The 2020 AI Smash Hits: Countdown by Bradley Center Brain Trust Members

https://mindmatters.ai/podcast/**ep157**/

Speaker 1:

There's a lot of exciting stuff going on in the world of artificial intelligence, and while it can be hard to pick a favorite, the Bradley Center Brain Trust rounded out 2020 by narrowing it down to their 10 favorite developments in AI. Now here's your host for Mind Matters News, Robert J. Marks.

Robert J. Marks:

Greetings. If you check out other top development lists in AI for 2020 on other services, on other websites, you'll see a lot of hedging. This or that, AI looks, "promising," or they say something like we expect that such and such will happen. This is hype born of what I've referred to as delayed scrutiny. The thing isn't happening now, but you know what? Maybe in a year or two it will. Most of these promises turn out to not come true. They don't come to fruition. Time passes and the fallen promises are forgotten. Then more promises are made. Hedging claims about AI come in part because of many preliminary results are what developers called proof of principle. An idea is shown to work on a simple problem.

Robert J. Marks:

Now, don't get me wrong. The idea of a proof of principle is the important first step of the development of technology, but extrapolating it beyond that achieved can be misleading. It's like a staircase to the moon has been promised, but in fact, only a few steps have been built. An ultimate test is whether a technology has been reduced to practice. Has it made a financial impact on the market? Has it been adopted by the very picky U.S. military? Has it changed lives? In our list, we'll also include cases where there is a clear path to reduction to practice. With this in mind, let's walk through the AI Smash Hits of the last year or so. I say the last year or so because some of these stories are from late or middle 2019.

Robert J. Marks:

To guide us through the countdown, we are joined by two members of the Bradley Center Brain Trust. Jonathan Bartlett is the director of the Blyth Institute. The Blyth Institute focuses on the interplay between mathematics, philosophy, engineering and science. Jonathan is the author of several textbooks and edited volumes which have been used by universities as diverse as Princeton and DeVry. He is a senior fellow with the Bradley Center. Welcome, Jonathan.

Jonathan Bartlett:

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

Robert J. Marks:

Yeah. It's always fun to talk to you and Eric. Dr. Eric Holloway, who likes to be referred to as Captain Eric or no, Captain Dr. Is a Dr. Captain. Eric, which one is it?

Eric Holloway:

Dr. Captain, sounds more prestigious.

Robert J. Marks:

It does. He works for the NIH and is a current captain of the United States Air Force where he has served in both the United States and Afghanistan, hence the idea of Dr. Captain. He's also a senior fellow of the Bradley Center. Welcome, Eric.

Eric Holloway:

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

Robert J. Marks:

Okay. Let's get started with the countdown of AI Smash Hits of the last year or so. Number 10, AI text translations. If you've ever taken a foreign language, you know that you just can't take a word in German and translate it to English, keep them in the same order and it be a proper translation. So something has happened in AI text translation which has really helped us. What's going on here, John?

Jonathan Bartlett:

So, text translation has actually been with us for a while. It's been a popular usage of AI, but what's really happened over the last year or two is the integration of these text translation engines into so many different parts of engineering projects. So for example, one of the things that makes this so successful is that when a lot of people think about AI, they try to think about something that just goes off on its own and you click a button and it does everything for you and that's it.

Jonathan Bartlett:

But what's actually been really successful about text translation is that people have started to recognize both the benefits and the limitations. And that's what's been key for making it very practicable. For example, I've worked with some companies that what they used to do when they wanted to translate an application into a foreign language, is that they'd go and they'd hire a translator and the translator would go and retype everything from one language to another. And that, it take a lot of time. Now, what they can do now is they can automatically translate these applications from one language to another. Now we don't expect it to be 100%.

Jonathan Bartlett:

And so basically what happens is that you write an application, you do an automatic translation tool, and then you have your translators go and simply correct where the AI has gone wrong. And so what's really made AI text translation work so well is the expectation that it's not going to be perfect, but that we're going to have a second pass. And the goal of the AI is not to replace the person, but to make the person more powerful.

Robert J. Marks:

Now, I have been using Google Translate. I talked to a person that used Google Translate a lot and she was very successful with Google Translate. And I said, "What's your means of success?" She says, "Well, I take Google Translate and I transform, for example, English into French. And then I transform the French back into English using Google Translate and see if I get roughly the same thing," because in translation, it's like the old game of telephone, where you whisper a phrase to a friend who whispers it to somebody else and you go down the line and the end it comes out just totally ridiculous of what it was initially set to. So what is all of this stuff done by Microsoft, Apple, in terms of this AI text translation? What have they done special this year?

