Artificial General Intelligence: The Modern Homunculus

https://mindmatters.ai/podcast/**ep159**/

Robert Marks:

Greetings. This is Mind Matter News, and I am your full-figured host, Robert J. Marks. Isaac Newton was the genius that founded classical physics, he also invented calculus. He also did other things. Newton was a student of the Bible, specifically Bible prophecy, and he wrote extensively on his research. Newton also dabbled in alchemy. Now, most think of alchemy as the quest to turn lead into gold, but there's a lot more to alchemy than doing this. Some in alchemy pursued creation of a so-called homunculus. Homunculus is a little person created in a test tube. If you watched the 1935 classic monster movie, The Bride of Frankenstein, you see a scene where the mad scientist, Dr. Pretorius, shows off his homunculi, I think that's the plural of homunculus. He shows off his homunculi to Henry Frankenstein.

Robert Marks:

No one, to date, has created the alchemist's dream of a homunculus. And if you exclude, maybe, cloning, I don't think that they probably ever will. The search for the homunculus today has been replaced by a search for artificial general intelligence or AGI, artificial general intelligence. Terms keep changing in a rapidly evolving field. AGI used to be called hard artificial intelligence. There are some that actually have a split definition, but we're just going to stick with artificial general intelligence. By any name, the search for artificial general intelligence will prove as useless as the search for the homunculus. Not everybody shares this opinion, that includes Elon Musk, that includes Stephen Hawking, but even so, there is a growing evidence AGI will never be achieved.

Robert Marks:

What does AGI do? AGI seeks to duplicate and exceed what you and I do. If artificial general intelligence is achieved, some say we will become pets of computers. There are some who worry that AI will begin to write better and better AI. The point where AI becomes superior to humans is called the singularity, by Google's Ray Kurzweil. If this happens, watch out, AI will write better software that writes better software that writes better software and an endless staircase of ever increasing intelligence. And there are smart people who believe this will happen, but AGI is not happening, and there's growing evidence it never will. AI can be written to mimic many human traits, but there are some human characteristics that will never be duplicated by AI. We covered this a lot on Mind Matters News.

Robert Marks:

Properly defined, these properties that will never be achieved include creativity, sentience, and understanding. In fact, AI seems to be going in the opposite way, more and more human expertise is being folded into the AI software. The added intelligence in AI is not due to AI, but is due to human creativity and ingenuity infused in the software, by the programmer. To talk about these things, our guest today is Dr. Justin Bui. Justin is a freshly-mind PhD from my research group at Baylor University, and he specializes in, among other things, artificial intelligence and deep learning. Justin, welcome.

Justin Bui:

Yeah. Thank you very much for having me on the show, Bob.

Robert Marks:

Great. Before we go into some trends in artificial intelligence, what I'd like to do is describe the playing field that you have been watching. And I think anybody that is involved in the development of artificial intelligence is familiar with these resources. There is, of course, the literature, and there's some fast literature in AI. But in computer software, there's a heck of a lot more incredible resources. There are incredible resources available on the web, many related to AI, and the incredible part is most of these are free. So first, let's talk about the software, Justin. AI software is widely available, it's free, and it's powerful. It's available to anyone on the net. Could you go through some of the AI software and some of the things that this AI software that's available for free does?

Justin Bui:

Yeah, sure. It's an interesting playing space. It seems like every day, there's a new tool that comes out that makes everybody's lives just a little bit easier. The big ones, of course, are our PyTorch and TensorFlow, both being driven by Facebook and Google respectively, and they make up... I believe, it's [inaudible 00:04:50] of probably 75% to 80% of a lot of the machine learning systems out there, if not more. They're very easy to use. Its tremendous development and the amount of available resources associated with these tools is phenomenal. And going hand in hand with that is the use of free web resources. A lot of systems out there provide free computational resources, basically virtual machines that anybody can sign up for and use. They can design, deploy, evaluate any machine learning model that they would like. And it's actually quite interesting to see how prevalent some of these systems have become.

