AI and You

Transcript

Guest: Olav Krigolson

Episode 65

First Aired: Monday, September 13, 2021

Welcome to episode 65! For those of you just joining us, I'm Peter Scott, and I love to help people understand AI. Starting with myself, because it is a bottomless topic. I will never be done with asking questions about AI and finding out the answers, and I suspect that it's going to evolve faster than I can catch up with it. And maybe I'll need AI eventually to help me understand it. Less than half kidding about that.

Today's interview is going to be a treat because you know that neuroscience is near and dear to our hearts here at AI and You, and the main reason for that is that so far we have not created artificial general intelligence, but our advances in AI in that direction, primitive as they may be, have mostly been along the lines of emulating how we think the human brain works. So the better we get at understanding what's on top of our necks, the better we might be able to make AI that can do general kinds of tasks.

Professor Olav Krigolson is a neuroscientist – his Twitter handle is @thatneurosciguy – at the University of Victoria, where he runs the Theoretical and Applied Neuroscience Laboratory, most commonly known as the "Krigolson Lab." He has developed a unique mobile electroencephalography (EEG to you and me) system to investigate what's happening in our brains when we're tired, stressed, oxygen-deprived, struggling with dementia, concussed—or on Mars, where it might be all at the same time. And he's done that at the Hawaii Space Exploration Analog and Simulation HI-SEAS, which unlike most of the rest of Hawaii is a bleak, barren landscape on Mauna Loa where the altitude and the rocks and the general lack of living things conspires to create an environment somewhat like Mars although obviously with much more atmosphere and much warmer. But it's a place where people can get together and be relatively undisturbed while they simulate what it would be like to live in in a Mars habitat. And Olav has used EEGs and analyses of their brainwaves to look at how people in that environment fared under conditions of stress, etc. So let's get right to the interview with Olav Krigolson.

Olav, welcome to Artificial Intelligence and You.

Hi, Peter. Well, thanks for having me.

Well, neuroscience is near and dear to the hearts of AI people as I was explaining in the introduction, and yet, when I was in high school, neuroscience wasn't on the menu. We were basic chemistry, biology, physics; neuroscience would be something that you'd have to decide to specialize in later on. How did your journey from, say, high school, take you into that field?

That's a great question. One of my claims to fame as a neuroscientist is, I never took high school biology. I actually did physics and chemistry. The first biology class I hit would be called a neuroscience class. So I was in grad school, and I was looking through the elective courses that were available to me for the program. I was actually going to do a master's degree in kinesiology. And one of them was a neuroscience course. And I took the course, and I was absolutely fascinated. I was interested in learning how people learn and how people make decisions. And here was this course that was telling me about what was going on in the brain when people learned and what was going on in the brain when people made decisions. I was just fascinated, it hooked me. So that led me to a PhD in neuroscience and a postdoctoral fellow in neuroscience.

Great story, and have you put your own brain under the microscope in the course of your work?

I have multiple times. I've had my brain scanned with fMRI a couple of times and I've had EEG measured in the lab multiple times. And now with the mobile technology we use, I measure my brainwaves almost every day and I've even fired magnetic pulses, transcranial magnetic stimulation into my brain to give myself a temporary brain lesion.

Wow, that sounds like it's carrying the commitment a little far. How do you know that it's temporary? And why were you injuring your brain in that way?

Well, it is temporary, there's a lot of research that was done on transcranial magnetic stimulation or TMS. In fact, it's used now clinically. Certain clinical neuropsychologists will use TMS to treat depression. So you could go to a clinic and it's a viable treatment. But basically, the idea is the magnetic pulse disrupts the brain function briefly, and then things return to normal. And why I did it? I had a fellow colleague that had one in his lab, and he was using it to study if you took out parts of the brain, how that impacts human movement. So he wanted to know if I wanted to have my brain temporarily fried, I guess. So I did. It was quite the experience.

This is your brain; this is your brain on TMS. I know there's no nerve endings in the brain, at least that's the common saying, so I assume it didn't hurt. But did you notice it?

