AI and You

Transcript

Guest: Daniel DeMillard

Episode 68

First Aired: Monday, October 4, 2021

Welcome to episode 68! And thank you for sharing about the show, our audience has picked up, and we're at record listenership right now, over 30,000 all-time downloads. It's really showing, so thanks and keep it up.

One thing that I think about AI a lot is that it is so versatile, that you can have an AI system that can be equally suited to performing as a defensive weapons system and performing medical diagnoses, and I didn't just make that up, I saw a system that gave exactly those two cases as applications. That seems to extend to the people in AI, that some of them range over really diverse fields and, what is it tying them together? AI, of course, but maybe something else, like curiosity. Well, and curiosity is certainly the gasoline that fuels the engine of this show. And you'll hear that a lot with today's guest: Daniel DeMillard, who has been with IBM's Watson division and is now CTO of Foodspace - and we'll find out what that is - but he described himself as wanting "to enable a future of work that focuses on more creative, critical thinking, and interpersonal skills. He uses his background in math and spirit of learning to explore the depths and possibilities of computer vision, natural language processing, and machine learning." And we've got a lot to talk about in those areas before we even get to his current position in this interview, which really opened up new thinking for me around the whole concept of search and what AI can do with it. Here we go with the interview with Daniel DeMillard.

Daniel, welcome to the show.

Thanks, Peter. It's great to be here.

Terrific. Now, I see that as part of your self-description, you say you're "passionate about leveraging artificial intelligence to automate away the boring parts of life." What are the boring parts of life?

Yeah, I think the boring parts of life are anytime you're doing something very repetitive and mundane, and don't require any creative thinking or deviation from that routine. I think that AI has the capability of freeing us up from those very mundane and repetitive tasks so that we can do the more creative work; and also reducing the cost of those things so that human beings can be more augmented and be more efficient to create, to really shine where human beings do, which is in that critical thinking capacity, as opposed to the robotic, you know, perform the same task over and over again. Let robots do what they do well, and let humans do what they do well.

Sure. Now, you've spent some time at Watson. What did you discover and learn about artificial intelligence there?

Yeah. So my time at Watson, I think was probably at the height of the hype of Watson. Back then, IBM, mostly the marketing department, was billing the Watson system as kind of, I think, a sentient being and it was almost being billed as a general AI. It certainly wasn't. It was one supercomputer that had been very specifically trained to play Jeopardy. And then also a collection of machine learning API tools that could do things like image recognition, natural language processing, like extracting entities, and OCR transcription, speech recognition. So definitely overhyped. It's sometimes difficult because people now have lost faith in AI to some extent and see a lot of the industry as charlatans. It certainly isn't. But the message that I would say is that we're more focused, I think. And I think IBM has made a transition to tone down some of that language as well in the last year, to really say, "It might not be general AI that can solve any problem yet, but we can solve certain problems, certain narrow focused problems really well now."

Well, they were very good at making headlines. Going back to Deep Blue, that produced a measurable jump in IBM's stock price when it beat Garry Kasparov, and so here they are pushing the boundaries, it gets cited all the time. And then they do the same thing with Jeopardy, which had looked equally impossible. And yet, I don't think they monetized Deep Blue - not a lot of market for a chess-playing robot. And it seems that despite their labeling the Watson product as their cognitive computing division, that you're saying they weren't able to generalize that to the extent that we would have thought that something that could win Jeopardy could do. Is that right?

Yeah. So the nuance of Jeopardy is I think it was trained on at least 24,000 questions. And those questions are in a very specific format, and they have a lot of repetition and structure to them across the questions. So although it was a huge leap forward in the ability for machines to understand language, and applying deep learning and neural networks to this domain, it's still a very closed-form problem. And just because you can answer Jeopardy questions doesn't necessarily mean that you can create Twitter chatbots that can talk to everybody or that you can automate the drive-thru process at a fast-food restaurant or that you can provide a companion for people who might want somebody to talk to. So that conversational AI piece is still very much an unsolved problem, even though in specific domains, you can train systems that do well like in Jeopardy, or virtual assistants, as we've seen with Amazon Alexa and Echo or Google Home, those work fairly well. But it's a very specific set of keywords that it performs well with.