Eric Holloway:

One thing I saw, not from this year but I think two years ago, Google came out with this really cool text to AI thing that also worked with optical character recognition. So the really cool thing is you could go to a totally foreign country and hold your smartphone up over some sign you can't read, and on the fly would translate the sign straight into English in the picture. So you hold it up to a sign that says bonjourno, and it will turn it into a sign that says hello with the same text format and everything, is really neat.

Robert J. Marks:

So that's nice, but I think we've had that technology with us for quite a while in terms of Google Translate.

Eric Holloway:

Yeah, yeah. It's a couple years old.

Jonathan Bartlett:

So what I've been seeing happening recently, for example, Localize is a company that specializes in translating application interfaces. And actually what they do is they help you manage your own translations. And so what they've done is it used to just be that you would create your app and then you'd send people to translate it. Now what you do is you create your app, you put in the original text and then that app will give you different sources of translation. So they'll say, hey, here's what Google says you should translate it to. Here's what Microsoft says you should translate it to. And look, here's actually what a human has translated this into.

Jonathan Bartlett:

So if somebody else in their system has had this and has had a manual translation, they'll actually tell you, "Hey look, here's a manual translation that someone did of this same very phrase. And so we believe that this is a spot on translation for what you're trying to say." So they mix in both the machine translation because it gives you somewhere to start, but also will give you exact translations. And then you can also go back with an expert later. And so this allows companies to very quickly, if you've got an international audience ... I know I've worked with companies where it takes months to go to the translator and get something back, especially if you're in several different languages, it can just take forever.

Jonathan Bartlett:

But with machine translation, you can actually have a first pass done overnight. You just go through and you just say, okay, we're going to go through and just accept all the automated translations first off. And

then the individual translators they'll get around to checking them when they have time. And so this allows you to iterate quickly, it allows you to be able to create user interfaces quickly. It doesn't slow you down, but it recognizes the limitations of the tool itself and recognizes that the humans are actually at the top of the intelligence barrier rather than at the bottom.

Eric Holloway:

That's a really clever business strategy too, because they get people in the front door by saying, "Hey look, AI's going to translate it for you." And then it turns out it's the people who they're getting in through the front door, who are doing the translation.

Robert J. Marks:

Okay. Very good. Very good. Okay. We're counting down the top 10 AI Smash Hits of the last year or so, 2020, roughly. Number nine, we're going to talk about smart cars a little bit and some advances that have been made with smart cars. So number nine is smart cars. And I was talking to Eric. And Eric might actually want to buy one of these. It's a human in the loop autonomous driving. What is this you want to buy, Eric?

Eric Holloway:

Well, I went to buy it because I'm a really stingy person. I don't like putting out a lot of money for things. And this thing, this is a little aftermarket kit, you can install in your car, that's way, way cheaper than actually buying a smart car.

Robert J. Marks:

Oh, so this is something which you can add onto an old car and make it self-driving or what?

Eric Holloway:

Yes. Yeah. And unfortunately, my stinginess also keeps me from doing this because you have to have a pretty recent car and that's way over how much I'd be willing to spend on a car. But more recent cars, they actually have the rudiments of a self-driving system already built-in. Lots of them now, like you mentioned with your daughter, they'll have some really simple cruise control that's a little bit smarter and can slow things down and speed things up. But what it gives is a bunch of hooks into the cars driving system, along with some sensors.

Eric Holloway:

And so Comma.ai, they've been really smart. Oh, and by the way, the founder of this company, he had his start with hacking Android phones. So he has a good background for hacking into stuff. But the idea here is that you take their system and you plug it into your car's computer and basically hacks the system and takes control of the little driving hooks and then makes them smarter. I've seen some quotes where people say, "Yeah, after I hooked it in, I pretty much only have to drive 1% of the time and the car can do most of the rest of it." But the other smart thing here is that it's not this level five driving.

Eric Holloway:

He's not making any grandiose claims. He's a very realistic self-driving car person. And he's only having the system give suggestions and helping maintain the state and stuff like that. Things that are pretty simple to do from a controls system perspective and never actually takes the human out the loop. And

the system also makes sure that the human stays in the loop. They have attention checks and stuff like that. So that's one aspect of this that I think is really cool.