You do notice it. The classic TMS test is they hit Broca's area, which is associated with speech production. So imagine I'm talking to you, and we're having a conversation, and all of a sudden, if you turn on a TMS pulse, a regular pulse to Broca's area, you lose the ability to speak. So your body is generating sound, but what Broca's area specifically does is turn words into a very specific pattern. So what you hear is sort of like a **[mumbling]** because your mouth is moving and sound is coming out but it's not the words you intend to produce. And what's very trippy about it, if you will, is that you were intending-- So classic, you could imagine you count from one to 20, so you go 1, 2, 3, 4, 5. And then they turn on the TMS and you go, **[mumbling]**. And then they turn it off and you go 10, 11, 12, and then they turn it back on and you go, **[mumbling]**. So it's quite an interesting experience because you're trying to speak and you just literally can't.

Wow, I had no idea we were going down this road. Just so we're clear for completeness' sake, people applying to be grad students in your lab: is this something they should expect to be asked for?

No, no, my lab is safe! We don't use any invasive techniques like TMS. I was just visiting a friend who had that particular technology, and I was always curious about it, so he played around with it with me for a bit.

Okay, so no giant yellow warning signs on your lab saying, "Warning. Transcranial magnetic field, keep head out."

No. Well, we do have one because we do have some biohazard material but that's related to do with genetic prototyping. So we take swabs from people's mouths to get a genetic expression, and that requires the university, at least in their world, would say that's a biohazard. So we have a special room where we have to keep the samples and protocols. But generally, what we do is just put some electrodes on your head and measure your brainwaves.

All right. Now in that, you've got some specialization, and we know from your TEDx talk, that you are able to look at people and tell whether they are too tired on an objective scale, much more useful than asking, "Are you tired?" Can you talk about that research?

Yeah, of course. It really started with the advent of mobile technology. So back in about 2012, I was approached by a Canadian company InteraXon, and they were developing a mobile EEG system. So most people think of EEG with a whole bunch of electrodes on the top of the head, and a hospital or at least a research lab. And they approached me with a little headband that only had four electrodes, and they were curious whether it measured brainwaves or not. So we did a lot of background work to make sure it did measure brainwaves, and lo and behold, it did. So we were quite surprised. But we've all seen virtual reality get cheap and affordable. We've all seen all manner of technologies go this route. So we started asking ourselves a question: "If you can measure someone's brainwaves anywhere in five minutes, what can you do with that?" And we've taken on a bunch of projects, but The Fatigue Project was an offset of that, which is, it's been quite well established in the EEG literature that you can measure fatigue with brainwaves. But what's not as well established--well what was impossible really--was to do it quickly, you know, to do it in a couple of minutes. So we thought, "Can we build a mobile fatigue assessment device? So that's what we did. Our first sort of study with this was just classic basic science, I'd call it. We measured people with the big system, our medical-grade EEG system, and we measured them with the portable system, and we got similar results. When we saw that they retired with our $100,000 system, we saw the same pattern and results with the low-cost mobile system. And then what actually happened is I went to the Consumer Electronics Show in Las Vegas, and I got to give a presentation on the main stage for the Digital Health Summit. And I was talking about the tech and that led to all sorts of doors opening up. So we've done some work in a hospital environment. So we looked at doctors and nurses in the emergency room to see if we could detect fatigue, and lo and behold, we could detect fatigue. And we worked with heavy industry in northern BC; I can't mention the specific company, but they were very interested in fatigue in the workplace, and we were able to deploy the tech there. And I guess the one that we're kind of known for is, I spent a week in the Mars HI-SEAS habitat as a proof of concept for a NASA grant project, where we wanted to show that we could measure fatigue in astronauts in a simulated environment.

Right. And we'll very much want to talk about the NASA experiment in a moment, but I want to backtrack to a Brainwave 101 kind of question because I look at my brain waves with one of the headsets, like you mentioned, and it looks like a spider had a heart attack on the paper. I get that there are patterns in there that you can find, and you can say, "These patterns look like this group. These patterns look like this group," and that that's how you decide someone's tired. But how did they get trained in the first place? How do you know that that pattern means tired and not aroused? Or not bored? Or another state? How did you decide that that general pattern means tired?

