control over the exoskeleton.

Robert J. Marks:

The brain and neuroplasticity is really amazing. And if you have a lot of your brain, which is dedicated to something such as body movement, you're not using it, it often adapts to other things. So I can see the neuro-plasticity adapting so that it control the exoskeleton. So the adaptation would not be in the exoskeleton itself, but it would be in the brain. The neuroplasticity of the brain.

Eric Holloway:

Yeah, that's actually a good point. I think maybe an even more effective route they can go with this is if the exoskeleton can feed him some kind of control signal, which he learns how to manipulate and learn to move the exoskeleton himself. Because what's even more impressive than say artificial neuro networks is the real neural network. The brain plasticity you're talking about, there's a doctor named Norman Doidge and he has a book called The Brain that Changes Itself. And these are fascinating accounts of what you can do with brain plasticity. There's one lady who was born, I think, without any sense of balance so she could never stand up. So he gave her a buzzing device that she had hold in her hand. And once she got off balance, the device would buzz and she was able to retrain her brain and actually regain a sense of balance. So she didn't even need that buzzing device anymore.

Robert J. Marks:

Yes. In fact, we have had a podcast with the neuroscientist Yuri Danilov who was one of the founding scientists of that. And also a gentleman named Sackier, Dr. Sackier who's head of the company that markets this. This is astonishing stuff. They mentioned that the tongue itself has more neurons per inch than any other part of the body. And so therefore, if you stimulate the tongue, you are stimulating a lot of nerves. And the other thing that was mentioned is that the tongue, when you develop is basically just pulled out of the brain. So the tongue has all of these nerves, which go directly to the brain. So that's the reason that these tongue vibration things work so well.

Eric Holloway:

So the tongue is actually part of the brain?

Robert J. Marks:

Yeah. Well, in a way it's kind of part of the brain that as you develop, it's pulled directly from the brain and you have a lot of neurons, which are a lot of connections, which go directly to the brain. Yeah. That's just fascinating stuff. You know, I got so exuberant that I skipped a number four didn't I?

Eric Holloway:

Yes. You skipped-

Robert J. Marks:

I did skip number four. Now there was no guarantee in this countdown that they would be in order. Did I say anything like that? No, I don't think so.

Eric Holloway:

I think that's implicit in a countdown.

Robert J. Marks:

Okay. Well, yeah. Okay.

Jonathan Bartlett:

It's a count meandering.

Robert J. Marks:

It's a count meandering. You're right. Count down means yes, count from the top to the bottom. You're right. You're right. Okay. So we're going to continue our count meandering with the top 10 smash AI hits of 2020. This story actually dates back to 2019 and it is deep learning for, I hope I can say this right, leukocoria or white eye. White eye is something which occurs in children and it's due to a pale reflection from the back of the eye and is a precursor to, here I go again, retinoblastoma, which is a fancy word for eyeball cancer. Now, the story behind this number for smash hit of 2020 is very local. Very, very close to me. There's a professor at Baylor University in chemistry named Brian Shaw. He had his son Noah who lost an eye to cancer, and he became dedicated to the idea that he wanted to develop something to prohibit this from happening in other people. Not prohibit, but at least have an early detection.

Robert J. Marks:

So he contacted a professor in the department of computer science of Baylor University, Doug Hammerly. And they applied a deep convolutional neural network to look at images of kids' eyeballs, to tell whether or not they had this condition called white eye. And they developed it. And it was very, very successful. And now is available as an app, a free app. They decided not to make it commercial and try to make a much of money out of it. But this is an app that you could get on Google Play or the Apple app store. And the name of the app is called cradle, C R A D L E. And what it does is you put the app on your cell phone and it scans the images on your cell phone of your child and tells you whether your kids, the pictures of your kids on your cell phone, have this white eye this indicator of eyeball cancer.

Robert J. Marks:

And I thought this was just an astonishing, wonderful application of artificial intelligence. So that's number four. So we continue to count down the top 10 smash AI hits of 2020. We've been to through number three. Now we're at number two. This is Carnegie Mellon and Facebook AI beats professionals in six player poker. This result astonished me. I ever heard poker players say that poker is not a game of rules. It is a game of bluffing, of psychology, and apparently not. AI was developed called Pluribus. And in the game of Texas Hold'em, Pluribus was able to beat one-on-one professional Texas Hold'em players. And it is interesting that this world series of poker that the same people show up year after year after year and Pluribus bus beat Darren Alias, I don't know these people. I don't watch the world series of poker, but he holds the record for the most world poker tour titles and Chris Ferguson. He was winner of six world series of poker events.