Eric Holloway:

The other aspect that I think is even cooler is that everything here is open source. So he is open source the whole operating system for his, I think it's called open driver, the operating system that runs all of this. So you could even augment your car even more and make it smarter in many other ways. And as we've seen with the open source movement, we've gotten huge amounts of innovation out of this. So I think this is the real direction to go with self-driving cars.

Robert J. Marks:

This sounds really strange. You plug it into your car's computer, it hacks your computer and takes it over. Is there a lot of danger of doing that?

Eric Holloway:

Yeah. Yeah. And you can also plug it into the internet. So I guess a hacker could take over your car too. So yeah, there's definitely, you don't want to trust this all the way, but the approach itself is pretty neat, I think. So this puts smart driving into the hands of the consumer.

Robert J. Marks:

So let me ask you, do you know offhand what the cost of one of these is?

Eric Holloway:

The unit is 1,000 bucks and then the cabling is an extra 200 bucks.

Robert J. Marks:

So 1,200 bucks you can take your recent car and make it into a pseudo, a self-driving car?

Eric Holloway:

Yeah. And then I could plug it into my brain helmet and then just drive everything with the car of my brain waves.

Robert J. Marks:

You do. By the way, you should know that Eric Holloway does have a brain helmet and he likes to put it on much of brain waves on the PC. Is that right?

Eric Holloway:

Yeah. You can see all kinds of cool stuff. Yeah, it really gets into the whole materialistic question of whether our mind is our brain. And you can see actually how your conscious experiences drive your brain state instead of vice versa.

Robert J. Marks:

That's facet. Well, that's a topic for another time. That would be fun to talk about. We are counting down the top 10 Smash AI Hits of 2020. We're down to number seven. Number seven, which I don't-

Jonathan Bartlett:

Wait. You skipped me.

Eric Holloway:

Yeah. You skipped John.

Robert J. Marks:

Oh, I'm sorry. I'm sorry, John. I didn't mean to skip you. We are counting down the top 10 Smash AI Hits of 2020. We just talked about self-driving cars. We're going to talk about self-driving cars again. We're going to talk about an article that John wrote for Mind Matters News, which says that Daimler, Waymo and GM make big gains in level four self-driving cars. Again, it's been a while since the last podcast. What is level four and why is that not the ultimate self-driving car? And then tell us what some of these breakthroughs are?

Jonathan Bartlett:

Yeah. So, basically when people think about self-driving cars, a lot of people think about cars that don't have steering wheels, you just get in it and you tell the car where you want to go and it takes you there. That's considered level five self-driving. That's where the car can take you anywhere and does not require any intervention whatsoever. But what we mentioned in the last podcast is that level five, there's a lot of roadblocks to level five, which at least in my opinion, we probably won't ever be able to achieve a full level five, but below that is level four. And level four is where you basically find certain parameters for which your self-driving works.

Jonathan Bartlett:

So for example, you might have an area of town that you've really well mapped. And so your car is able to do self-driving in this area or below certain speeds. You feel comfortable that your car has enough sensors on it that it can drive itself. So basically, as long as you are within a certain set of parameters that are well defined, then the car does operate by itself and doesn't require any intervention. What we've seen over the last decade or so, is that the people who have been pursuing level five self-driving are nowhere. They've sometimes come up with little tricks and techniques that are interesting, but they haven't gotten any closer to the actual dream of level five self-driving. But what's really taken off is level four self-driving.

Jonathan Bartlett:

And when you see companies that have aimed at having defined areas for their cars to drive in, this has actually worked out really well. And the first company to do this was Dahmer. They had a self-driving parking garage, and basically they had cars that would go and park themselves within the parking garage. And that's because the garage itself had an intelligent infrastructure that had sensors, and so the cars weren't limited to just what they could see. So the cars could locate a parking space with the help of the infrastructure and could navigate there. And then over the last year, we've had lots of developments with tax services like neighborhood taxi services that have been able to drive in low speed neighborhoods and things of that nature.

Jonathan Bartlett:

GM just got a permit to do driverless cars in San Francisco. And they're supposed to be having driverless cars there by the end of this year. I don't know if they actually succeeded or not, but that was their goal. So level four self-driving has actually really taken off. And it's, you can think of level four as being an engineering project and level five as being a philosophy project.

Robert J. Marks:

That's fascinating. It's a great analogy. Okay.