Was it, and I want to stay on this for a minute because this is really interesting, was it that they were too far ahead of their time? I mean, here they are spending incredible amounts of money in both cases on Deep Blue and Watson and making these headlines. And yet, they *didn't* come up with AlphaGo. They *didn't* get into deep learning, which has really opened up the space of game playing in one go and showed how it could learn how to do any game. They didn't generalize it in that way. Deep Blue was optimized for one thing - winning at chess. In fact, at the time, it was optimized specifically for beating Garry Kasparov at chess. And now, here's Watson that appears to have been optimized specifically for winning Jeopardy. If they were too early, to generalize that, well, here we are 10 years later. And so the possibility is that now the generalization that could win Jeopardy if it wanted to - no one's going to get to try, so that question is open - is available. Have we turned a corner in AI that makes that kind of generalization of the Watson capability possible?

Yeah, I would say certainly not. And I think the next five to 10 years of AI is still what I would call boiling the ocean. So current monetization of AI boils down to accuracy of your problem. If you don't really care about an extremely high level of accuracy, take recommender systems, for example, you might be okay with a little bit of error, because if you happen to recommend a movie to somebody that they don't want to see, you're making hundreds of recommendations all the time, it's not that big of a deal. If you fail to recommend a movie that they really want to see, it's fine, you've recommended 15 other ones that they do want to see, so the accuracy isn't the biggest issue. The other thing that you can do is if your ML system isn't super accurate, you can get humans in the loop. You can look at fraud detection algorithms. You alert somebody if there's cybersecurity threat, or your credit card is now being used in the Cayman Islands or something, and will alert you for fraud. You might get a false alarm every once in a while, where you're actually using the credit card for something legitimate but as long as you are not missing the false negatives, as long as you're always catching any fraud instance, that's okay. The third option is basically you need very high accuracy results. And in order to achieve that, generalized models are not going to do very well. So generalized models can do fairly well for the other two scenarios but if you have mission-critical systems where you need 99.99% accuracy, you're going to have to solve every single little corner case that you encounter in order to achieve that extremely high accuracy.

Sure. Self-driving cars.

Yep, self-driving cars is a great example where it's like I think Teslas can probably currently drive in a vast majority of scenarios. We had self-driving cars that could drive themselves on a closed track in the '80s, but the difference is once you are off that closed track, once you introduced it to nighttime, once there's a little bit of precipitation, once you're in a different country, once there's an upside-down street sign because a bolt came loose, once you're in some strange construction scenario, or there's some strange municipality that has weird lines on the road, any of that, you have to now account for every single one of those scenarios. And the current state of AI is not capable of generalizing and handling those things. The current state of AI does not reason or have knowledge about the world. It is literally just trained on large, massive amounts of data combined with very specific rules that the human engineers have embedded into this system. And that combination is basically trying to address the very long tail problem of all of those weird little scenarios that it might find itself in.

But these narrow AI applications are very good pattern matching, pattern finding in narrow domains when we can specify them, so we've turned a corner in that. What sort of relationship with artificial intelligence do you see us having at say, the consumer level as AI evolves in its use in these narrow applications?

Yep. So the things that we've really revolutionized since 2012, and the advent of deep learning, is being able to see and understand things and images, given enough data, understand text for specific problems, and understand human speech and really transcribe that into text in a fairly accurate way. For specific scenarios, we are very good at those things. Certainly not in the general scenario. You take any one of those text to speech recognition, for example, it works well if you know what the set of keywords are in the general conversation, and somebody has a regular Midwestern accent, and they don't deviate from that. So I think what we'll see for the consumer side is definitely automated virtual assistants are going to be much more prevalent. Some of those mundane repetitious tasks like picking up a phone call to take a reservation as a hostess at a restaurant, that's something that can absolutely be automated via a virtual assistant. Checking out through a drive-thru while you're ordering fast food, you'll be interacting a lot more with an automated system in that case. But humans always are going to be in the loop there, where if you deviate from the script in any way, if it's instead of, "Give me a second, I'm checking my order." "Hey, I would like this, this, and this." If instead, you're like, "How was your day? How are you doing?", it might be able to have some canned response but it's not going to be able to really deviate from the script at all. So as long as you have the intent very well mapped out of how anyone might interact with your system, those automated virtual assistants, I think we're going to see a lot more of that. And then certainly, for what we do, using AI to automate the process of extracting product information from grocery images, there's a lot of things that go behind the scenes for the consumer that you might not realize that AI is involved in. But we're speeding up the processing time to make sure that that data can be up to date and available and complete. And then we're also dramatically reducing costs. Again, I don't think that human beings should be looking at a product image, writing that data into an Excel spreadsheet, and transcribing it. They should be doing something a little bit more interesting and that's a great candidate for automation.