Well, this is an area that needs a lot more research, actually. One of the problems with EEG findings, if I'm being honest is that in some instances, there are multiple interpretations for the same pattern. Now, as we move further along, and we build better classifiers -- When I started doing this, we didn't use machine learning approaches, it was all just, we would look at specific bumps in the EEG signal. So a classic one we look at is called alpha. There are neural oscillations between eight and twelve hertz, and generally, when people are tired, you see an increase in alpha power. But the problem is alpha power can increase for any number of reasons. So where it's heading is the classifiers are able to separate patterns that look on the surface to be very similar, but there are subtle differences. So what we do is we look for truth datasets we call them, so a dataset where we know that all things being equal, the only difference between dataset A and dataset B is fatigue. So we match for any number of factors to control for other things - arousal, what you had for lunch, how much you slept last night - and then we build the classifiers on those datasets. And they're not perfect, but it's a step in the right direction. So, as the technology improves the hardware, the software is going to improve as well. So people are constantly improving the mathematics behind the classifiers, but we're also going to constantly improve the data we put through them.

Right. And so you mentioned a study or research that you did about using this technology for potential Mars explorers. Tell us about that.

Yeah, well, the Mars mission is going to be a very unique thing for NASA. So NASA is working under the assumption they're going to Mars, they've been told that by two sequential presidents, so they're planning for a Mars mission. And the biggest problem on the Mars mission is the crew is fairly autonomous. When they went to the moon, the reality was, they were in live communication the whole time, and everything could be troubleshooted from the ground. On a Mars mission, they're going to be either out of touch or at least have severe communication lag for extended periods of time. So it means the crew has to operate independently, to a large extent. And one of the things they're worried about is you can put a heart rate sensor on and know someone's heart rate is higher than normal, and you can measure blood pressure. But what's the gold standard for measuring the brain? How do you know if someone is ready to go outside, say, and do a Mars walk? So there is an open call on the NASA website. They reach out and they have research calls for solutions to brain health and performance. So how do we monitor brain health and performance? So me and a couple of colleagues realized that this mobile tech could be used potentially to do that. So we made a pitch and we ended up with a week in the HI-SEAS Mars habitat on the Big Island, Hawaii. It's used by NASA crews for Mars simulations. It’s used by crews from the European Space Agency working on projects, and we got in there for a week. It’s basically a dome on the side of a volcano. It’s built to the specifications and the size that they think that they would be able to build a housing unit on the surface of Mars. But it doesn’t look much like one. There are two by fours that are visible inside and it’s essentially a large tent in some things, but it does simulate the size and space. And it also, more importantly, simulates the psychology of it in the sense that every time you go outside, you have to wear a spacesuit, your communications are lagged, you live off of freeze-dried food. One of my particular favorites is you get about 60 seconds a day of water for showering, but you’re exercising four hours a day because that’s what astronauts do. So we spent a week in there and the whole goal was to see if we could use our tech to track our brain performance over the week and to see if we could detect fatigue in particular. We’re really focused on cognitive fatigue for this first one.

And are you heading down a road where NASA will end up incorporating EEG sensors into the spacesuit helmets?

That’s exactly where we’re heading. So the follow-up to the initial mission was going to be a longer one-month mission to really get into it. COVID got in the way of that like it got in the way of everything. So we’re going to go back as soon as we’re able to schedule it and get back to Hawaii. And that’s the hope is that that data will then be used to get a larger NASA grant. And then the idea is that what they do is they get lots of submissions for these projects and then much like a science fair, they pick the best one, you know, “This is the tech we want to go with.” But my hope is that the tech and the algorithms we’re developing will end up on the ISS and then hopefully one day on Mars.

NASA is not going to be the only people on Mars. Have you had any talks with SpaceX?