Robert J. Marks:

Each pro separately played 5,000 hands of poker against five copies of Pluribus and Pluribus won. Now this in itself was an astonishing result. What was more astonishing is the Pluribus won in a game with five different pros at the same time. So there were six players. One of them was Pluribus and they played a total of 10,000 hands. And again, Pluribus emerged victorious. This to me was an astonishing result and lets me know that winning at poker is very highly algorithmic. There's some randomness that goes in here for sure, but the fact that artificial intelligence can win at Texas Hold'em, to me is astonishing.

Robert J. Marks:

One of the things they used, which I hadn't seen used before was something from game theory called the Nash equilibrium. Possibly you saw the movie a beautiful mind, starring Russell Crowe. It was a story of John Nash who had some mental problems, but was a genius mathematician and came up with the idea of a Nash equilibrium. And he won a Nobel prize for doing this and the Nash equilibrium, again, as a game theory sort of concept, which was applied to winning this Texas Hold'em and was a technique used by Pluribus in doing this defeat. Again, I think that's an astonishing result. John, Eric have either of you heard about this?

Eric Holloway:

I think I saw previous results where they'd beaten two player games, but yeah, I haven't seen this one with five players before. It's very interesting. And I'm also very curious about these, all of these results where AI beat humans. Because you you're only ever told the end result that the AI beat the human and you're given some kind of insight into how their play style differs from the human, but you don't really know much about what exactly is going on under the hood.

Robert J. Marks:

That's very interesting because we saw nothing about Pluribus playing Pluribus, right? What would happen then?

Eric Holloway:

Yeah. Well, I think that's actually part of its training. It plays itself a number of times to... It's like the AlphaGo Zero that plays itself a whole bunch in order to develop its strategies.

Robert J. Marks:

Yes.

Eric Holloway:

So I think that's what's going on here. Yeah. It's just very interesting to think about, because like we were talking about with the AI hacking, when you have these really complex AI models, they have these blind spots where if you know where the blind spots are, you can poke those spots and make them do what you want. So I'd be curious if down the road with all these game-playing AIs, if people start finding out these blind spots in the AIs and figuring out how to control the game AIs.

Robert J. Marks:

Yeah, that's a fascinating observation. We will see, I suppose. We are counting down the top 10 smash AI hits of 2020, and we are down to number one. This is maybe the most powerfully impacting artificial intelligence result of last year. Protein folding. AI has cracked a problem that stumped biologists for 50 years. And it's a huge, huge deal. Jonathan or Eric, elaborate on this a little bit.

Jonathan Bartlett:

Well, protein folding has been a tough problem for biology for a long time, just because there is all the interactions. It is hard to predict exactly how a protein is going to actually fold. So typically, or historically they've done it by x-ray crystallography where they basically shoot x-rays at proteins and watch them bounce off. And then guess what the protein looks like based off of these x-rays. But what they really always wanted to do is be able to guess what the structure is just from the sequence. So for those who don't know, you have DNA and DNA is basically this long strand of what are called bases, which are basically the letters of DNA. And so those letters of DNA get translated into proteins, which again, there's just a long strand of proteins and there's again, basically letters of proteins that are just connected all along.

Jonathan Bartlett:

But unlike the DNA, the proteins actually do things. They connect with each other. They have interactions between the individual amino acids. The problem is, is that, there's all sorts of interactions that might happen between these different amino acids that would cause the protein to go into different shapes. And so the question is, is which way will it actually fold? And being able to suss that out has been a difficult problem. And the biologists always want to just be able to see the sequence and infer what the final shape is going to be. So that's what they've been... That's been the problem and humans have generally been bad at coming up with rules for this. And so this is why they put it to AI is to try to get the AI to develop a system that can take a sequence and predict what the final structure is going to be.

Robert J. Marks:

Way before the artificial intelligence, I think I learned this from you, Eric, was that there was a game called foldit. Now, again, this is way before this artificial intelligence, but foldit was a demonstration of humans creativity. Could you walk us through that? And then I think that there's another point I want to make.

