AI and You

Transcript

Guest: Bryant Cruse

Episode 101

First Aired: Monday, May 23, 2022

Hello, and welcome to episode 101! My guest today is Bryant Cruse, who has been a pioneer in the application of AI technology to difficult real-world problems. After serving for eight years as a Naval Aviator he returned to school for an MS in Space Systems Engineering from Johns Hopkins. While on the Mission Operations team for the Hubble Telescope he found a personal mission to change the way spacecraft were operated by seeking a way to capture human knowledge in computers. This work led him to a six-month residency in AI (MS equivalent) at the Lockheed AI Center in Palo Alto. He went on to found two successful AI companies, and now New Sapience, which offers tools like companion AIs. We'll learn more about that in the interview. Just to give you an explanation of a couple of things referred to in the interview, Minsky's Society of Mind, that was a theory of Marvin Minsky, one of the founding fathers of AI, if you will, and also the title of a book, where he said that our minds are built up of putting together simpler parts, starting with agents, which are themselves built up of even simpler parts, which do not themselves have the quality of intelligence of consciousness. So a theory of emergence, if you like, the whole is greater than the sum of the parts. And then there's a reference to stochastic parrots, a paper by Timnit Gebru and others, for which she was fired by Google in a notorious incident. And that phrase is her description of transformers like GPT-3, which are sophisticated language models but the term stochastic means relating to probability, and so this was saying that those language models are driven by random numbers applied to the language inputs they had, which are gathered from the Internet and hence already biased racially and in other ways. Let's get to the interview.

Bryant Cruse, welcome to AI and You.

Glad to be here, Peter, thank you.

And you started out your career flying planes for the Navy. I can't resist saying *Top Gun* and Tom Cruise but it wasn't quite that and plus different spelling of Cruise. Tell us something about that so we have an idea of where you came from?

Well, it was kind of unusual to find myself in the Navy because I went to school across the wall from the Naval Academy in Annapolis, Maryland. But it was St. John's College where I studied the great books, it was a liberal arts curriculum. But after four years of discussing philosophy and things like that, I was kind of hungry to actually do something rather than talk about it any more and so I found myself as a naval aviation officer candidate, ended up being stationed back in Patuxent River, Maryland, with strategic communication, flying the C-130 Hercules, which was about as far as you might think, as you can imagine, from discussing the great books, but it was a great experience and one thing about the way the Navy operated its mission, since we were a 16 person crew on a big airplane and we were told basically to go out and for two weeks and operate, more or less independently to do our mission. So it was a great education, and practical reality, it was no longer theoretical, you had to do things, right and you had to be situationally aware of all the factors, the environmental factors and so it really had a profound effect on me going forward about the mental and cognitive skills you need to solve real world problems by being focused on what's in front of you. I always call it the skill of an engineer. An engineer looks, as opposed to the scientist who thinks about what might be there, what should be there, what are the underlying causes and formulates theories, an engineer has to be very focused on what's there and I was just talking yesterday with a friend of mine about the difference between being a scientist and an engineer and that--we can talk about theories all the time and they are maybe the best knowledge we create, I believe about the universe--but in the engineering realm, a bridge either stands up or it falls down, in the end, if there's a hard test in reality. So I thought that part of my education, to add to the theoretical background was very important to me.

And after that, you moved to the Goddard Space Flight Center and Lockheed for the Hubble Space Telescope. What did that bring to your education in this respect?

Everything kind of connects together, there's so many threads, the warp and weave of a person's life. I really enjoyed aviation, but I felt a very strong connection to space and to being part of that great adventure. So when I left the Navy, I went and entered a program in Space Systems Engineering at Johns Hopkins, a master's degree, but I only got matriculated, then I got my first job in the industry and found myself very quickly on Hubble and working on it and my job there was, I was in charge of the operating an operational analysis and the training of the control center operators for the onboard data management system. And when I walked into the control center and saw how they were operating spacecraft, which was basically people, individual, highly trained technicians, basically, on the onboard data management system, sitting at a console with a big screen and watching numbers on screens, which were tables of numbers, which were telemetry. I know, Peter, you probably are very familiar with that from control centers and things that go into JPL. But from my perspective, as a Hercules pilot, I was kind of horrified because I said, "Well, this is a science project. How could you fly the vehicle" -- and we're talking about a vehicle that hadn't been launched yet by the way, we were three years from launch at that time, bigger than a Greyhound bus, up in high Earth orbit--by looking at numbers on screens, which are telemetry values, with mnemonics like, CBAT 6VLT's you had to know that meant battery six voltage, okay and then you see a number you know, 28.2, 28.1, 28.24, you know and based on all these numbers at the same time you're supposed to know what's going on with the vehicle and you're supposed to know that if you send a bunch of commands, which are also encoded in hexadecimal, that the vehicle is responding normally.