We haven’t yet but it’s on our to-do list is to start exploring who else is interested in the tech. And to be fair, this idea of fatigue assessment, we’ve had a lot of conversations with airlines, heavy industry, trucking, any field where cognitive fatigue might impact decision making or performance, people are kind of interested. One of the differences with our approach is that a lot of the existing approaches rely on continuous monitoring, so the sensors are on the whole time. Our assessments take about two to three minutes. So it’s a case where you put on the device, you do the assessment, and then you’re deemed good to go or not. So if you think of the airline example, the pilot would do the test before they get on the plane, and if they get a pass, then they’re deemed to be not fatigued, and they can fly. If they are fatigued, then potentially that’s a down check and another pilot flies the plane that day.

Sort of like a breathalyzer for the brain.

That’s exactly the way I envision it and think of it, is it’s just a quick brain breathalyzer.

All right. Now, I want to talk about the links between human cognition and computer cognition here to bring in the AI angle some more, because there’s this pervasive notion in AI that we are modeling the human brain, in fact, it’s almost a subconscious assumption. And I’ve questioned that assumption at times. And yet, our models of physical computers look radically different from everything I know about the human brain. The physical computers have a processing unit, a separate memory area, and a way of getting information in and out. They have separate programming and execution phases. Human brain, at least to me, looks like one large lump of jelly where there is no separate memory and processing, no separate programming and execution phases. Horrible oversimplification, I know, but do you have anything to say about perhaps untangling some of those parallels?

I do. And we do some work with this in my lab, as you know, we actually use computational models, a form of AI, a very simple form of AI to make predictions about neural responses. But where we're going is not so much the hardware side, but the software side, which is building neural networks that are designed to mimic the brain. The problem, of course, is that the size of the neural network you would need to capture the entire human brain would be massive. So you know, very simple neural networks are on our phones that learn the pattern of our face and know to unlock the phone. We use them in the lab for very simple things, but to mimic the whole brain is quite a massive undertaking. Now there are groups that are looking at that problem, that are trying to actually build a neural network that would have the same number of neural connections as the brain itself. I don't think we're there yet, but we've seen the power of these things. Now, where it gets interesting is how accurately it compares to the human brain. I call it biological plausibility. I actually supervised a PhD student a couple years back, who was in computer science who had come up with a computational model that learned to perform a task. And it learned a lot better than a human. He'd made some changes to the model and this thing could learn extremely quickly to perform this task. But my question to the student was, "Well, this is an impressive technological programming feat, but it doesn't tell us anything about the brain because the brain can't learn that quickly, we know that. Humans doing the same thing can't do what your algorithm can do." So the big question on the neuroscience side is - and there is an entire field called computational neuroscience – the big thing is always biological plausibility. How well does your model align with what happens in the brain? And there are people that do that. Definitely, the neuroscientists that are using computational or AI simulations at the lowest level, so single neurons, have put a lot of work into making sure that their model holds up to what you actually see in terms of a neuron firing, if you will. Where it gets trickier, of course, is how does that scale up? And as I mentioned earlier, current computers can't handle enough simulated neurons to capture the entire brain.

But even at the lowest level, there is something called the OpenWorm project, I've heard of, which is attempting to simulate the neuronal structure of a nematode worm, which has 302 neurons. And the last I heard they were still not connecting that to anything that demonstrated the behavior of a nematode worm, which doesn't bode well for going to larger scale organisms.

No, no, it's a challenging problem. Some things do work, like one of the things we do is, is when people learn, one of the ways they learn is through what I would call a prediction error. It's the difference between your expectation of what's going to happen and the actual outcome. So imagine you hit a tennis ball, you have an expectation of what's going to happen, and then there's what actually happens. And if there's a difference, we call that a prediction error. It's a key principle in reinforcement learning. Now, it turns out that if you measure with EEG the response to feedback in a simple gambling task, so you get some people to play a little gambling game, the pattern of the EEG results parallels the pattern you see with the computational model that uses prediction errors. So when the model predicts you should see a large prediction error, given the outcome, then we see a large neural response. And when the model predicts we should see a small prediction error, given the results, we see a small neural signal. So we can come up with these sort of computational models that parallel what we see, they're not perfect, but at least the pattern is there. So I do think we're on the road to something. And in fact, this prediction error thing is big in our lives. If we have a robot vacuum, I bought a Roomba recently and that's literally the key algorithm in these things is when you bump into something, that's a problem, so change your behavior and do something else.