Sure.

So I think in the next five years, we're going to see more of the ways that we interact with some of those wrote exchanges. Call centers are another great opportunity for automated virtual assistants to be automated, but then also behind the scenes, reducing costs and providing an improved consumer experience. And then certain high-profile opportunities like self-driving cars, or in the construction industry, automating some of the construction processes using robotics.

Well, let's talk about how that interaction with the consumer shows up for Foodspace. How would someone - a consumer - see the impact of your app?

Yep. So currently, if you go to walmart.com, around 40 to 60% of the grocery products have at least one error in the ingredients or the nutrition label, on average it's two or three errors. That reduces trust, certainly. And then additionally, Walmart's main page, you can't filter by important attributes like vegetarian diet, or peanut allergy, or gluten-free. And so to find your products, you've got to type that into the search bar. But the problem with that is unless that keyword search shows up somewhere in the title or the product description, you're not going to find that product. And additionally, even when filters are available, if only 40% of your products have those attributes, if only 40% of them have the vegetarian label populated, you're going to have a poor experience with that filter where you might not trust it, and then you might go back to the search bar, which leads to a poor experience anyway.

Because if you're searching on "vegetarian," it'll give you vegetarian burgers, but it won't give you cream of broccoli soup, because it doesn't say "vegetarian."

Exactly, or like a couple of months ago, we just bought a new bed, a California King. And we were shopping on Wayfair for bedframes to match that. And in their little sidebar, they have a little button that you can click to say, "I'm only looking for California Kings, only show me those options." And when you do that, they only have around 100 bedframes or so for you to choose from when you click that filter. But if you just type in "California King bedframe" into the search bar, there are thousands of options. So as the consumer, this filter isn't working very well; most of their products are not attributed. So you don't want to miss out on all of those other options so you need a, we call it a critical mass of labeled products in order to have trust in the filter. And even if you're talking about things that are assumed to be vegetarian, I think even if you type in "plant-based meats", or "vegetarian", and then you type "meats" into the search bar and vegetarian as the search filter, if only 50% of your plant-based meats have the vegetarian label, if you're a plant-based meat manufacturer, that's terrible, you definitely want to show up in that filter. So that's really where having your data digitized, extracted, and well attributed is super, super important.

I think this is why I and others find that the search engine on some company's site isn't adequate and go to Google instead because it knows about synonyms, and then put in `site:`, the place we're looking for, and then the search and it will match on that. So we know that then in food products, we've got a lot of data available that's not being used, and you can get at it with text recognition, natural language processing, that is an opening. What does exploiting that opening, moving into that opening, what sort of new experience does that make available for someone like me?

Yep. So I think in the near term, it's having trust in the underlying data, having that data be available, producing those nice little filters in the sidebar so that you can select your allergies or certain health criteria that you're looking for, like "I am on a low sodium diet so only show me products that have less than 10 grams of sodium" or something, as well as anything that coincides with your values or things like non-GMO, or fair trade, or it's recyclable containers. So really optimizing that experience to say, "How can I find the products that I'm looking for?" and increasing that discoverability as opposed to having to go to Google or something to try to find the best products. It should happen at the retailer website.

So that's a B2B model, and Walmart's got to buy it. Is there an opening for a B2C application here?

Yep, so we've definitely floated ideas of bringing that to a consumer app where you can kind of have a Google search for food products, and you can discover things that way. We've also had conversations with smart appliance manufacturers, where you might have in your smart fridge, some of this data digitized and being able to interact with your retailer, or even recognize certain things like your milk is running low, so we will alert you that, "Hey, you might want to order some new milk," automatically recognize the nutritional information and the profile brand name, product name, and then add that to an Instacart order and automatically say, "Hey, would you like to order this?" We've also talked about kind of automating microwaving or oven cook times where simply by taking a picture of your product, it will automatically set all that time for you on the stove or the microwave or the oven, or even preheat your oven based on the product information that you're saying that you want to cook. So it's like saying "I want to cook Stouffer's lasagna before I get home," and it automatically sets the preheat setting on your smart appliance for that.

