

AI Smash Hits 2020 Part I (<https://mindmatters.ai/podcast/ep116>)

Robert J. Marks:

We count down the top 2020 AI smash hits 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, 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.

Robert J. Marks:

Hedging claims about AI come in part because of many preliminary results or what developers called, proof of principle, an idea is shown to work on a simple problem. 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 beyond that achieved can be misleading. It's kind of 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 US military? Has it changed lives?

Robert J. Marks:

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. 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:

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 Doctor, is it 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 is 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. 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.

Jonathan Bartlett:

For example, I've worked with some companies that what they used to do when they want 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 re-type everything from one language to another. And that'd 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 a hundred percent, 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 tech 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 at the end 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, but 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 it would translate the sign straight into English in the picture. So like if you pulled it up to the sign that says buongiorno, and it will turn it into a sign that says hello with the same text format and everything, it's 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 their app will give you different sources of translation. So let's 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." 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."

Jonathan Bartlett:

So they kind of mix in both the machine translation, because it gives you somewhere to start, but also we'll 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, and especially if you're in several different languages, it can just take forever. 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 is 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 on to 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 kind of 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. 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 it basically hacks the system and takes control of the little driving hooks and then makes them smarter.

Eric Holloway:

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. 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 control system perspective and never actually takes the human out of the loop. And the system also makes sure that the human stays in the loop, they have attention checks and stuff like that.

Eric Holloway:

So that's one aspect of this that I think is really cool. The other aspect that I think is even cooler is that everything here is open source. So he's open sourced 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 it?

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 off hand what the cost of one of these is?

Eric Holloway:

The unit is a thousand bucks and then the cabling is an extra 200 bucks.

Robert J. Marks:

So 1200 bucks you can take your recent car and make it into kind of a pseudo self driving car?

Eric Holloway:

Yeah, and then I could plug it into my brain helmet and then just drive everything with the car with my brainwaves.

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 and watch his brainwaves 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:

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 number eight.

Eric Holloway:

Yeah, you skipped Jon.

Robert J. Marks:

Oh, I'm sorry. I'm sorry, Jon. I 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 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 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.

Jonathan Bartlett:

But below that is level four, and level four is where you basically define certain parameters for which your self-driving works. 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, it doesn't require any intervention.

Jonathan Bartlett:

What we've seen over the last decade or so is that the people who have been pursuing level five self-driving are kind of nowhere. They've 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. 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 Daimler, 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 find locate a parking space with the help of the infrastructure and navigate there.

Jonathan Bartlett:

And then over the last year we've had lots of developments with taxi services, neighborhood taxi services that have been able to drive in low speed neighborhoods and things of that nature. 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 kind of, you can think of level four as being an engineering project and level five is being a philosophy project.

Robert J. Marks:

That's fascinating. That'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 goalposts moving, because they'll say, "Oh, back in the day no one thought an AI could win chess. But now it wins chess, and now, oh, that's no longer an AI problem." So they are saying, "Oh, all this kind of... It's not general AI, it's just goalposts 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 on 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 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:

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, 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 kind of 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 vector 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. 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.

Eric Holloway:

And so we have the same kind of situation with modern AI products, because there's such huge parameter models you don't really know what the AI does outside of its dataset. Now, a lot of the times you've kind of 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, is an inescapable 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 the 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 in its 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 subtly manipulate traffic signs and make the AI make very disastrous decisions. Like 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 meets fighter pilot. This is the headlines, in a dog fight it's clear DARPA gets it right. Eric, you're a captain of the US Air Force, you know about planes and dog dogfights and things like that, what in AI is going on here?

Eric Holloway:

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

Eric Holloway:

Also the parameter space is a lot simpler than say what you're dealing with as a soldier on the ground navigating 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. And you don't even need a very fancy AI for this, this is using probably control systems that have been well studied for the past couple of decades or so, we'll 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 can not. 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.

Jonathan Bartlett:

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

Robert J. Marks:

Exactly.

Jonathan Bartlett:

That's a line from Top Gun.

Robert J. Marks:

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

Jonathan Bartlett:

Your ego is 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 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 kind of 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 sort of 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 wars are just going to be giant robots fighting each other, like in Japanese cartoons driven by 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 that 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 pilots, but it never innovated. It was pretty much just following and dogging the pilot until it got in for the kill shot. Now the pilot on the other hand, he was always coming up with new innovative moves. So there is an aspect still in which the AI is still very different in terms of fighting than the pilot. And so once the AI got it locked into the very small area that it knew how to get the kill shot in, 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'll replace all our fighter pilots with a bunch of AI pilots, but it's going to be more a 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. Well, we've been working our way through the top 10 AI smash hits of 2020 with Bradley Center Brain Trust members, Eric Holloway, and Jonathan Bartlett. We're up to number six, the top five is for next time and so until then be of good cheer.

Announcer:

This has been Mind Matters News with your host Robert J. Marks. Explore more at mindmatters.ai, that's mindmatters.ai. Mind Matters News is directed and edited by Austin Egbert. The opinions expressed on this program are solely those of the speakers. Mind Matters News is produced and copyrighted by the Walter Bradley Center for Natural and Artificial Intelligence at Discovery Institute.