So was your reaction based on "this is a cognitive overload situation," like you're coming at this from being a pilot, where reaction time is everything and so the avionics systems are engineered to make that information go into your brain as smoothly as possible – and here's someone looking at a console where they've got to translate what that acronym means and then they've got to translate the digits into an analog representation in their head, of, "is this good or bad?" and that's an overhead, but they're also they're not flying something where they could be shot down or they have to react in milliseconds: how much does it matter?

It matters hugely because the Hubble Space Telescope was a national asset; still is, it was irreplaceable. I forgot how many billions of dollars have gone into it from the beginning, it was irreplaceable. No, lives will be lost; but spacecraft have been lost from sending the wrong commands, from misinterpreting telemetry: it's the wrong way to do it.and it became a mission of mine, to change the way satellites were operated. I came in there and you know I was trained as a space systems engineer, I knew what those numbers meant, but human beings aren't good at keeping up and doing that translation in real time when you have to keep up we have a real time stream, there were 4000 telemetry parameters -- space station has hundreds of thousands, I think, but 4,000 on the Hubble at that time -- that's engineering only not science and so we had 6 or 7 people trying to keep up with those numbers and I can recite you cases where spacecraft were lost because they couldn't interpret the data.

We might get into that; but let's go over, what does it mean to keep up with those numbers. When they're looking at those displays, what are the humans doing with them?

Well, you want to know what the state, say, of the battery is, you know if that's part of the data or the state of the onboard tape recorders at the time and so you know that if this parameter was in a state one and this one is between that number and that number, and this fourth one has a certain value and a fifth one has a certain value, then the state of the recorder is "playback," okay. That's a lot of cognitive processing and it takes time to do that. So what you want is like you go to your stereo system and that says it's in playback is gonna be in that sense, it's the graphical thing but I saw the problem is I know what the telemetry means, give me some enough time and I can get there. But what we wanted, what I saw was a problem, I wanted to use a computer, that's what they're good at, is data processing. So I saw the problem as one of putting the knowledge that I already had in my mind, my expert knowledge into the computer and let it use its speed and repeatability and deterministic processing loop to unequivocally and effectively and accurately translate those numbers into vehicle states, which could be presented to me just like the cockpit of a C-130 actionable information that I can really understand and take action on. So that became my mission.

So I think the question for me here and maybe the audience, is what level's the human in this operation center occupying in a command and control hierarchy or loop the data instruments reflecting this stream of telemetry back to them that's just labeled with at the level of this instrument with this code has this value right now. That isn't what they're there for. At some point, they've got to accomplish mission parameters like "point at this star and open the shutter to take this image" without pointing at the sun in the meantime or running out of propellant or all kinds of other things, which are pretty algorithmic but hadn't been instantiated anywhere except that humans brains and...

They weren't really algorithmic, they're procedural and in order to do a science experiment, for the Hubble, it the process actually starts months in advance. When you pick out the target, you pick out the instrument, you pick out the instrument setting and then you start working on ephemerides, where does it have to point and so all this gets worked out in this set of procedures eventually gets converted to a command load. Okay, which then you load into the vehicle and then on the appointed time you kick off that load and the onboard commands, execute the procedure that the humans had to meticulously create. So a lot of that is because of the way the system is designed, you can't fly it like an airplane. It's not designed like an airplane; it's designed like a science experiment and, of course, you would still have to do a lot of planning, just like to fly an airplane, you have to do mission planning and you have to sit down ahead of time and figure out where you're gonna go and how much fuel you're going to need, apply your expert knowledge. But the control mechanism for spacecraft and the interfaces and all the things are extremely expertise-intensive, extremely data-intensive and extremely dependent on humans in the loop. For instance, in a C-130, if you got a fire you got to handle lights up red, got a fire, pull the handle. They can't do that because of their control system for spacecraft. So they put it in safe mode and says to stay safe, we'll figure out, we'll analyze all the telemetry and then we'll figure it out. So that's how I got into artificial intelligence. Yes, I want the computers to do this for me, how can I get knowledge, that my knowledge into the computers so they can process the kind of knowledge I have, but they can do it with their speed and reliability and ability to crunch the data?