Robert Marks:

Yeah. This is really interesting. One of the fascinating things is the free available computation AI like deep neural networks, for example, can take a long time to train. And so you're crunching the computer again and again and again and yet, there is available fast software resources, available on the web, to allow you to do this in cloud. And that, to me, is just amazing, that people are making this available for free. There's also something called fast AI. What is fast AI?

Justin Bui:

So fast AI is a wrapper for PyTorch with a lot of pre-built models. So it's meant for rapid proof of concept testing, if you will, it takes advantage of a lot of transfer learning techniques where just a normal, we'll call them everyday person... But really, anybody can pick up a Jupyter Notebook or a little bit of Python code and follow along on one of their tutorials and effectively deploy a classification model or regression model. It's really meant to help speed up the initial proof of concept for a lot of these model development processes.

Robert Marks:

It's an interface in a way, is that right?

Justin Bui:

Yeah. I think a good way to classify it would be like a high-level wrapper almost, but it... Again, it lets you take advantage of some of the work that's already been done and that ultimately cuts down on somebody's development cycle.

Robert Marks:

So PyTorch, the Py is for Python. Python is a computer language that's available for free, everybody can use it, right?

Justin Bui:

Correct. Yeah, it's a... Torch itself is actually built on Lua which, for those that are perhaps more intimately familiar, is a scripting type language, and so PyTorch is... You're right though, the Python high-level wrapper for the Lua interface, that is Torch.

Robert Marks:

Okay. What sort of stuff can you do with all this free software? Maybe, specifically, some of the stuff we see in the news today?

Justin Bui:

Oh yeah, sure. I mean, all of these tools have high-level code wrappers for doing custom layer developments. So of course, you've got convolutional layers which, for those that are familiar, go into convolutional neural networks. You have transformer layers which are gaining popularity and...

Robert Marks:

Can I interrupt you just for a second? What does a wrapper for the general audience serve?

Justin Bui:

Good question. So a wrapper is just like a high-level function call, it's a chunk of code that ultimately makes deploying something more complex very easy. You can think of it as like a super function in a way.

Robert Marks:

I see. So you might have software and you go build the pyramids, and you click yes and the pyramids are built, something very big happens.

Justin Bui:

Exactly. It'd be something like build a pyramid, and all the hard stuff is done underneath the hood, so to speak.

Robert Marks:

Okay. So you were talking about some of the stuff that you can do with all of this free available software and all of this free available computational space.

Justin Bui:

Yeah. So yeah. Of course, like I mentioned previously, you have convolutional layers, you have recurrent layers, which are things like LTSMs, now that add a little bit of memory, so to speak, to the neural network transformers, as I mentioned previously as well, and a whole bunch of combinations in between, and it's with... These high-level layer calls, if you will, it really lets you get creative with the architecture. You can combine different techniques into this... I guess, you can consider it an [amalgamy 00:09:14] of different neuron types with different inputs and outputs. And you can create this hydra-

looking system, if you will, where it can take various inputs and create various outputs. And it's really great because it lends itself to this creative model development through its flexibility.

Justin Bui:

And both of these tools allow that to happen, and you see this competition back and forth. It's been interesting to follow along as these tools develop. When I started doing a lot of my research, TensorFlow 2.0 was still relatively new. I believe it was still in beta actually. And most recently, I believe they're up to stable release 2.6. PyTorch, I believe, in a similar fashion was at about 1.2 at the time when I started my research, and most recently, their stable release is 1.9. So you're seeing some pretty heavy iteration improvements in these tools, and it's great because it's driving a lot of the AI machine learning development, going hand in hand with deployment of these tools as you're seeing more and more of these free resources that you've mentioned before becoming available. And I view them as a mix of things. One, it gets the tools in the hands of people to experiment and to learn.