You mentioned prediction there several times and I've pulled a quote off one of your papers, or websites, where you said when people make emotional decisions, you get a lot of activity in emotional parts of the brain and that correlates with irrational decisions, while a more active prefrontal cortex is associated with more logical rational choices. You can actually use this in a predictive fashion to predict what?

I'll give you an example. There's a classic game that economists like to play, it's called the ultimatum game. And the ultimatum game works very, very simply. I have $100, and my only job is to share some of that money with you, Peter, so I have to make a decision. You know, I'll give you $50 and I keep $50. So we walk away with whatever the share is, assuming you say yes, your only job is to say yes or no. And one of the key rules to the ultimatum game is you only play one time. So it's your one crack at this. So the original research was done by a group from Princeton, in Grand Central Station in New York, and they would just pull two strangers together, make sure they didn't know each other, and then they'd give them the money, and the offer was made. And if it was yes, you took your money and walked away. And if it was no, then no one gets any money. That's the other feature is if you say no, no one gets anything. So most people do 50/50 but there are people that say, well, they think, "Well, hey, I'm going to keep 90, you take 10." Now, logically, you should take $10 because $10 is better than nothing. But what does human behavior show us? About 70% of people say no to that offer, they go, "No, I'm not doing this." Now, if you were going to play more than once, that's actually a rational choice because you're basically saying, "Well, if you screw me over like this, I'm going to keep doing this and no one's going to win." So you're sending a message, you know, treat me fairly or else. But if you only play one time, you should take the money. Now, the neuroscience of this is interesting. If you do this in an fMRI scanner, the people that accept the $10 have more activity in the prefrontal cortex than they do in the emotional parts of the brain. And the people that reject the $10 and reject the offer and get zero, they have more activity in the emotional part of their brain than they do in the prefrontal cortex.

They're going, "It's worth $10 to teach this guy to not be a jerk."

Exactly. So what you can see is this pattern of results where the brain activity predicts what someone's going to do. And of course, the brain activity is in advance of them actually responding, it's not after their response, it's before.

So you can see from the brain scan what they're going to do before they say it.

Yeah, yeah, it's a very interesting field. There's a lot of very interesting work that's being done with real-time fMRI. People are using it for image prediction. So they show you a picture of a house or a car and they know what you're looking at.

Well, you can go straight from this to more controversial things like lie detection, someone strapped to an interrogation chair and has one of these on their head. And now you said it needs an fMRI, so they've got to be lying on the table.

There's work being done with EEG in this area as well, but the lie detection has already been done. There's a published study with they were showing people pictures of a room and either you'd been in the room or you hadn't been in the room. And they could tell whether you'd been in the room or not. Now where I'm going with this is you could imagine taking this to a court of law, and you show a picture of the murder scene to the potential murderer, and their brain is saying, "Yes, I've been in the room." Well, the idea is, that would mean you're guilty.

Right. And I remember a fiction book by Daniel Suarez where a guy was kidnapped by some people with advanced tech. And the point was that this wasn't that advanced, but they wanted information out of him. And they put him in an fMRI scanner and asked him for his password and said, "Is the first letter A? Is the first letter B?" It was just a machine doing this and it noted when the scan spikes, "Okay, now is the second letter this?" and from what you're saying, we could do that now.

It's close. The biggest problem is accuracy rates. To publish a scientific paper, your accuracy rate doesn't have to be 100%. It's got to be about 70% or lower, possibly. But it's getting there. And we actually just published a paper with a colleague of mine, Alona Fyshe, she's at the University of Alberta, where with EEG we were able to determine what words you were seeing on the screen. So we'd put up a series of words. And we were able to separate that using, again, a machine learning approach. But there was enough information in the EEG signal to tell you that you were looking at the word turtle versus the word dove, or something like this.