So I have a question about data ownership because this could equally well generalize to any other industry. Who owns the information that Foodspace extracts from these images that belong to Walmart and the other companies? You have pulled out of that, data that they didn't have, you have aggregated it to create meaning. Walmart would presumably argue that they own that because they paid you for that. And yet that it would be useful in many other contexts because you could combine it with other images of the same product that other companies have and go even bigger. Is this a lawyer's paradise for arguing these points?

Yeah, so the current industry standard is the companies that are performing the data extraction own the data. And if they want to, what we call, syndicate it-- So our primary customers are the brands themselves. We have found that the retailers, because the e-commerce sales are a smaller portion of their overall sales and they tend to be a little bit slower to adapt to some of these innovations, we've found it easier to just go to the brands themselves, which it's interesting that a Nestlé might not have insight into their own chocolate, for example. But the brands themselves due to this large network of third parties, and you're outsourcing the people who are creating your product information and the product labels, and the boxes, the brand managers, if you're a marketing manager, you don't actually have access to your own data. So if you want to figure out the nutritional information for your product, either you're searching very disjointed arcane databases, or you're actually going to the physical store and purchasing it. So that experience is certainly not great. So the current model is that the company that's extracting the data owns the data. And so if you're Nestlé and you want to provide that data to Walmart or to Target or to Albertsons, you're going to pay a separate licensing fee, because you don't own the data, to each one of those retailers. We don't do that. We share the data with the brand. So we will provide that data back to you, you can always access it whenever you want. But we will also syndicate it for you if you'd like for a small fee. So if you want us to send the data to Albertsons, we're happy to do that. If you want to get the data out of the system and send it to Albertsons and not pay a fee, that is also absolutely a happy path. So the current model really allows us to derive those insights and provide that data, that very complete and accurate database to a large group of people such as smart appliance manufacturers, but also maintain that partnership with our brands to ensure that it's mutually beneficial.

And once again, I would say to our listeners: generalize this beyond the application we're talking about because here's an issue that occurs to me is applicable across a wide set of domains. We've been talking about data, and the obvious example of that would be what you've done, OCR on the labels to extract. But then there's meaning, and we want to end up at models, engines that have been trained to answer things like, "Plan me a diet that's got the optimal things, but low sodium, low magnesium" or something like that. That's more than data.

Now, you've got to have an engine that's trained on that. Is Walmart going to buy that or are they going to rent access to it? Do you see how we would get to that?

Yeah, so definitely down the road, we would like to be more involved with the recommendation side of things and the search and discovery side of things and provide those insights. We would use a software-as-a-service model. So instead of selling the model itself, it would be you're accessing it every time you want to interact with it. Does that answer your question?

Yes, I think it does. So when you're constructing that - and maybe that's a little way in the future, but if you think about when you get to that - would you need to train it on specifically, the different types of question I might ask, like, it's reasonable to expect someone might want low sodium, so do you go in there and tune it to answer the question, "What's low sodium?" But can you also code it so that it can handle things you didn't anticipate, like low magnesium diet? I'm just making up that that might be useful for someone. And if you haven't thought about all these kinds of questions but can you generalize it to be able to answer "low something diet," "plan something that has got nutritional targets, but is low in X" for any value of X?

Yeah, those are really great questions. So I think on the search and discoverability side of things, we would really love to be able to help in the parsing of the search query itself and optimizing that based on attributes. The way that current search works is mostly the Elasticsearch, inverse index model, where you're basically matching on this term frequency-inverse document frequency model, where you're counting the number of tokens that show up in a particular title or product description, and then you're normalizing by the number of times that token shows up in all documents. So if you take a very common word like "the" and it shows up a lot in a single document, but it also shows up in a lot of documents, that's going to be a very low score. But if you have something very specific like "gluten-free", that's going to be a very high score. The problem with that model is it's very fuzzy. And so if you type in "plant-based meat", for example, the plant-based is a hard filter, you should not return any meats that are not plant-based. But in the TF-IDF model, plant-based might be kind of an obscure term. So you might see plant-based results show up at the top, but you're also going to see other meats that are not plant-based as that fuzzy match, which is not optimal.