Robert Marks:

Now, the interesting thing, this is available to anybody in the world. All of our adversaries in the United States, at least politically, militarily, like China, Iran can plug in, get this free software, and do all this artificial intelligence and do it all for free.

Justin Bui:

Yeah, that's right. It's a double-edged sword in a way. But I think in an ideal world, anyways, what you're doing is you're providing the masses with the tools and the opportunities to push the envelope forward. And I think it's a good thing because it makes the accessibility and the learnability of the techniques much more grounded. Whereas before, it was pretty heavily academic and very computationally intense or required a lot of subject matter expertise. A lot more of the innovation now is who can get to the finish line first [inaudible 00:11:30], so it should, in a way, encourage some more competition.

Justin Bui:

There is one caveat, of course, to free resources, is that they are constrained. Most systems, you're typically limited to a fixed number of training or running hours, you get a fixed amount of memory which, if you think about... I think most systems provide between two and 16 gigs of RAM to use, which sounds like quite a bit. I mean, most people probably have 16 or 32 gigs of RAM on their personal computers. But if you're loading a data set that contains, say, 150 gigabytes of DICOM data, for example, or other medical data, well, that's not going to fit in memory, and you'll find out very quickly that these systems break.

Robert Marks:

However, if you do have the resources of computation and memory by yourself, you can download the software and run it in your system with, basically, limitations which are totally dictated by your resources that you have locally, right?

Justin Bui:

Yeah, that's correct. That's one of the nice things about the open source tools, is that if you want to build yourself a small supercluster and with a couple of terabytes of RAM and whole bunch of processors, you're, of course, welcome to do that, and you have almost no limitations other than making sure that you have compatible drivers and that all of your system software plays nicely.

Robert Marks:

I was thinking of some of the specific things in the news that could be done with this software, one of them is deepfakes. Can you do deepfake images and videos with this software?

Justin Bui:

You can, yeah. In fact, there's a system out there called Kaggle which, I believe, we'll probably mention again in...

Robert Marks:

Oh yeah, let's talk about, what is Kaggle?

Justin Bui:

So Kaggle, it's actually owned by Google, I believe they were acquired somewhat recently, and it's a open source platform that provides computational resources to data scientists and machine learning engineers. But of course, anyone has access to it, if you have an email, you can get access to it. But it's a website that allows people to post competitions, so there are a lot of design competitions of various types. So there's image classification, stock price prediction, housing price prediction, a couple of different things that just highlight the industry in general. And they have modernized the open source resource sharing, community-driven AI push, if you will. And it's been very interesting, I've spent some time perusing on their forums and reading through the discussions and taking a look at some of their competitions, and people get very creative. It's fun to watch as a spectator and see how people approach problems and what techniques do they try if things aren't successful. Do people share that experience, or do they brush it under the rug and move on?

Robert Marks:

I remember from, gosh, a long time ago that Netflix put out this competition to come up with software that, when given user data such as data from you or me, could figure out the sort of things we would like to watch, in other words, things to suggest for us to watch. And they offer big cash prizes for that. Kaggle does this also, doesn't it?

Justin Bui:

Yeah. So it's interesting because they're more of a host system where anybody can throw a competition up there. One example is the NFL has a helmet detection competition going on right now. It's in cooperation with Amazon Web Services.

Robert Marks:

Okay wait, NFL helmet detection?

Justin Bui:

Yes. Yeah. So what they're trying to do is develop a system that can detect and track helmet locations for players, and what they're really getting at is being able to detect illegal hits like targeting, for example, by tracking helmets and detecting when there's helmet-to-helmet collision. So part of it's player safety, but they're looking at ways to automate this because if you think about a human in the system, a referee has to watch how much of the field, well, really all of it, so they miss some things from time to time. And when you think about it from a player safety perspective, you want to be minimizing

or, ideally, completely eliminating some of those rough shots. And the thought is that if you developed an AI system to be able to do that, you could shift that burden, so to speak.

