22.12.10 Real Robots Real Life Hour

NS Nathan Sprague
SM Sarah McConnel
LF Lauren Francis
WL William Little
DY Daphne Yao
AS Amarda Shehu

Speaker 1
Both of these are up. But we should both be when we're when we're driving.

NS
Okay.

Speaker 1
Yep, power on power on. Down disables the remote kill.

NS
Is this the remote kill?

Speaker 1
That is the rear kill.

SM
There's a new robot in town.

NS
Maybe if we're lucky, we'll see some of the food delivery robots on campus. So that's awesome. Yeah, that's one of our challenges is we have to always be interacting with these other autonomous vehicles, which they don't stay out of our way. So we have to stay out of their way.

SM
James Madison University's X Lab has created an automated driverless golf cart. They call it the JACart. They're hoping that one day a whole fleet of driverless golf cart taxis will operate in closed retirement communities.

NS
Normally, there's a satellite view, but that doesn't seem to be working. So here we are, the little yellow circle indicates the location of the cart. And these are all locations around campus. So this is labeled the mall and this is labeled the cafe. There aren't actually malls and cafes there. Those are sort of fake destinations that we've set up to sort of simulate the environment of maybe some sort of retirement community.

SM
So for now, the JACart lives on the James Madison University campus. With Good Reason Associate Producer Lauren Francis went along for a ride.

LF
This was really cool. That was like an amazing moment just then too. Like you had students skateboarding beside us, walk towards us, and then students walking really closely beside us. And nobody seemed to panic about this autonomous machine, and the machine didn't panic about the humans.

NS
Nope, so we haven't hit any humans yet. And I'm confident that that's going to remain true.

SM
So is this the future? from Virginia Humanities this is With Good Reason. I'm Sarah McConnell. Today, will humans matter in the future?

LF
I'm having a lot of fun, Nathan, we've almost hit a wall, we've almost hit a truck.

NS
This is science, this is what it is all about.

SM
Nathan Sprague is a professor of computer science at James Madison University. He's part of the team creating and improving on the driverless golf cart. And he's an expert in automation. Nathan, why did you start trying to make a self driving golf cart in the first place? Who did you imagine would be using a golf cart?

NS
So I have to say when we started out, we didn't have the end user in mind. It was more a matter of can we do this? And it started as a course. So, we had noticed... so this was back in 2018. So you know, the whole self driving vehicle movement was really starting to take off. And we noticed that there were some companies popping up in Virginia, where we're located, that were hiring people in the field of self driving automotives. So we thought, well, this could be a good opportunity for our students. So we just dove in with about 15 students and said, "let's see how far we can get."

SM
You know, I saw a story in today's Wall Street Journal that said investors are getting impatient with how long it's taking the car companies, the big ones, to deliver fully autonomous driving cars.

Elon Musk said he'd have one that could go from LA to New York by the end of 2017. That never happened. Do you now appreciate why this is so hard?

NS
Yeah. If you pulled people aside in 2018, or 2016, engineers who are working on the problem, and ask them if they thought that we would be now where the business people were saying that we would be, there would have been doubts about how possible that is.

SM
So tell me where you are, and the students are, on the golf cart that they're retrofitting to be an autonomous driving vehicle.

NS
At this point, the golf cart knows where it is on campus. It has predetermined routes that we've given it access to. So we've sat down and looked at a map and said, this is a reasonable place for the cart to be able to go, this is a reasonable place for the cart to be able to go. And there's a user interface that allows the passenger to select destinations. And then as long as that destination is along one of those routes, the cart will navigate itself to that location. It's also able to pay attention to the route ahead of it so that if, for example, a pedestrian walks in front of the cart, the cart will stop. If there's another vehicle moving in front of it, the cart will slow down so that they're moving at a reasonable pace together. And that's that's where we are.

SM
What do you still need to do?

NS
What we're involved in now is really asking those questions about, now that we have this autonomous system, what questions do we need to answer to be able to make it a system that my grandparents could sit down in and be able to use in a way that makes sense to them, that is comfortable for them? The biggest questions that we're interested in are how should the person interact with the cart? Should it talk to them? Should it pay attention to what the passenger is paying attention to? We have a pretty good handle on the level of autonomy that we're interested in. So the cart actually operates pretty reliably.

SM
First, this was a class, can we make an autonomous driving golf cart? And then you started meeting more funds, got some grants from some outside organizations that were interested in the project. Who is interested? Who thinks you might help move the needle a little here?