And just to interject, I used to write search engines before Google, so I know that when you have a query like that, that doesn't have Boolean operators in, you have to ask yourself, "Is it plant-based *and* meat? Or is it plant-based *or* meat?" And typically, what you would do is you would say, "Well, I'll give you the results that are the **and** first but if there aren't many of those, then I'll give the results to the **or**."

Exactly. So what we would like to do is identify which of those attributes are more structured queries that should be a hard query versus the soft, fuzzy, TF-IDF match. So if you're adding gluten-free snacks to your query, we want to return snacks, and also terms that are semantically similar to snacks like cookies and crackers. But that gluten-free, we know is a hard query, treat it the same way as if somebody clicked on that filter on the left side so that it really structures that query. And so parsing any search query by Boolean hard query operators versus that fuzzy string match operator. And moving beyond that to kind of optimizing health and diets, we are very cautious about that area. One of our nutritionists really hates to say anything is healthy or unhealthy and be prescriptive or proscriptive in that way. And that's because depending on who you are, what your goals are, what your values are, healthiness can very much be subjective. So we want to be very cautious about telling you what to eat or what not to eat. We want to just make sure that that data and those tools are available to the consumer so that they can make the most informed decision as possible.

One of the things that I don't see in any of these engines because it's a really hard problem to solve, I wonder to what extent people are thinking about it, is a dialogue. Like if I go into a library to the reference desk and ask a question, I don't just bark at them, "Avocados Mexico," and expect them to come up with a result. Number one, it's rude. Number two, they will say, "Well, what specifically did you mean about that?" They will have a dialogue with me that gets me very rapidly to a level of understanding that Google does not because it is a stateless transaction. Do you see any avenue, do you have any plans for dialogue that would say, "Do you mean plant-based *and* meat or are you looking for something plant-based and *something else* that is meat?"

Yeah. So I think some of our partners, the easier and more structured way of doing that is basically allowing you to set up a profile for yourself as you sign on to the platform so that you can prespecify those values that are very important to you. Whether you're only looking for Fairtrade products or pescatarian, or you follow a gluten-free diet, or you're on a South Beach diet for right now and being able to really just structure those things in a query. I think if what you're referring to is that more conversational AI, virtual assistant solution, it's a hard technological problem right now and that's why we don't see the dialogues. I think, for anything that's conversational, and being able to follow up outside of anything for a script, so if you have that very well-defined script, where you can say, okay, you're looking for plant-based. "Well, did you mean low sodium plant-based? Or is that not important to you?" If you can write that kind of tree of decisions that a consumer might ask, then that is certainly something that we could provide in that virtual assistant framework. If instead it requires the nuance of reasoning and pulling from an external data store and taking those external elements, and building inferences from them, I would just say AI is not there yet. I'm very interested in conversational AI myself; I've attended the little subgroups at NeurIPS and there's a little group that gathers like 100 people. So I'm very excited about where conversational AI can go in the future, but we're just not there. And I think a good illustration of that is GPT-3, everyone's very excited about that. They generated a bunch of fake news articles, and they asked humans, "Can you tell the difference between the fake news articles that were generated from GPT-3, versus actual news articles?" And in general, the humans were not largely able to tell the difference. But the one that they were most able to tell the difference was a news article talking about Joaquin Phoenix wearing a paper bag over his head that says, "I'm not famous anymore." I'm not sure if you're familiar with the story but actually Shia LeBeouf did this. It was kind of like a freakout moment for him. He was at some red-carpet event, and he put a paper bag over his head and said, “I’m not famous anymore.” The nuance of that problem where GPT-3 picked up a story, all of the elements of the story were there, but then it associated it with the incorrect person. And anyone who’s familiar with that story knows that that’s the incorrect person. But I think that why that was picked up as an error illustrates why conversational AI is so hard. GPT-3 is doing statistical correlations between words, and it has associated that Joaquin Phoenix is a person, probably an actor, and then they have this story about an actor doing something kind of freakout and weird. So to the model, swapping out those two things is entirely okay because it’s a statistical argument. It probably didn’t see a lot of those examples of those news stories so it’s easy to do that. But a human knows that those two things are so incongruent because we would know if Joaquin Phoenix did that.

Right.