Robert Marks:

I see. So this is an ongoing competition right now. And do they supply data to train the neural network or the AI with?

Justin Bui:

Yes, they do. AWS has provided several gigs of video files, of image files, they've even provided some example code from previous competitions. And the prize is $100,000 dollars split across... I believe, it's the top eight. But it's a pretty large prize purse, and I think if you follow the competition's history, because this is several years in the running now, they wanted a system that could do everything. But they pretty soon realized that getting a system that could do everything was pretty challenging, so they said, "Okay, let's dumb down the problem, let's start with helmet recognition and helmet tracking." If you can start with that, eventually, you could build up to a system that could detect helmet-to-helmet collisions or stuff like that, and so it hearkens quite nicely to the AGI competition. So I think the thought was that the system would create this master referee that could watch every player on the field, track locations, detect illegal hits, et cetera, but people are realizing, "Well, it turns out that's a lot harder than we thought."

Robert Marks:

I can also see this being used by people such as neuroscientists to study the impact of these collisions on brain development. We had a guest in the podcast a while back. I mean, Yuri Danilov, who was a neuroscientist, did just fascinating work and he said... His indication was that all football games were just terrible, and he refused to let his kids play football, until his oldest son finally did get on a team. And I said, "Well, what happened? I thought you forbade it." He said, "I was outvoted," so his kid literally played football. But I could see tracking this in real time would be really interesting because you could measure, for example, the acceleration of the helmet, you could do the... Let me get a little nerdy here. I think in beginning physics, everybody talks about distance, velocity, acceleration. And then I learned, when I was working for Boeing, that each one of those is related by a higher derivative in calculus.

Robert Marks:

So you start with the distance, you get the velocity, you get the acceleration. And then what is the derivative of acceleration? It's something called jerk. And if your acceleration changes really quickly, you have a jerk associated with you, and I could see being used with AGI in order to monitor jerk, which I think that neuroscientists would find very interesting in terms of tracking potential brain damage. And then the cool part is the derivative of jerk is snap, the derivative of snap is crackle, and the derivative of the other one is pop. That sounds really, really strange, but they could also monitor snap, crackle and pop. But the price for this is $100,000, that's not minimal. Who is involved in this, is it universities, is it companies, is it both?

Justin Bui:

That's the one thing about platforms like Kaggle, is it really is anybody. Anybody who wants to participate can join, so I think from my observation, it's a lot of individuals. You can actually join teams and coordinate across the world, really. If you'd like, there's several teams that are multinational. But it's really, anybody's open to it, and I think the larger thing to take away from that is it's crowdsourcing the

development, so to speak. So you can, in a way, fork up what sounds like a pretty significant amount of money, but in the grand scheme of things, from a company perspective, is relatively small and get, basically, unrestricted access to the IP that's developed basically for cheap.

Robert Marks:

Wow, that is really interesting. These are companies which are, if you will, outsourcing their R and D to competitions and probably getting results a lot cheaper than hiring a bunch of experts and trying to tackle the problem locally.

Justin Bui:

Yeah, exactly.

Robert Marks:

Wow. So that works very, very well. One of the things you mentioned to me, Justin, which I appreciate... By the way, Kaggle is K-A-G-G-L-E, so it's kaggle.com for anybody that wants to take a look at it. You mentioned to me that in monitoring these things on Kaggle, that you saw not an advancement of AGI but, in a way, a reversal of the AGI. Could you repeat what you told me about that?

Justin Bui:

Yeah, sure. I think to summarize it, what we're seeing is, like you said, it's a 180. You're really seeing almost this hyperspecificity in a lot of the applications. If you go through and you observe a lot of the competitions that have closed where many of the competitors have shared their code, you see a lot of evidence of transfer learning so, of course, there's some network reuse and stuff.

Robert Marks:

Wait, just elaborate, just a second on transfer learning.