Eric Holloway:

Yeah. The basic protein folding problem, it's what's known as a MP complete or NP hard problem in that, just looking at the basic structure of the DNA, when you have a string of 10 DNA, you get two to 10 different possible ways that could fold. And when you have computers just trying to go through all their permutations, it takes way too long to do it just by brute forcing it. So they found that complete amateurs, people who have no understanding of biology whatsoever when they saw these folding algorithms doing their thing, they could easily spot optimizations that the algorithms were missing. And so the researchers just turned it into a game and they started making breakthroughs.

Robert J. Marks:

I remember. Yeah, Foldit used to ask you, "Hey, you're not using your computer today at 2:00 AM to 5:00 AM. Let me use your computer."

Eric Holloway:

Yeah, that was the original approach. That's where they got this insight from. First they just were trying to use people's spare CPU's. And they were very successful at that. They had thousands and thousands of people donating loads and loads of CPU time, but still even with all the free CPU time, they made very, very little progress. But part of the software was a screensaver that showed the computer owner what the software was doing with the protein folding. And that's where people started contacting the researchers saying, "Hey, I could do a much better job than this algorithm is doing." And so they just turned it into a game.

Robert J. Marks:

Indeed. This is, I think, a very, very interesting insight into the foldit program. This breakthrough was made by, I believe DeepMind and DeepMind is famous for reinforcement learning, which was used for example, to beat the world champion, Lisa Dell in the game of Go. And it's a game. Protein folding is a game. And the game was proposed to users and the users before the AI were able to solve it much more quickly than the artificial intelligence.

Eric Holloway:

But that being said, I do believe that AlphaFold does even better than the human players does.

Robert J. Marks:

Oh, AlphaFold, yes. I was talking about the old technology and now we have the new AI, which is the AlphaFold.

Eric Holloway:

Yeah. Although that being said, I don't know what exactly it means when it says AlphaFold has solved protein folding. Are they testing this on entirely new DNA sequences? And if so, how do they actually

know it's telling them the truth? I think they're actually still using a known data set as a reference data set. And so they're... Yeah. I don't know if it's actually completely solved protein folding.

Robert J. Marks:

Biologists really are excited. One of them, Andrei Lupas said this is going to change medicine. And this was an article in Nature, a very, very prestigious magazine. Not always right, but very prestigious. It says it will change research. It will change bioengineering. It will change everything.

Eric Holloway:

Yeah. Although there's another thing too, this Andrei Lupas, they try to hold him up as an example of how incredible AlphaFold is. And they say, Oh, he spent a decade trying to figure out the shape of one protein and AlphaFold does it in half an hour. But why was AlphaFold able to do it in half an hour? Because it depends on decades and decades and decades of researchers just like Andrei Lupas trying to figure out the shape of one single protein. And so really Andrei Lupas, his research has accelerated because he has better access to all these other scientists research through AlphaFold.

Robert J. Marks:

Let me ask you, Jonathan, Eric, what is going to be the impact if AI is able to solve the protein folding problem? Where is it going to be used? Will it affect me?

Jonathan Bartlett:

Basically, what it's going to allow you to do is model drugs before we actually test them. So for example, if some drug company has a drug that they want to put out, they're going to be much more able to test its effects against various proteins, because it's going to be able to have a model for them. So it can estimate what it thinks is going to happen and run a lot of those tests in Silicon rather than in life.

Robert J. Marks:

Wow.

Eric Holloway:

Although also there, I would be a bit cautious too, because this again is one of those deep learning models. And I'm not quite sure how they're verifying the results. So you could have a lot of corner cases that deep mind is just totally off on with the folding. So I think it can probably accelerate them by showing them where to look. But I think they probably can't completely replace real experiments, too.

Robert J. Marks:

Fascinating stuff. Well, there you have it. We've worked our way through the top 10 smash AI hits of 2020, and thereabouts with Bradley Center Brain Trust members, Eric Holloway, and Jonathan Bartlett. Eric and Jonathan. Thank you very much. It's been a blast and fun to talk to you. We are wrapping up this top 10 list and we're going to return to our regular podcast next time on Mind Matters News. Until then be of good cheer.

Announcer:

This has been Mind Matters News with your host, Robert J. Marks. Explore more @mindmatters.ai. That's mindmatters.ai. Mind Matters News is directed and edited by Austin Egbert. The opinions