Jonathan Bartlett:

And that's why the level four projects really work really well, is because they're defined, they have a scope, you know what the difficulties are, you know what's going in your favor, and you can really design around a space. But level five is the idea that we're going to have machines that basically think like humans and can participate socially with them with no problems. And that's more of a philosophy project. And when people confuse their philosophy projects and their engineering projects, we have problems.

Robert J. Marks:

Very interesting. The other thing that level four is, is a much more narrow AI. AI has been successful in applications to narrow problems, where the idea of general artificial intelligence has remained elusive and I think possibly might be unachievable. So this narrows the scope and I guess guarantees a greater degree of success than the more general problem.

Eric Holloway:

Yeah. Although one thing that pro AI people will say for that is that that's goal post moving, because they'll say, "Oh, back in the day, no one thought an AI could win chess, but now wins chess and now, oh, that's no longer an AI problem." So they're saying, "Oh, all this kind of, it's not general. AI is just go post moving. We're incrementally moving towards general AI." Have you heard of that sort of thing before?

Robert J. Marks:

Yes. Absolutely. And that's a fair criticism. I would call that AI of the gaps or algorithm of the gaps. And indeed, it's a fair thing to claim, but there are fundamental limitations in computer science that give a ceiling on what artificial intelligence can do.

Jonathan Bartlett:

Well, one thing that, I think what makes it convincing is that, so basically as computer programmers get better, we invent new axioms that help us to understand the world better. And basically, the way that I frame it is that, that humans are able to invent new axioms and computers basically aren't. And so what happened, well, the reason why AIs are able to do more and more and more and more is because we keep on feeding them more and more axioms. And that if the axiom well ever dried up, the AI improvements would stop, but it's because humans are able to rethink and reform the world in different ways by developing these new axioms to understand life. And that's what allows us to create AIs that can do more and more things, but it's always the humans that are leading on the axiom side.

Robert J. Marks:

And in fact, that goes back to the fundamental premise that AI itself cannot be creative. For generation of these axioms, there has to be creativity.

Jonathan Bartlett:

Exactly.

Robert J. Marks:

Excellent. Yes. Okay. We are down to number seven then. The top 10 Smash AI Hits of 2020 and thereabouts. This one I don't quite understand, but Eric says it's interesting. Hacking AI, exposing vulnerabilities in machine learning using adversarial examples to control AI powered products, that's a lot to unpack. What's going on here, Eric.

Eric Holloway:

Yeah. This actually follows from a fundamental result we talked about a couple sessions ago, that a lot of AI suffers from a problem known as under specification. So with a lot of these real world AI products like deep learning or extreme gradient boosting, or support factor machines, they generate massive, massive models. And even though they have huge amounts of data generating these models, the models themselves still have way, way more parameters than there are data points.

Eric Holloway:

And to use an analogy I used before, let's imagine you had a grid and you had a single dot on that grid and you're supposed to draw the best fit line through that dot. Now there are infinite number of different lines that are actually equally best fits for that dot, but the lines themselves are very, very different from each other. And so we have the same situation with modern AI products because they're such huge parameter models. You don't really know what the AI does outside of its data set.

Eric Holloway:

Now, a lot of the times you've interpolated between data points. So between those points, maybe you can know what's going on, but there's a lot of unknown areas in there. And hackers can prod those unknown areas and nudge the AI models in directions that the hackers want the models to go. And that's just, I think as unescapable symptom of our AI systems, because to make these things work in the real world you have to have these really high parameter models to fit really complex data. But the paradox of the situation is that they become very brittle and much more easier to manipulate.

Robert J. Marks:

Well, in fact, you hear about deep convolutional neural networks trained on images, and all of a sudden you change a pixel or two in an image and the deep convolutional neural network is totally wrong. So they are incredibly brittle. So it's this sort of thing that you're talking about, right?

Eric Holloway:

Yeah. And it's not just the result is completely wrong, but the machine's confidence is result is complete certainty. It's absolutely certain about the wrong result. And in this particular example, they took, I think a self-driving AI and they could just suddenly manipulate traffic signs and make the AI make very disastrous decisions. For example, they gave it a sign that said speed limit 35 and they changed the number three slightly so the AI thought it was 85.

Robert J. Marks:

Fascinating. Okay. We are counting down the top 10 Smash AI Hits of 2020, and we're down to number six. Number six has to do with the military. AI beats fighter pilot. This is the headline in a dogfight. It's clear DARPA gets it right. Eric, you're a captain of the U.S. Air Force, you know about planes and dogfights and things like that. What an AI is going on here?