NS

I mean, in terms of who's who has ponied up the money, the two funding agencies that we've gotten, we've gotten grant funding from is The Jeffress Trust. So this is an organization that supports interdisciplinary research, often research related to health problems. So the real motivation for us in terms of the application that we envision is having an autonomous taxi system that's going to operate in like a retirement community, a closed retirement community. So that's an environment where the people living there often have mobility issues, they often can't drive. So just being able to get around one of those campuses, without asking someone, without burdening someone else with giving you a ride, would be a real benefit to those people. So that's the application that we've had in mind. And I think that's the, that was sort of the selling point for that original random grant funding. We've also gotten funding from the State of Virginia, there's something called the Commonwealth Cyber Initiative. So they're interested in autonomous systems, they're also interested in security. It's a pretty big project. So you know, my piece is sort of the autonomy, making sure that the cart is able to navigate correctly. We have other people who are focused on the sort of human computer interaction component, and we have other people who are focused on the security component.

SM
Do you think that actually, what we may need to do to have driverless vehicles, is completely revamp roadways? That we reconceive of roadways as limited-access smart corridors, right? As opposed to things where untrustworthy people with dangerous signs might show up?

NS
Yeah, I mean, this is a big debate. So that's, certainly a lot of the problems become easier if you have complete control over the environment, right? If I can instrument the roadway with, to make sure for example, that I always have high speed connectivity for my vehicle so that I'm never sort of out of network range. If I can instrument all of the vehicles so that they can communicate with each other, that makes a lot of the problems, all of a sudden, easier. If I don't have to worry about other vehicles on the road making unexpected maneuvers that my vehicle has to respond to, then all of a sudden the problem gets gets a lot easier. So of course, certainly people are thinking about all of those issues. I think it's been surprising. If you went back 10 years ago, and asked how far can we get without doing those things, right, without carefully instrumenting all the roads, without, you know, having electronic signage that the car is able to interact with, without having the vehicles able to interact with each other, I think people would be impressed at how far you can get just turning an autonomous vehicle, instrumented with sensors, loose on roads that were not designed for that system. And like I said earlier, 95, 98 percent of the time, we can deal with, with the same kind of road environments that a person could deal with. I mean, our cart can't, right? That's one of the advantages of of our of our scenario. We're

operating a low speed vehicle golf cart, right, we're only going a few miles an hour, we have a fixed set of routes on a known closed campus. So a lot of those those edge cases, a lot of that long tail of very unusual circumstances don't come up for us.

SM
You have a very hilly campus, though. And it's possible like if, if there is one misstep, one driving error by this autonomous vehicle, it could go down a steep embankment, right?

NS
Oh, absolutely. And that's something that we talk with our students about a lot. It's dangerous to have a, you know, it's 1000 pound golf cart. It's operating on its own. Like you say it's a hilly campus, it could definitely go off the path and go down the road. If we're not paying attention, if the cart doesn't operating correctly, it could hit someone. It's interacting with pedestrians all the time. So we take the safety considerations pretty seriously and for us, you know, because it's a research vehicle because, you know, we're not trying to get it to that extreme level of reliability that you'd have to have in order to deploy something really autonomously in a, in an unstructured environment, we always have a safety officer on hand with a kill switch.

SM
You know, you mentioned earlier the user interface. And I saw in a promotional video that the students created something a little bit like Google Maps, where there's a list of destinations on campus and you can pick one and go. And I thought that was such a cool idea that you could make a directory with something that we're already familiar with. So in Google Maps, I can get directions to the place I'm driving, and it makes me think in the future, could my autonomous vehicle just take me there instead of tell me how to get there, right?

NS
That's absolutely the vision, and that ease of use, that making it a very comfortable interface for the passenger, we think about that a lot in this project.

SM
Have you had a moment that you loved where the students, together, sort of crowdsourcing among themselves, have come up with a plan you might not have thought of?

NS
Oh, many of those cases. So I can't overstate how much of the sort of real legwork on this project has been done by our undergraduate students. I mentioned that we have a navigation system that allows us to create routes on campus and then for the cart to autonomously do the path planning. I had a student propose a system for for taking

that on, disappeared for a month, didn't know he was doing anything, got worried about him, and he came back, and he had this fully functional system in place. And it's really beautiful the way it operates, you can drive the cart around a new place that it's never been located before. It's basically laying down waypoints along that route. And there's a system in place where we can then go in after the fact and make adjustments to that waypoints. And every new place that you drive with the cart is another place that it's able to navigate to next time you operate the cart.

James Madison University is not primarily a research school. So we're, we don't have PhD students. So the way academic research often works is, the professors are writing grants, the bulk of the actual research work is done by graduate students. And maybe undergraduates will come in and sort of have some involvement in the lab. But here, all of the actual legwork for the research is done by our undergraduate students. And that's, they can do some amazing things, but that does mean that we have to be very willing to sort of reuse pieces that we find in the community. For example, our cart has a system of localizing itself, which is a technically hard problem. So we use LIDAR, which is a laser rangefinding technology that gives good information about where objects are relative to the cart. But the hard technical problem is, how can the cart recognize where it is on the basis of that LIDAR data, recognize that it's been in that same location before and use that for localization? So that's a that's a tough problem, right? So that's a problem that teams of engineers have worked on for years. And there are some very good open source solutions for that, that we were able to take advantage of. So we haven't had to build every component of the system from scratch. And we never could have gotten as far as we have, if we did.