Because you saw a roomful of people acting like machines? Carrying out machine-like operations, and you want machines to do that. And that's one side of an equation with AI; and another one is, are we getting machines to act like people but that's further down the road. We don't have to get into that. I think well, that will be the second thing we talk about. I'm also thinking that what you describe is analogous to a nuclear reactor control room, right? The human goal is produce power without melting down and you've got 1000 dials to maintain to do that and we know that that created cognitive overload in some well documented situations...

And their actual control interfaces are in some ways more straightforward and better than what we have in our control room. Although, at least over the years, they've gotten somewhat better at presenting graphical schematics and things like that, where the telemetry is not just in tables, but they've become more advanced. But yes, the problem of cognitive overload based on human beings are data processors and we can do calculations, but we don't do it as reliably and fastest computers can do it. So whenever you have that situation, you want to push the equation toward the machine and let them do what they do well and let the humans do what they do well.

When we put it like that, when you say it like that it sounds obvious; it sounds like well, yeah, of course, why wouldn't we and so: why didn't we, why did it take so long?

Oh, they tried, I tried, my first company was designed specifically to solve this problem by using expert system technology at the time. Expert system technology, now called symbolic AI, is based on the premise that you can solve these kinds of problems by manipulating symbols using a computer's ability to do logic in a very deterministic fashion and that you could do it by giving it an X number of rules. And the problem I had at that time is the expert systems engines that they had were computers weren't very fast for them at that time. So they took up a whole computer just to run the inference engine. So what I did at my first company, NASA sent me out to Lockheed AI center to study AI and helped solve this problem and I took up with some men engineers out there and we found a project to do to the first what we called the real-time expert system. In other words, an expert system engine, a rule base engine that had enough performance to keep up with the telemetry stream. So in order to do that, we had to put the user interface on one computer and the inference engine on another computer and the data interface on another computer and then we had to write the software to get to computers to talk because the Ethernet hadn't really been invented yet. But we did it, but we ran into the same problem that everybody ran into with symbolic AI and that the notion was we've relying on an algorithm logic to do that and I remember being told when I went through my residency in AI at the Lockheed AI center and we had the AI professors come up from Stanford and they said, "yeah, this is how you think we actually have rules and you applying these, you're just not aware of them." And I was kind of skeptical at the time that I thought of everything in terms of if-then rules. But nonetheless, that was what the wisdom was, in that case. But it doesn't scale did it? I could analyze my first experiment, I successfully analyzed 100 telemetry points and the program took about 15 minutes to run, it was better than a human could do, in terms of its reliability, but it wasn't quite real time. And when we tried to scale it, it just didn't scale and so actually we abandoned that and my next company, we came up with a way that actually, arguably was the first time we succeeded in putting knowledge into a computer.

Well, I want to get into this relationship between the symbolic AI and the connectionist AI and where we go from there. And just to fill in the gaps, perhaps for the audience on what you were talking about, you're using PROLOG or something like it, it perhaps for that symbolic rule processing and I remember we were trying that at JPL in the navigation section just experimenting with that. And although I wasn't central to that, my impression of that was that you could easily demonstrate simple examples, like you could put in the rules to play the animal game - guess the animal - and the system was one way, instead of having to write out a whole set of if then else blocks and get them in the right order, you would write them out in any particular order you wanted and the system would figure out how to connect them together to get a result. But either you were solving something so simple, you didn't need it in the first place or the number of rules that you did need was too great for you to figure out what they all were.

Exactly, they all started out well, but then the performance versus effort flattened out and became asymptotic and so you could never get to the level of performance you need and that was universally proven.

And this is the point in AI history where AI took down to where people said, well, this is never going to amount to anything and it was rescued by the deep learning and neural networks and the hardware coming up to the point where we can execute those. And now we have systems where we have no idea how they work. But if you train them, they can surpass human performance on pattern recognition type of tasks.

Pattern recognition, yes, I mean, but they're statistical beasts, aren't they? If you have a problem whose solution lies in statistics, it's very powerful.

But not the other way, like to a machine learning expert, if you have an engine and you ask it, what's two plus two and it says 3.97, they say great, we're done. That's well within the bounds. That's not the class of problem that should be solved that way. Facial recognition sure, would be a great result. So now, at some point, you started getting into this field of AI deeply enough to see the limitations of the good old-fashioned AI, the symbolic AI and the network-based AI. So what was that train of thought like for you?