Eric Holloway:

Yeah. I think this is actually, in one of the few high headlines, it's actually a pretty legit pipe. I think the fighter's pilot job, they like to claim they're irreplaceable and so on, but I think they're actually pretty replaceable by a decent AI. And for nothing else, then the fact that AIs are not humans. You don't have all the constraints of a human body and AI can pull 20Gs and not break a sweat, whereas a fighter pilots is going to pass out at around 9Gs. So just in pure performance, an AI is going to be able to do things that a fighter pilot cannot.

Eric Holloway:

Also, the parameter space is a lot simpler than say what you're dealing with as a soldier on the ground navigating through a dense urban situation or a deep jungle in air. You're just flying through air high above the clouds where you're not having, I guess you have some jet streams and stuff, but it's a pretty simple environment. You're moving in three dimensions, but that's really not too tricky for different various algorithms to figure out.

Eric Holloway:

And you don't even need very fancy AI for this. This is using probably control systems that have been well studied for the past couple decades or so. Well, probably since, even before the 50s. So this is going to be a kind of AI that can rely on really well established control systems and do things that fighter pilots cannot. So, I definitely think this is the way to go for the air force.

Robert J. Marks:

That is going to be very interesting. And you're right about the ego of fighter pilots. I've also found this in other professions. I think that heart surgeons really have incredible egos. And as far as fighter pilots, if you watch Top Gun, you see big egos from Tom Cruise and Val Kilmer. And this isn't far from the truth.

Eric Holloway:

Your ego is writing checks your body can't cash.

Robert J. Marks:

Exactly.

Eric Holloway:

The line from Top Gun.

Robert J. Marks:

Oh, is that a quote from Top Gun? Could you say that again?

Eric Holloway:

Your ego's writing checks your body can't cash.

Robert J. Marks:

Okay.

Eric Holloway:

Yeah. I think this also moves ... I mean, essentially once you have a AI fighter pilot or a fighter jet, you just have a really smart missile. And I think you can also probably scale that down and get really smart bullets and stuff like that. So I think you have a lot of interesting possibilities out of this, narrowly constrained smartness in the weapons.

Robert J. Marks:

Well, the fascinating part is, as you mentioned, Eric, is that if we do come up with these AI fighters, that they are not constrained by human weaknesses. They can stand these high G forces and can do things a lot quicker.

Eric Holloway:

Right. Just in general, I think the military can rely a whole lot more on robots and AI because robots can go many places that humans cannot go. And plus you don't lose human lives. Hopefully in the future all of our words are just going to be robots fighting each other like in Japanese cartoons are driven by a bunch of eight-year-olds and 13-year-olds.

Robert J. Marks:

Okay. Well, I hope not. I hope not. But one of the things about this, I watched the, they did this AI beating the fighter pilot in real time, they streamed it. And I watched part of that. And of course, the AI did beat the fighter pilot, but much of this was meant to be psychological. It was supposed to be a gut punch to the egos of the fighter pilots to make them realize that AI was going to be viable. So I don't think this is something which is going to be reduced to practice tomorrow, but I think the DARPA did get it right in taking these first steps to make sure that the fighter pilots understood the power and the utility of artificial intelligence.

Eric Holloway:

Yeah. Now, one thing I did see, yeah, I watched the whole casting of the fighter pilot fighting the AI, but there was a qualitative difference between the two techniques. So the AI was just really good at optimizing really tactical moves. It could take tighter turns and follow the pilot, but it never innovated. It was pretty much just following and dogging the pilot into a got in for the kill shots. Now the pilot on the other hand, he was always coming up with new innovative moves.

Eric Holloway:

So there is an aspect still in which the AI is still very different in terms of fighting than the pilot. So once the AI got it locked into the very small area that it knew how to get the kill shot, then the pilot was toast, but if the pilot is able to broaden out the domain he's fighting in so there's a lot more options and he can take a longer term strategy, there might still be ways that pilots can defeat AIs.

Robert J. Marks:

This is part of the history of warfare. If you have an enemy AI fighter pilot, they are such that you can probably, after a while, predict the limitations of their movements, then you can gain them to defeat them.

Eric Holloway:

Exactly. Yeah.

Robert J. Marks:

So it gets back to the idea of the creativity of AI. So you need the creative human aspect in there in order to win the day.

Eric Holloway:

Right. And I think the bigger vision for this is not just we replace all our fighter pilots with a bunch of AI pilots, but it's going to be more hybrid approach where you have the fighter pilot, and then he has a bunch of robot wingman that he can control.