SM
Are you done are still going?

NS
Oh, we're very much still going, though the focus of the project has changed a little bit. You know, the first few years was very much about getting an autonomous system in place. Now we have a reasonably mature autonomous system, and what that has allowed us to do is have a research testbed for asking all of these other questions that we're interested in. So for example, we have a group of collaborators at the University of Alabama, Huntsville. And they're very interested in autonomous vehicles from the psychology point of view, and in particular, they're understand, they're interested in understanding the degree to which humans trust those systems, what causes them to trust those systems. And that's interesting to us, because we're very interested in having the cart be able to respond to our passenger in real time. So if we have a passenger in distress, for example, we'd like the cart to be able to recognize that and call a remote operator and be able to, so that they can intervene.

So one of the studies that we completed recently was to put a passenger in the cart, have them operate the cart around campus while recording heart rate data, and giving them a survey both beforehand and after afterwards about their general level of trust in technology, their level of knowledge about autonomous systems, hoping to kind of tease out what is it about our system that causes people to have faith in that system? What is it about our system that that tends to act against that? So the next round of studies, we're actually introducing a set of intentional defects in the way that the cart is going to operate. So we can intentionally cause the cart to go off the path, overcorrect, come back on the path, so that we can look at well, how does the passenger respond when the cart doesn't behave as it's supposed to?

SM
Nathan, this has been fun. Thanks for talking with me, and good luck on the project.

NS
Thank you. It's been a real pleasure talking to you.

SM
Nathan Sprague is a Professor of Computer Science at James Madison University. A filmmaker brought the late celebrity Chef Anthony Bourdain back from the dead in 2021. In Morgan Neville's documentary "Road Runner," an AI-generated Anthony Bourdain speaks words that he only ever wrote but didn't actually say. William Little is a Professor of Media Studies at the University of Virginia. He says this voice from the grave is the least of our concerns for the future of AI and film.

William, there was controversy over the new documentary about the late celebrity chef Anthony Bourdain. This was right after the film director casually revealed that he had used AI for Anthony Bourdain to read a note that Bourdain had written, but never actually voiced. So this wasn't Bourdain's voice, but sounded like it.

WL
Sure, the the line from the film reads like this, it's Bourdain's voice, quote unquote, speaking and he says, "You are successful, and I am successful. And I'm wondering, are you happy?" And he's speaking again, in quotes, to a friend. And the controversy arose from the fact that Anthony Bourdain didn't actually speak those words, he wrote them to a friend. And then the director of the film took data samples of Bourdain's voice and synthesized them using AI technology, and had Bourdain then in a post mortem fashion speak that line. I have plenty to say about this, I think the the ethics of this matter, they're not for me the most important issue here. I understand why some folks are troubled by that. But I'm more fascinated by the kind of

philosophical, and even religious dimensions of this subject. So I'm happy to speak about that.

SM
So let's get the ethical side out of the way. It would have been worse had he never written the words, and then fake words are programmed to come out of his mouth, right?

WL
Exactly.

SM
So people however, felt betrayed. You actually sort of tricked us about what was coming out of his mouth.

WL
The thing is, documentary film involves manipulation always. So there is no pure, unfiltered, raw, if you will, presentation of reality in a documentary. This, this is a, I suppose you could say a kind of extreme example of that. I also think that there's a kind of investment in the idea that a celebrity, a celebrity's life is one that is to be honored, in a special kind of way, and so to manipulate a celebrity's voice carries all the more sort of impact.

SM
So tell me, what is fascinating you about the use of AI in film that has a more religious or spiritual quality?

WL
I mean, I think this directly relates to this case, because in a way, what the filmmaker did was he resurrected the dead. You know, that, to me is just absolutely fascinating, because that is happening as we speak in many different quarters of the AI technological economy, if you will. So for instance, there's this company called HereAfter AI that is creating what are called digital clones of a loved one, created so that when that person dies, you can then speak to that person in a hologram form, in a digital form, and that person can respond to you. I know this all sounds really wild.

SM
No, I saw, I saw a piece on this using Holocaust survivors on "60 Minutes." So that there is a vast database of things that you could ask them or their hologram after they're gone.

WL
Right.

SM
How did you come across this company, that's actually making a profit doing this, to create these databases of sound with loved ones?