Well, I actually went a step beyond symbolic AI and I mentioned that I had a second company still focused on telemetry analysis. We took a different approach in that we said, well, what is knowledge? If you look at the functionalality we're trying to achieving computers, we talk about intelligence. But you know, if Albert Einstein's brain was in a Cro-Magnon, theoretically had the same intelligence, the same IQ but without knowledge, he wouldn't have figured out relativity right? So you have two aspects, you have intelligence and you have knowledge and we talk about what intelligence is and we know intuitively what it is, in a sense, we don't know how it works. But it's that characteristic, that quality that human beings have in such a degree compared to other animals. So it mounts to a difference in kind. But specifically, we recognize it by its results; because what really makes us unique is our ability to alter our environment to an extent that no other species can do that we know. So what is that? Well, it's knowledge, it's the knowledge of the world. So it's really, the approach that I took after my experience with the failure of symbolic AI was, maybe we don't need a science of intelligence so much, artificial intelligence, we need a science of artificial *knowledge*. So the failures of the past and the failures of past epistemologies were basically well, knowledge is a bunch of facts, right and even in classical epistemology, they would talk about the whole thing as to prove the truth or falsehood of assertions. But what if that isn't what knowledge is? So this is all happening, the first step I took toward this was back in the '90s, we didn't call it AI because you couldn't because you're in the middle of the AI winner and if anybody said, AI, they laughed you out of the room, we just called it advanced automation technique for spacecraft analysis. But, going back to my undergraduate notions of epistemology, I realized that this knowledge that we have, about the world about the universe or I should say about nature, to be more specific, is to the scientific method. And those theories are models, right? They're not a bunch of facts, they are carefully constructed models. For instance, the Ptolemaic model with the Earth at the center and all the planets and the stars going around there. That was perfectly scientific, they had certain different principal premises, they started with that they looked at the phenomena, they looked at what they saw, they formulated a theory about what was going on underneath of it and then they see if they can make some predictions based on that, which is straight scientific method, even though it was before Francis Bacon came along and articulated it better. And it worked for certain things very well. You know it worked for predicting when the eclipse was going to come, it predicted when the spring tides were going to be come and it predicted when you should when you should plant your crops and it and it told you how to navigate a ship around the surface of the globe, which was understood to be spherical and in fact, it did that very elegantly, with a very simple instrument, called a sextant. Now later, Newton came along with a with another model and we said that Greeks didn't know what they were talking about, it's the Sun that's in the center and this was different premises and it's really about force at a distance and mechanics and all this stuff and it was a better model because it explained more stuff, right? The same model now explains when the Moon rises and when it sets, and why an apple falls from a tree. That's cool. But you know, and this is a great insight for me. So when we were out there flying C-130s over the ocean, one of the modern navigation tools we had is something called an inertial navigation system and it was an entirely Newtonian animal. It was this big thing that was strapped in the rack and it had inertial stabilized platforms and accelerometers and gyroscopes and you start with you on the ground, you punch in your lat long and you take off, and it tells you where your lat-long is all by solving the integration of the third differential and it was really cool because it was really complicated and it would break a lot and it would fall down. So you know what we did? We had a little hole on the top of the airplane that you could open up the chip up this little thing, which is essentially a telescope. No, it wasn't, it was a sextant and we cranked in the starshot. We looked at the tables and we figured out where we were. Those tables basically were the same ones that Ptolemy put together in Alexandria. So what did I learn about scientific theories and knowledge? It's not belief about absolute, it's about utility.

And utility is where I want to go with this because you weren't just sitting behind the desk in some tenured university department, thinking about this, you have gone into business to make this work. That requires a whole higher level of empiricism and making it work on the ground. What is that experience of putting your ideas to the test been like?