That's the end of part one of the interview. Part two will be next week. At the time I'm recording this, it's the centennial of the birth of Gene Roddenberry, the creator of *Star Trek*, and NASA just aired a panel on diversity with his son, and George Takei, who was Sulu. And with me being a Trekkie, I was very much into that. But you may ask, "What's that got to do with AI?" Well, if you've listened to much of this series, you know that there's a lot of *Star Trek* connections to be made with AI. So I thought I'd mark that Centennial by bringing up a few of them. And this is getting further out into the future here in the speculation, but just because it's a long way off doesn't mean that there aren't some things that we can predict fairly accurately. We know that one day, the Sun will turn into a red giant; we know that with complete certainty. But long before then, we will have to solve the problem of how we get along with AI.

Once it becomes superintelligent, of course, there are projections that it might wipe us out. But some of the more optimistic ones include it taking care of us like pets, like in, say, a *WALL-E* movie, where we are just laying back on our loungers while everything is taken care of for us. Now, if you're Homer Simpson, that sounds pretty good, but for most of the rest of us, it's rather horrifying. And yet, it is hard to come up with a plausible scenario that is better than that once artificial intelligence becomes superintelligent. What are we going to do when AI can take care of everything for us? And I like to think that *Star Trek* shows us the best, most fully fleshed-out, most relatable example of what that best alternative could be. That we're learning, growing, doing good things in the universe. And that most of all, we had earned the right to be there doing those things because *Star Trek* showed us having evolved to be a humanity that valued life in all its infinite diversities and infinite combinations. That very much contrasted with where we were when the series was conceived and where we are now. You could trust the universe and the species in it to Kirk and Picard and Spock and McCoy and Scotty; but we haven't gotten there yet. Only in a few places, not enough.

Now, yes, you can say that's not exactly an accurate future because, for one thing, it doesn't include artificial superintelligences at nearly the scale that is 99% inevitable by that point. So yes, even in the most optimistic scenario, it's going to look different because of AI. But the point is, we need to be the kind of people that can be trusted out there. One of the ideas about our future with superintelligent AI is that it might choose to keep us in a simulation as a way of satisfying our desires while keeping us safe. But what if it were not just keeping us safe from ourselves, but keeping the rest of the universe safe from us? And if you really want to go full tilt tinfoil hat, then you could say, well, what if that has already happened, and we are in a simulation right now because the AI is keeping us there until we've demonstrated that it's safe to let us out? No, I don't really think that. Or, at least if you think that you're trapped in a perfect simulation, then your actions and beliefs should be identical to what they would be if it were reality. It only makes a difference if you think you can escape. But if you do think we're in a simulation, then maybe you should work on helping the human race evolve to the point where it's mature enough to be trusted around other species, because maybe that's the criterion for being let out.

Now all of this talk about *Star Trek*-like futures is obviously way off in the future, but we have a space program right now. And periodically, we hear these same ludicrous arguments about "why are we spending all this money on space when we need it on earth?" And firstly, just because it's concentrated in a few spectacular rocket launches, and other things, that doesn't mean that it's a lot of money. The United States has spent less on NASA throughout its entire lifetime since the program was started than it spends every year on the military. The NASA budget is less than twice what Americans spend every year on snacks and supplies for Super Bowl parties and you don't hear them saying maybe we should cut back on that and give it to the poor instead. There is more to our existence than just taking care of our most pressing problems and nothing else. Otherwise, why climb Everest? Why write a symphony? Why hold the Olympic Games? And many of us don't do any of those things but we don't begrudge the people who do. And the same is true of space. I suspect I'm preaching to the choir here, of course. But even though every dollar invested in space development is returned seven-fold in advantages from developing technology, because virtually everything that is done in space is the first time, a one-off of some kind, not like building the same bomb over and over again to blow it up, some people still think it's frivolous. And yes, there are more reasons to go into space than just return on investment in high technology, because if you wanted that, you could probably find other ways of doing it. We do this because we need to become a multi-planetary species, we need to be independent of the Earth. And that sounds so very science-fictiony because it's so very far off, but if we don't start, we'll never get there. This is where we need to start.