So we have this data store that says that is impossible, those two things can’t happen. I’m quite certain that GPT-3 knows who the first President of the United States is. That’s something encoded into the knowledge, into the neural network because there are so many examples where George Washington and the first President of the United States are probably co-occurring. But it doesn’t know that as a hard fact, it’s just a statistical correlation with high confidence. And I think that’s the problem with conversational AI, where it doesn’t have any internal knowledge store or these hard inferences, where it’s like, “This is something that is entirely incongruous. This person just told me that they were a vegetarian, but now they’re asking for beef. Wait, do you mean plant-based beef? Because that doesn’t make sense to me.”

Right. Or maybe they’ve got someone who’s not a vegetarian coming over. Although, I’m going to have to look up the example that you cited there. When you introduced it as fake news, I thought that meant it was tasked with producing stories that were false but looked plausible. And in your case, it was supposed to be ones that weren’t necessarily false. Because otherwise, that’s a classic fake news kind of tactic to take something that’s mostly true and change the names and attribute it to someone else. And that way lies Bill Gates conspiracy theories and things like that. Now you’re getting to something here, where I made a note earlier, where you were talking about AI is now able to understand the various kinds of things, text, and so forth. And using that word is loaded, the word *understand*. And I wonder if we can find a better one because it produces a near reflex action in many people who say, “No, it doesn’t understand anything, it doesn’t have common sense. It has no idea what the real world is.” All of that is true but it’s very convenient to say that it does understand it. Is there a better word?

Yeah, that’s a great question. My knee-jerk reaction would also be to say that it does not understand. It correlates, I think it’s pattern recognition. And as long as there is a pattern to be learned, machine learning can shine very well in large data-rich environments where you have a large number of those examples, and you don’t have to worry about deviating from a very well-defined path. That path can be extremely complicated, like generating a news article. And sorry, when I said fake news, I didn’t mean just generating any news article, those indistinguishable. So we are talking about very long paths of generating hundreds or thousands of words, but it’s very well defined and there are well-defined grammars that allow that to succeed. The moment that AI fails is when there's a single bit of information that completely makes everything incongruous. And it's really the fact that there are no hard data stores in these deep learning models to allow you to be like, "Wait, that makes no sense." Just because you can swap out an entity doesn't mean that it's going to make sense. There are some approaches and research that has been trying to address that, things like neural Turing machines or neural memory networks that basically try to encode knowledge as hard encodings. And then basically, you can look up those facts in the neural net itself, and then you softly attend across them. Those have seen marginal success in certain scenarios. But still, that problem of taking two pieces of information that are very far apart from each other, that don't just show up in the text somewhere and actually understanding that, that's a challenge. I'm sure that GPT-3 has understandings of two plus two equals four, and can do addition, but only so far as it has seen those examples. If instead, you try to extrapolate to what is 22 plus 24, if it hasn't seen that specific example, it's not going to be able to expand into that generalization.

Right.

There are models that try to address that specific problem of arithmetic, and it will get to the point where it can do three digits plus three digits, but then when you try to go to four digits plus four digits, the accuracy starts dropping. Or once you get to five digits plus five digits, it can't do it at all. And that is so different than how humans work where either you're going to be able to solve addition for all of those problems or not really for any of it. It's not, "Oh, now it's more digits, I can't solve it."

And to some extent though, we also do a lot of that in patterns, like we memorize times tables, so that we don't have to do it on our fingers or think about it. And I think one of the insights here is that it's interesting how much of what human beings call "understanding" can actually be replaced by pattern recognition. Even driving a car most of the time is operating on what think Daniel Kahneman called System 1 thinking, you're just pattern recognizing. That's why you're able to tune out and do something else, and you get there and have no idea how you got there. But if somewhere along the way, a rabbit runs across the road, suddenly you are summoned and your pattern recognition system is very good at saying, "I don't know what to do here straight away. Help." And you show up.

I love that reference to Daniel Kahneman's System 1 thinking. System 1 thinking tasks are probably something that can be mostly solved by machine learning in its current state. System 2 where you need that reasoning, you have to stop and think and take a step back; that's, I think, where we're going to see that really great division between AI automation, and then humans doing what they do well. I haven't heard that before, but I really like that comparison.

Well, this is a good place to look then into the future. Ten years from now, where do you think artificial intelligence will be? And where do you think you'll be?