Well, it's been very rewarding, really, I mean, in my case, I kind of from the beginning, saw myself as a scientist or natural philosopher, but I was very interested in the theory and epistemology, if you will, but I had this very practical - through the Navy and being a Space systems engineer - trying to solve real problems and as I said, these are engineering problems and it either works or it doesn't work. When, for all the theory, whether I was big on the theory of the fact that or the thought that symbolic AI was theoretically sound and that it was or not when it didn't work, it didn't work. So I had to look for something that worked. And I did go back to theory; or later, the theory informed what I did, but what I call what I'm doing or what we're doing at New Sapience, right now in AI, which is working, I call it sort of, like, I say, it's like Roman engineering. Well, what's Roman engineering? Okay, well, some point somebody in early Roman Italy, someplace, was trying to get across the stream and they piled up rocks in a certain way and they made an arch and they didn't know how, why it stood. But they saw that it did stand. And once you can build an arch, you can build an aqueduct and if you can build an aqueduct, you can build the Pantheon by spinning that arch in three dimension. And the Pantheon is still standing. And recently, our modern day scientists went in there with all kinds of strain gauges and interpretation with all the knowledge of resolution of forces and statics and dynamics and they measured and they found that why the Pantheon had those interesting little ridges around the edges. It turned out, if you took them off, it would fall down, now we have the theory that understand that now, but that's kind of what we're doing. We went in and we did start with a theory when we do get New Sapience because we're looking at knowledge. So we started with this conjecture. Okay, so we look into our minds and we look at ideas, we don't know how our brain does it, we have no idea really. Yeah, it's a neural network, but layer upon layer of complexity and this and that the other. 6 million years of cognitive evolution and biological evolution. So wow, that's really something but well, can't build an artificial brain. But we can introspect the knowledge as it's useful to us in our minds and we see that we have ideas correspond, we think, to the world, at least as our knowledge allows us to predict a phenomenon. So we have enough knowledge about the world, we have ideas about what will happen if I jumped off the Empire State Building and we can test that and so we know not to. And so the knowledge has utility. Our models, our ideas have utility, in helping us change the world, which is what we're trying to do, right?

And these ideas in our brain and I think this is the central problem of general AI, are in encoded are represented in our brains in some structure that we have no idea what it is, but the only way that they come out of that brain is in this linear form of words through our mouth, which is (a) linear and (b) ridiculously compressed to then try and construct from that the original representation where we don't know what that original representation was, is been the reason we've been chasing our tails for decades.

Well, yes; but what we did we backed off one from that, let's not get to the language problem, let's not get to the communications problem, because if you look at it, you can communicate knowledge without language. The first Cro-Magnon chipping a stone arrow head and "come over here "and chipping like they could show they had show but no tell, let's say show until that show. So clearly, you can have knowledge, and this is one of the insights that is fundamental to what we're doing at New Sapience. Back at St. John's, reading the philosophers, always this big discussion about whether you could have knowledge and intelligence without language, or whether it was dependent on language. And we've distinctly come down on the I believe or in the practice that no, they're separate. Language is a communications protocol. You have to already have knowledge in one mind and a commensurate store of knowledge in the other mind in order to be able to communicate ideas because the words have to relate to ideas that already exists in each other's minds, so knowledge clearly precedes language. Once you have a minimum core of ideas, only then is language possible. And you see that, as a child is learning the world around them initially, to sight for direct sensory input and only gradually do they start to associate words with their sensory experience. So the knowledge precedes that. So what we did at the New Sapience is we looked into our minds and we tried to say, kick the language out of it, but what do you see? Well, we could see that our ideas, the more complex ideas were composed of simpler ideas and that it wasn't just a bunch of random facts floating around; that our knowledge was models of the world and models are carefully constructed, in this case, ideas that are related to other ideas in specific ways. So this looks like the way, imagine, you know the way material atoms and connect together to form if simpler atoms to form more complex molecules to form more complex materials. So we took this and said, well, since that looks the same, let's imagine we take our ideas and we start breaking them down into simpler ideas and let's make the conjecture of Democritus and say, let's keep breaking them down to the point you can't break them down anymore.

It's like Minsky's Society of Mind.

Maybe, but in the end, what we end up with, is there are then atoms of thought, that is fundamental ideas, that are the building blocks of all knowledge and all other ideas below which you can't go and if there is, do they have properties that control their connections? Okay, so that they are constrained to connect together to create more complex ideas that are fundamentally such that what they create is knowledge of the world and not nonsense. And it turns out, yes, both of those things are true and that's what we have done at New Sapience, we have identified atoms of thought and we've identified the rules of combination and we've classified them such that in a way that would correspond to say, analogous to the periodic table of the elements, except it's not periodic. But the point is that given a cognitive atom of this type and a cognitive atom that type, how make they combine or not combined and then as you build up more complex ideas the molecules and up into the minerals, how did those things connect or not connect. And by doing that and - we have to do this manually. Okay, let's take these atoms and connect them together in such a way as we build a model of the common-sense world.

Now, there you use a loaded term: common sense. Everyone likes to use that word. We all feel like we know what it means. But it doesn't seem possible to break it down. I guess, if it weren't common - the thing that makes it common sense is what makes it impossible to define it rigorously.