Justin Bui:

Oh yeah, sure. So with transfer learning, you take an existing system, an existing neural network, and you basically discard some of the weights and biases. So you take some of the trained network, and you let it forget some of the information, and then you apply it to a new dataset. It's really common in object detection and image classification networks where, with these very deep neural networks, you have, say, maybe the bottom four or five layers. The ones closest to the output, they wipe their memory, so to speak, and train it on the new data. And so you have the core detection layers up top, which have been trained and tested and verified, being reused but on a new set of data. It's a pretty common technique, and it produces some really great results. And that's one of the things that you see a lot with some of the Kaggle competitions, is a VGG or a resonant being used for top layers, then maybe a little bit of customization on the bottom side.

Robert Marks:

Here's the way I understand transfer learning. Suppose that you had a neural network that was trained on dogs, that you trained this neural network to detect dogs, and you would have to spend a heck of a lot of time figuring out this neural network and training this neural network to recognize dogs. Now, you want to come along and you want to classify cats. Well, it turns out that classifying cats is similar to classifying dogs, so why would you have to go back and start again with scratch, why couldn't you use

part of that dog neural network to train the cat neural network? And the art of doing that is referred to, I believe, as transfer learning. Is that fair?

Justin Bui:

Yeah, that's a great example. It's one of those things that a lot of people are like, "Hey, why reinvent the wheel when I have a system that gives me 85% of a wheel? So yeah, you're spot on.

Robert Marks:

Okay, good. Despite all of these challenges with AGI and your observation that it's going the other way, maybe we're waiting for a new theoretical breakthrough, which I don't think will ever be achieved. But nevertheless, there are people that believe that we are making steps towards AGI. And there are those that believe that, indeed, this is going to happen. Now, George Gilder, who is one of the co-founders of Discovery Institute and just a genius in terms of economic and business commentary and forecasting, says that this dream that these software engineers have is something which could be called rapture of the nerds. I like that because it takes a lot of faith to believe that we're going to get there. One of these companies, which is just overtly into promoting this, is OpenAI, that's the company that brought us GPT-3, and they claim they are pursuing AGI. And I looked at their mission statement, and it included the following, this is a quote, "We will attempt to directly build safe and beneficial AGI, but we'll also consider our mission fulfilled if our work aids others to achieve this outcome."

Robert Marks:

So they definitely believe in this, they have faith that computers will eventually develop AGI. I always thought that was very interesting. Now, you and I had talked a little bit about why these software engineers believe that AGI is achievable. I mean, these are guys which are really, really intelligent, and they believe that the AGI is indeed achievable, and one of the reasons, I think, is because in terms of AGI operations such as understanding and creativity and sentience, that they don't understand it's not algorithmic, they haven't gotten to the computer science. And the word that I used for this was a so-called keyboard engineer. These are people that when they're looking for a solution don't sit down and look at the theory, rather, they go directly to the keyboard. You had some interesting comments on that. Could you elaborate on that?

Justin Bui:

Yeah, sure. It's one of those things that some of my colleagues and I have jokingly referred to as Stack Overflow engineers. It's a very similar concept, but it's a...

Robert Marks:

Okay. Stack Overflow, that's a website, right?

Justin Bui:

Correct, yeah. It's a forum where people can post errors or issues that they're having with their code, and it's a community-sourced solution house, if you will. But it's pretty funny because some of the colleagues I've had throughout the years have joked about, "Okay. Hey, we just got this problem, let me go check Stack Overflow really quick. Chances are somebody's done it before. I'll just reuse it." And so I think that feeds into some of the AGI belief as well. "Oh well, open AI has produced X, Y, Z neural networks," and, "Oh hey, Google and Google's Brain team have published on A, B, C works, and if we can start merging these together, the system will just become super intelligent," and so I think in some

regards, it's fed a lot by what people are observing from the major companies and some of the major influencers like... You had mentioned Elon Musk before, you have some very popular people that are promoting, dare I say, preaching some of these ideas and beliefs, and people latch onto that.