Yeah, I think 10 years from now-- I am certainly very bullish on AI. I don't think that there's anything fundamentally different between the human brain and a computer system and being able to model that. The human brain is an incredibly efficient system. There are 3 billion neurons and quintillions of connections between those neurons. In order to replicate that in a machine, you are talking about many massive supercomputers that are an order of magnitude more efficient than the current supercomputers. But I think fundamentally, the limitation there is just a hardware one, getting to the point where we can do enough compute to arrive there. I think certainly, more generalized tasks, there will be general models that are floating around, and you'll have conversational AI models that you can basically download from the internet and fine-tune to your task. Or largely, you can have a reasonable conversation with this conversational AI out of the box. And you could fine-tune that to your domain, say, to be a therapist or a call center operator, or even more nuanced things like I don't know, a podcast host, maybe. Those opportunities, I think are going to be more there in 10 years from now. We will certainly see a lot more of our machinery, automated things like self-driving cars, construction vehicles, trucking vehicles. Even airplanes might be automated, as scary as it seems to fly in an airplane that's driven by a robot. But if the statistical error is less than human error, I think that's a strong argument in favor of that.

Right.

Will we have general AI in so far as the singularity and being able to compete with humans on all intellectual tasks? I'm not sure. I mean, it's too difficult to predict that. Andrew Ng has a metaphor about the politics on Mars when we haven't ever been.

Right.

So it's like, it's speculative. I think the world will be very similar to the way that it is today, but also very different. And I think that if we looked back 10 years, if you thought that you would have a little speaker in your house that you could talk to and have it set timers and control lights, you wouldn't have thought that was possible. But at the same time, the world isn't so dramatically different that it seems science fiction.

Sure. She's sitting right next to me. So in your business, do you have a goal or a vision that's like something on the wall, a question that you would like AI to answer that would tell you, "Hey, we've gotten to the ultimate goal of what you're pursuing at the moment"? Do you have a little dream in that respect?

Yeah. So my dream is to create an AI company that can adapt very quickly to these narrow, specific problems, that can see and read and understand the problems that are specific to their domains, and really decrease the time to deployment of those solutions to solve these very mundane and repetitive tasks. So to me, creating the training pipelines, the generalized models that we can use transfer learning for, things like BERT to transfer learn to the NLP task, or the image net models like ResNet, to transfer learn to identify key regions in a product image. Building that system so that we can target one industry after another to really reduce the expenses and allow human beings again, to be freed up to do that more creative work, that's my dream for the next 10 years probably. I'm not necessarily trying to solve the general AI problem. I'm trying to apply AI to enough specific scenarios very quickly and efficiently so that human beings can do more creative work.

Got it. Fantastic. Well, thank you, Daniel DeMillard. How should people find out more about you and what you're doing?

Yeah, go check out foodspacetech.com.

All right. Thanks for coming on the show. It's been a fascinating discussion.

Yeah. So much fun, Peter. Thank you so much.

That's the end of the interview, and, wow, I hope you love this stuff as much as I do, because I could ask questions about it all day. And probably have done so for many days by now. And of course, if you know someone else who loves this stuff as much as this too, tell them about the show, because it's not like they'll have seen the *AI and You* blimp advertising in their neighborhood. But that was a great example of going from general, wide-ranging topics about AI to a specific, practical application, of using AI to find food that meets your requirements. It also gives you a good idea of what sort of application is potentially monetizable in today's AI economy.

In today's news ripped from the headlines about AI, *Wired* magazine reported that the Pentagon is improving its cybersecurity by using AI to hack itself. A so-called Red Team – that's the generic term for a group that's commissioned to find what's wrong with an organization's product or systems – used machine learning to look for weaknesses in their own AI models. Gregory Allen, director of strategy and policy at the Joint Artificial Intelligence Center – does that sound cool or what – said, "For some applications, machine learning software is just a bajillion times better than traditional software." But he added, machine learning "also breaks in different ways than traditional software." If you've been listening to much of this show, that won't come as a surprise. We'll have a future show where we talk about explainability of AI and how AI systems are being made more robust and reliable, or at least, the opportunity to do that is being created for some of them.

But next week, I'll be talking about disinformation and misinformation. Me, myself, and I. Because I promised that way back in episode 1, and it's about time we tackled it. AI has a massive role in that, not always good. So we'll dig into that next week, on *AI and You.*

Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

http://aiandyou.net