No, not at all. I would have to disagree with that. I think we make a big deal about common sense right now, because it's something that machine learning can't do even a little bit. So it's all mysterious because if we think machine learning is on the road to artificial general intelligence, then it's got to be able to do common sense and it can't and machine learning as we understand it, I believe never will. But common sense for the most part is just common knowledge. That is everything things about the everyday world that we all know is true and if you model the things in the common-sense world and model their characteristics and what they can do and what they can't do and you start that as a baseline in your computer, that is common sense. It's really common knowledge.

Well, I think we use the term as a Shibboleth or a catchphrase to distinguish what humans know from what machines know. Two plus two equals four is not common sense; it's axiomatic mathematics.

It's a calculation but if you follow our line of work and our line of reasoning, knowledge is a model. It's not composed of symbols, it's a structure. It's a structure that recapitulates or resembles something in the world such that by looking at the model, you can predict the behavior of the thing, of which it is a model. It's not a symbol. So a picture is kind of a very simple model, I can look at a picture of something and it'll tell me a lot about the real thing, Ssay one of the things that we use in our, when we're trying to explain this, I put the formula for glucose up there, it's, you know, C6H-whatever and it's a bunch of symbols, but you have to know the syntax, you have to know the encoding, conventions and eventually, it will tell you something about it. And then I put a picture of the model of a glucose molecule up there and you take 10 times more stuff about it, you see it that fast. That's why we say a picture's worth 1000 words because you don't have all the encoding, symbolic encoding. So that's why the symbolicists failed. You know they got wrapped around in the symbols, they were stuck in a communications protocol, instead of modeling things directly. So we model them directly. So what does a machine know? By our standpoint, machines don't know anything because they don't have any models, they can calculate, they can manipulate data, they can manipulate information, they can take one type of information and they can turn it to another type of information, oh, basically, by mathematical calculation, but no one has given them the structure until now that we've done, now they have something to think about, now they know something.

And looking at the site for New Sapience, this shows up in the form of, for instance, a companion that you describe as being an agent on a smartphone, that you can communicate with, that does a lot of things that I would like to have. An assistant – can't afford a personal assistant to do all this stuff for me but maybe this would do that kind of thing and it would be affordable and it irresistibly sounded like Samantha in the movie "Her".

I haven't seen that movie...

Scarlett Johansson off screen voicing an AI assistant on a smartphone, very thoughtfully done, so I can't ask you if that's what it's like but maybe you can describe it.

I know enough about those things. You know the difference is and I think in the movie, the man had formed a romantic attachment and, of course, people do form romantic attachments to chatbots even though they know it's an illusion, they know it doesn't have comprehension, it doesn't have emotions, it doesn't understand a word they say, it does not understand a word it says they know this; but such is the power of human theory of mind you know that we throw it away anyway because it feels good. In the case of a sapien, that's what we call these things that we've created sense we needed a new common noun for them because each instance of the software as it learns and extends its knowledge it's born with and of course, in our case, it's born already, kind of with a common sense world model on the order of a five year old, four year old or at least it will when we get it ready to go to market, we're getting close to that, it's comprehending language in the same sense we comprehend language, which is completely alien and different to what, say, GPT3 does or machine learnings do when they process language. And thankfully, there's still a little rigor with language left in the field that they say GPT-3 it generates a text as opposed to talking but then they sometimes say talking but it generates text. It's not communicating because it has nothing to say.

The term was "stochastic parrot" from Timnit Gebru, which is pretty good.

Yes, that's exactly what those things are. Their parrots I use that same term, all the time and but what a Sapiens does, it has a model of ideas that are interrelated that are implemented in software, you know, they intuitively, a structure, a graph-like structure, we can think of it as a graph structure, it's not composed of symbols, it's a series of relationships and nodes, if you will and relationships. But when you give it words, its words will have a reference to someplace in that model as an entry point. So, as you communicate, it goes through the same communications protocol that humans do and I think what that is, so you have an idea, you want to explain what a parrot is and to somebody who doesn't know what a parrot is and you say, well, it's a colorful green bird with a big yellow beak and they can mimic human speech. So in order to tell them that you had to take that your idea and ending your parrot and you had to break it down into kind of simple component ideas, it's a bird, a bird has feathers so you describe it. So you encode all that in just a few words and you embrace the new grammar, which is say, look up these references in your mind to these and then put the ones if you've got references to them, then bring those references in and decode the grammar, which just kind of telling you how to put them together to create this idea of a parrot, based on ideas that you already had, it's now put together. So you already have to have knowledge in the machine on both sides of the communications question to have communication. So GPT-3 isn't communicating. It's not using language in the same sense but a sapiens is doing exactly that same thing.