Justin Bui:

It was funny because when I think of AGI, I think of Hal 9000 or Skynet or, for those of you that are more into movies more recent, Ultron, these systems that seemingly have limitless resources and infinite knowledge and, obviously, evil intentions. I think that's one of the things that helps capture people's attention and their creativity as well. But I think at the end of the day, Bob, like you said, people just... They go straight to the keyboard, they don't sit down thinking about how to approach a problem. How do we solve it from the theory perspective and then start deploying it? It's really more, "Well, okay. I need to go make a classifier that tells me the difference between kumquats and giraffes," and they just sit down and start coding.

Robert Marks:

And so they import these things and download the software and use the software as a black box, without looking at the deeper theory of how it is created and the computer science of where it came from and the possibilities of doing AGI in the future. They don't address some of the things we talk about on Mind Matters News, they don't address the Lovelace test for creativity, which has never been demonstrated in artificial intelligence. They don't talk about even simple counter-arguments, like Searle's Chinese Room, about understanding, and as a result of this, there... I don't know. We're guessing here, aren't we?

Justin Bui:

Yeah, in a way.

Robert Marks:

Yeah. We're guessing, but it seems to me that they possibly don't understand or at least, maybe, they've just blocked out this idea that AGI can't be achieved. Really great points. Okay, any final comments?

Justin Bui:

Well, yeah. Actually, I did want to build a little bit on that, too. I think in some regards, AI and machine learning, they've become catchphrases throughout the world where I used to joke that AI is very similar to the word synergies in the business world. Synergies, everybody wants synergies. The new thing is everybody wants machine learning. They don't necessarily understand what it is, it's... Like you had said, it's a black box, let's wave our hands over it, let's see some results. Are they the results we want to see? Great, we now have machine learning. And it's not that easy, but I think a lot of the drive and reason why a lot of keyboard engineers have a lot of success in their careers and gain a lot of influence is they're able to produce those results, which businesses see and they like, and it feeds into this system of like, "Okay, Hey, this person's achieve these results, this company is using machine learning. Well, maybe our company can use machine learning, let's go do the same thing."

Justin Bui:

And the focus, all of a sudden, becomes the material goal, a little bit less about, "Let's create a perfect system," more about, "Hey, let's create the system that provides us the best benefit." Yeah, I think that that's one of the things that really feeds into the keyboard engineer mentality as well, is sometimes, you

don't get the freedom to sit down at a whiteboard and say, "Okay. Hey, how do I approach this problem from a theoretical standpoint, from a high-level concept standpoint before I start writing code?" Typically, it's, "I was just handed an assignment, I've got two weeks to do it. I'm going to go to my keyboard and start writing some code."

Robert Marks:

Okay. That's fascinating. I think that there's a little bit of walk back on the idea of AGI. I did a interview with George Gilder who is neighbors with Ray Kurzweil and good friends. Ray Kurzweil, of course, is the one that introduced the idea of the singularity and was a big proponent of AGI. And Gilder says he's noticed recently that Kurzweil has begun to backtrack a little bit on AGI and its implementation, and the problems... I don't know, there's a bunch of problems associated with AGI. Number one is in the arguments about it, people are using seductive semantics. They say AI will be creative or have understanding or be sentient without really defining what these mean. They're using seductive semantics, and in order to discuss those, you have to be careful in defining them. So yeah, we'll see what happens.

Robert Marks:

And I have little faith that AGI will ever be achieved. And what I was going to say is that AGI has to be defined, and the way that we're defining is the ability to duplicate what human beings do. There will be a lot of stuff AGI or artificial intelligence will do which is a lot better and a lot more impressive than humans would be. Heck, that happened when they came out with the calculator. Calculator does a lot that I am unable to do and does it much more quickly. Okay. Thank you, Justin. Those were really good points. Our guest today has been Dr. Justin Bui, and I thank him for his valuable insights, and I've learned a lot from him. So until next time on Mind Matters News, 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.