Robert J. Marks:

Exactly. 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 and the hall of presidents, most of the 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:

Well, so 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 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 for you have the reface app that takes your face and puts it on some movie stars and clips in special movies. And that's 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.

Robert J. Marks:

Okay.

Jonathan Bartlett:

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:

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

Jonathan Bartlett:

Yeah. 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. And they do make it hard for people to recognize truth from reality, but there's also a lot of practical applications people can do as well. It speeds up animation. You can think of animation, you can as a giant deep fake project.

Jonathan Bartlett:

And so the ability to do real-time defakes is helps people do some filmography and some special things. But another really interesting advance in defakes is for compression. You basically, what some people have figured out is that you can basically deep fake yourself. And it 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 transmitting 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, as kind of 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, 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:

Wait a minute. 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 bitstreams that we have in our files are not just any bitstreams. 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 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.

Jonathan Bartlett:

So each of these levels of expectation allows you to compress your signal to some degree. And so, basically what defakes 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's 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. Anyway, so that's what defakes do, is they take a small amount of data and it 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 in entertainment as a way to capitalize on not having to hire a 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 there's just throw a cartoon suit on them virtually. So there's a lot of possibilities here.

Robert J. Marks:

This is really interesting. I wonder if we'll ever have a deep fake superstar. And I say that-

Eric Holloway:

Oh, yeah. Yeah.

Robert J. Marks:

Yeah. 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 in 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 is actually a music band like that. The music band it consists of four cartoon characters and they've made a couple of hits that have gone 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. 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. You'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'll 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 could see the neuroplasticity adapting so that it controlled 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 even more effective route they can go with this is if the exoskeleton can feed him some 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 neuro network, like 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.

Eric Holloway:

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 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 her 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 a neuroscientist Yuri Danilov, who is one of the founding sciences of that. And also a gentleman named Dr. Sacker 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 developed 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 part of the brain. As you develop, it's pulled directly from the brain. And you have a lot of connections which go directly to the brain. Yeah, that's just fascinating stuff. I got so exuberant that I skipped number four. Didn't I?

Eric Holloway:

Yes. Yeah, you skipped.

Robert J. Marks:

I did skip number four. Now there was no guarantee in this countdown that they would be an 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 a count meandering with the top 10 Smash AI Hits of 2020. This story actually dates back to 2019 and it is a 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 it's a precursor to, here I go again, retinoblastoma, which is a fancy word for eyeball cancer.

Robert J. Marks:

Now, the story behind this number four, Smash Hit of 2020, is very local, very close to me. There's a professor at Baylor University in Chemistry named Brian Shaw. He had a 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 prohibited, but at least have an early detection. So he contacted a professor in the department of computer science of Baylor University, Doug Hamerly 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.

Robert J. Marks:

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 can 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 the pictures of your kids on your cell phone have this white eye. This indicator of eyeball cancer. 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.

Robert J. Marks:

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 beat Darren Elias.

Robert J. Marks:

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. 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 a poker is very highly algorithmic. There's some randomness that goes in here for sure, but the fact that artificial intelligence can win a 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,

it's a game theory 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 all of these results where AI beat humans, because you're only ever told the end result that the AI beat the human and you're given some 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. What would happen then?

Eric Holloway:

Yeah. Well, I think that's actually part of it's training. It plays itself a number of times too. It's like the AlphaGo Zero that plays itself a whole bunch in order to develop it strategies. 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 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 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've 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, when you have DNA, and DNA is basically this long strand of what are called basis, which are basically the letters of DNA.

Jonathan Bartlett:

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. 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.

Jonathan Bartlett:

And so the question is, is which way will it actually fold? And being able to assess that out has been a difficult problem and the biologist always want to just be able to see the sequence and infer what the final shape is going to be. So 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 is what's known as NP-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 a two to 10 different possible ways that could fold. And when you have computers just trying to go through all the 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 CPUs 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:

So, indeed this is I think a very, very interesting insight into the folder 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 dataset. And so there, yeah, I don't know if it's actually completely solved protein folding.

Robert J. Marks:

Well, 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 bio-engineering, it will change everything.

Eric Holloway:

Yeah. Although there's another thing too. This Andrei Lupas, they try to hold them up as an example of how incredible AlphaFold is. And they say, oh, you 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? It's 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 DeepMind 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. So until next time, be of good cheer.

Speaker 5:

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.