What is the tall tent pole, the challenge in constructing this? In GPT-3 the cost is in training. You've got to feed it terabytes of data, it costs millions of dollars to train it in electricity and so in your system, what is it to the effort, where does that go?

Well, the original effort of the past 15 years was the philosophical breakthroughs in design, creating a science of knowledge as it were and figuring out what the atoms of thought were and what the connection properties were and that was the hard part. Now, it's much less, it's building the model itself, it's fairly intricate. So that has to be done by hand to give it that initial core. It's like when fitting DNA and we're putting DNA together because it's this core, this what we call the cognitive core, which is knowledge about knowledge at the root of all this has to be handled and then on top of that, you build a model of the common sense world as your starting point. And the utility of this model, this knowledge that we're pretty good because all models have a utility that they're designed to support. So what we call the common sense world model, is to support the comprehension of everyday language and so once you can support comprehension of everyday language, you automatically have a very powerful tool that you can tell it things that it will remember them for you and then you can build on that to give it more and more knowledge and give it expert knowledge and more and more intelligence in terms of not one giant master algorithms, but lots of little algorithms, which can be again, discerned by how we see ourselves solving problems. So how big is that common sense model has to be? Much less than you might think because there was a rough correspondence with the everyday concepts or everyday ideas that we need help with, right? As we go through our lives, like that you'd want a personal assistant to give you. So there's a rough correspondence, it's probably closer than an order of magnitude. But it is rather rough between everyday vocabulary words because they point to ideas in the model or entry points in the model that are in common use. But it turns out and you've probably heard this statistic, that about 70% of all the words on the internet are the same 1000 words, you get to 2000 words, you get to 80% of all the words of the internet, they're all the same 2000 words. Then of course, it goes up eventually asymptotically; English has about a million words, but we don't need them for everyday common sense interactions and transactions. So our world model today that common sense for a bottle underneath of it the cognitive core there's maybe 100 of these " atoms of thought" arranged in a careful architecture and then we've got about 3000 ideas that built together in this model and they're all connected. But what we're doing now which is the tedious part is we're building out those rooms. Imagine it's a skyscraper with 3000 rooms. We've got the skeleton up, we've got the curtain walls up and now we're putting the plumbing and heating electricity in and furnishing to each room each idea and if you interact with a sapiens today, I'll certainly be talking to it goes, wow, it really understood what I was saying and then you'll talk about something else, it'll be surprisingly ignorant because we haven't built out that set of rooms yet. But as we complete that, it goes faster and faster. And here’s the cool thing about it, Peter: expert systems had this flat curve and so does machine learning, has that flat curve, you try to get better and better and you need orders of magnitude more, more training, more data. They both had that unfortunate, reverse curve of performance versus effort or resources. We’re just the opposite. The better we’ve got something modeled it’s slow at first and then it gets faster and faster and you know why that is?

Network effect?

Well, yes, but it’s like a jigsaw puzzle. I imagine when you’re putting the other pieces, it’s really slow at first, but at the more of it, you see the fracture goes because our individual ideas are just like pieces of a jigsaw only certain ones can fit into certain ones, right?

Putting the last 10 pieces in a jigsaw goes much faster than putting the first 10 pieces in. What do you use as a figure of merit for measuring how well it’s doing to know how good it is and how good you want it to be?

Well, we actually, it’s rough, but it’s quite useful, we use something called Bloom’s taxonomy of learning, did you know what Bloom’s Taxonomy is?

Tell us.

Well, it’s something that educators do to assess the cognitive or the comprehension skills of human students. You know starting early on, they’re happy if they can just kind of process language to the point where you can tell them, give them a simple statement, and they can answer a simple question about it. Okay, what is a cat? A cat is a mammal. What is a cat? It’s a mammal. That’s level one, they sometimes call that memory, or rote learning, as it were, maybe not a deep level of understanding. But it goes up to more levels. The next level is being able to translate from language to ideas, so that you understand what something means, regardless of how it was said, That’s level two. Level three is to apply what you’ve learned that way and its ramifications to what you already know about the world, that’s level three. Right now we’re definitely on two and into level three.

And the place that today’s AIs like GPT3 fall down is usually in somewhere called - the example is called a Winograd schema, where it requires actual understanding of the terms to be able to answer the question, you can’t do it with the stochastic parroting that they’re doing so just pulling an example from Wikipedia I got here, example “the city councilman refused the demonstrators a permit because they feared violence” and then the question is, “who is the they in there?” and does it refer to the city councilman or the demonstrators it requires understanding to know which one? Have you use the Winograd schemas for benchmarking what you’re doing?