And you have to have a good imagination to see where it might lead us because at the moment, we're restricted to trying to build a few domes on Mars, maybe an outpost on the Moon; neither of which are very appealing places to live on. Sometimes I wish we'd grown up in a globular cluster where the stars are so much closer together, and there would be a chance that an interesting planet would be a lot closer to us. I've read that although it was thought until recently that there was a small likelihood of there being inhabitable planets in a globular cluster, that is no longer thought to be the case with some new observations. So we're stuck with the closest star being over four light-years away. It's a pity, it would have given us a lot more incentive to develop space travel that could go further if we only had a third of a light-year to go to the next solar system. But some people think that the goal of establishing an independent colony on Mars, for instance, isn't a good idea. And they say that we should instead focus on putting that money -- again, with the argument about the money, I want to bring up that meme that says, "Why not do both?" -- but to put that money into keeping Earth habitable. And they think that as long as we have another colony somewhere else, then that removes the incentive for us to take care of the Earth. This is a silly argument.

Suppose you went to your bank, and you got to chatting with the IT officer there. And he said, "You know, we don't keep backups of your data. And we're not into this cloud thing either. There's just one copy of it - the live copy that's in our central mainframe, but we think that because of that, it incentivizes us to take really good care of that data. It's only one place that we have to look now, and so we're going to take extra special care of that one place to protect it much better than would happen if we were spreading our resources." Are you going to keep your money there? No. The next words out of your mouth are going to be, "I'd like to make a withdrawal, all of it."

So if it's not true for your bank account, why should it be true for the human race? And the people who make that argument about needing that motivation, aren't thinking about *who* needs to be motivated. It's not them that needs the extra motivation, they've already got enough. They're thinking about other people that need to be motivated to take better care of the Earth. But those people don't look at anything that they can't see beyond a year out, so it won't make any difference to them. But let's say that you did take all of the money that we're spending on space and spent that on the environment and mitigating global warming and pollution and every other kind of man-made damage to the earth. You could do all that only to have the human race wiped out by something that wasn't even our fault, like a massive coronal mass ejection, or the explosion of the Yellowstone supervolcano, or the reversal of the Earth's magnetic field. There are several ways that the human race could be wiped out on Earth by something that we have no ability to control. But there are very many fewer things that could do that if we were also on Mars. So again, it only makes sense to have a backup location for the human race and God bless the people who are looking forward to living on Mars because it's going to be hard, and it's going to be rather monochromatic, but we need them.

So today's space program is the very first steps on a very long road towards a possible *Star Trek*-like future. And I thank Gene Roddenberry who would have been 100 years old now, for giving us that vision. I actually went to his memorial service, 1991, Forest Lawn. No, I was not one of the Trekkies that showed up in costume. I wore a suit: It's a memorial service. But it was very tasteful, the service itself. There was a bagpiper because he was an ex-police officer and there was a Missing Man flyover with Beechcraft Starship experimental light planes of a rather radical design. Because at least as far as I'm concerned, the *Star Trek* that he and of course so many others gave us was the clearest goal for what we need to do to coexist successfully with advanced AI, however many years in the future that happens.

In today's news ripped from the headlines about AI, Google has introduced a conversational search engine called LaMDA. There's no "B" in there. It stands for the "Language Model for Dialogue Applications", and it is, as they say, their first steps in taking search from a stateless to a stateful conversation, where you ask a question, it gives an answer and asks other questions to clarify what you want. That, of course, is huge. And they did that by training LaMDA on dialogue rather than single question-answer exchanges. Of course, it's using transformer-based language models as just about anything to do with NLP does these days.

Next week, we'll conclude the interview with Olav Krigolson when we'll be talking about many more things about neuroscience, including the structure of the human brain and how it's organized. 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