We have not, well, they are difficult in a way and that you need to have quite a bit of detailed knowledge and you have to have it in your model and you have to be able to also have the conversational context modeled in many cases. So it’s kind of the hardest things to do in terms of comprehending language. It’s really comes down to finding the reference to common nouns or pronouns, which one are you talking about. And that’s one of the hardest things to do in language comprehension and to do it based on knowledge you have to have it built up pretty well. So the answer is there, we know how to do it, but we won't be demonstrating Winograd schemas commonly until we get our model further built out down the road. I mean, we could concentrate on building up some of that in certain areas, but we're really aren't that interested in having a Winograd schema test, we're interested in having the broad foundation, that that kind of knowledge-based reasoning and that knowledge, applying knowledge to language and probably before along with that, we have a big task ahead of us to apply knowledge to linguistic analysis. Right now, we are using a machine learning based parser provided by Google, which it guesses what the grammar and syntax of the sentence ought to be. And it does okay but it does things like all of these statistical tools, it does things that are just stupid, it comes out with a language parse and we say what, that's not a now that's never now but, you know...

Time flies like an arrow, fruit flies like a banana...

Somebody asked if we would tell us a sapiens that it was very amusing so we said, sure, I'll type it in and because it's more sophisticated now, but the first thing, it said, I have a problem with what you said time can't fly and then the next minute it went into a processing loop...

Well, some humans find themselves in the same situation.

I mean, the notion that time can't fly was a good one, I've recently been went back to common sense but what knew about what the word *fly* meant, it knew that that wasn't appropriate to time because we haven't gotten to the level that we started adding the metaphor...

Well, at least if you're at the level where you can now explain metaphorical language, then...

It will be able to do that, once you get the literal understanding down, then you can recognize a metaphor or you can recognize humor. That's one of the things that's gonna be fun about sapiens, when you go down the road is everybody's always believe that AIs won't get jokes, well, actually, ours will; because most humor is based on a violation of some epistemological rule or if some violation of common sense. A snail walked into a bar - snails don't walk into bars, oh, this must be a joke. I recognize the pattern that'd be a joke, all right, go ahead.

That's amazing stuff here. We're actually running out of time. So I want to ask you to tell our listeners, how they can find out more about what you're doing, get in touch with you or follow your work and discover more about sapiens.

Thank you, I appreciate that, Peter, the best place is to learn specifically about our approach and our technology is on our corporate website, newsapience.com and I also have a blog site personally, which I talked about the future and called forwardtothefuture.com which I would, for instance, one of the things we've been in recently is we have something called the AI Hypo Meter because there's so much hype going on that about AI these days that somebody has to say the emperor has no clothes. But anyway, a lot of good material on that and then we have a website, we've recently put up specifically aimed at telling you about our first product, the companion sapiens, which you can imagine to be something like, what you would like Siri to be someday, but isn't when you have a digital personal companion that actually understands what you're telling it?

Well, I will be first in line or one of the first in line for that. So appreciate you letting me know about that telling all our listeners about that and really looking forward to seeing where you go with this. So Bryant Cruse, thank you very much for coming on AI and You.

Peter, it's been a pleasure and thank you so much.

That's the end of the interview. I was particularly interested in how Bryant got into this by looking at people in an operations room basically all acting like little cogs in a machine, and thinking that this wasn't the best use of humans' unique abilities. And also that so much energy is going into cracking the artificial general intelligence problem.

In today's news ripped from the headlines about AI, a team at DeepMind challenged the accepted wisdom that transformer language models need to get bigger, as in number of parameters, to get better. And there are some huge ones right now, Google just came out with one with 540 billion parameters, which is three times the size of GPT-3, but the DeepMind researchers built a small model, called RETRO, with 7 billion parameters, and showed that it performed as well as a larger model, one called Gopher, which had 280 billion parameters. It doesn't need as much training and it has an optimization for improving its responses by looking through a database of 2 trillion text tokens for similar language. Sort of like cribbing from real examples. It'll be interesting to see how this sort of research plays out and how it compares to what New Sapience is doing.

In next week's episode, I'll be talking with Dr. Richard Ahlfeld, a PhD in Aerospace Engineering and Data Science and founder and CEO of Monolith AI, for improving the efficiency of engineering decisions. That's 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