WL
I'll tell you how I stumbled across it. I teach a film titled "Marjorie Prime," which stars Jon Hamm. It was based on a play, and it features this technology in the film. So you see a hologram version of a dearly departed figure who's sitting on a couch in a living room and speaking to the surviving family member and is a source of consolation, and comfort, and joy even. I was drawn to that film because it represents AI in a way that's not exclusively dystopian.

SM
Right. Why are you interested in AI in film? What launched your interest in this, because there are a million examples from dystopian to the sort of hereafter uses of AI?

WL
I wanted to understand how important it is to be able to think about the cultural implications, the ethical implications, the social implications, the political implications. And I thought, well, film is a medium that is already telling stories about this technology. And can I think about those stories in a way that will help me and my students make sense of this just rapidly emerging technology?

SM
You know, they're becoming adults in a world where AI holograms of deceased people and deep fakes are the norm. Can you tell whether they have a different response to AI manipulation in film than even you do?

WL
I'm not sure it's, I'm pleasantly surprised. But I am surprised by how often it is that they are rather surprised by some of these developments. And I think that's because they're happening so quickly, that we're all sort of, you know, in a position where our eyes get rather wide when we hear these reports about new kind of iterations and wrinkles in this technology. I mean, another example is, there's a company called MyHeritage. And they developed a technology called deep nostalgia. I mean, this is all within the last 18 months, maybe two years. So even since I began teaching the course, the company invites you to submit photographs of deceased loved ones, and they convert that photograph into an animated image. So it begins to look more like a short film. So again, that's another version of resurrecting the dead.

SM
Oh, I forgot about that. There are these even old pictures they can take, where somebody in the picture turns to you and winks or gives a little smile, right? And those are charming, charming iterations of AI, wouldn't you say? I want to see my mom and dad turned to me and give me a loving smile.

WL
Yes, there is a charm. There's also what you know, is described as the creepiness factor. It's that, it's that uncanny valley, you know, experience where it's, it's very close to being human, and then on the other hand, it seems mechanical, or off, or alien somehow.

SM
Our greatest fear of AI fakes, you know, it's one thing to say, hey, there's a way to render Tom Cruise where it looks, and sounds, and seems like the real guy, because there's so many portraits of him in the archives that you can make a real AI synthesis. It's another thing if you think political leaders are having words come out of their mouths.

WL
I would imagine folks in the film industry are also worried. I mean, there's this possibility in a, you know, perhaps not too distant future, that all actors that you see on screen are synthesized versions of real people. That that's pretty dystopian. And I think there would have to be some kind of contract that would say something to the effect that the actor has willingly provided all that data to be used, as, you know, an organization or a film company, or what have you, or an advertiser sees fit.

SM
There's any way for us to avoid allowing ourselves to be manipulated in this way?

WL
One of the problems is that AI offers a fantasy of efficiency, and sort of what I call a kind of drag-free experience. And Amazon uses AI to deliver its packages as quickly as possible. Then it's awfully hard to kind of push back against some of that, and to say, no, no, no, I'm concerned about the implications. Yeah, and I mean, another example of this, you were asking earlier about recent iterations of AI technology, there's at least one influencer that I know of who is a, an entirely digital creation. So it's not a real human being. But that figure has tremendous clout. And one of the reasons it has so much clout is that it looks so perfect and behave so perfectly. So it's like a reified version of a real influencer so that that figure just seems to offer a kind of dream life almost.

SM
And why shouldn't influencers all be fake? Right? They're just influencers.

WL
Yeah, and isn't, I mean, isn't influencing culture itself a kind of expression of the promise that you can lead a drag-free life.

SM
And so then by extension, why shouldn't fictional representations in media be AI or AI manipulated? Because aren't we just there for that fictional experience? What we want it to do is convince us and tap into our pleasure zones, right?

WL
Yeah, I mean, the thing is, if it actually worked that way, it doesn't make for very good drama. Because, you know, drama depends on conflict. So, it's, it's much easier to depict AI in a dystopian way. But I do think that that, that seductiveness of the friction-free experience, again, just sort of overlooks the kind of basic, important hardships and struggles of being a human being. So that's what, that's kind of what troubles me about this dream of a sort of drag-free experience with the help of AI.

SM
It depends on who it's being drag-free for, right? What if down the road, actors are holograms? And what directors are doing is manipulating all that to create not drag-free storylines, but storylines with conflict and fascination, et cetera. But using AI more thoroughly.

WL
That's true. Where the drag-free part of it comes in, you can see is the producers, the director, what have you, they don't have to worry about an actor getting tired, an actor contesting the contract or the salary, actor not showing up. None of that drag.

SM
All right, we're gonna look back on Meta's collapse the metaverse's collapse and realize it was actually just the beginning.

WL
Yeah, I mean, this goes back to the issue of touch. I think, in a very simplistic way, that part of the reason the metaverse has not taken hold is that while it may be attractive at times to sort of drop in to virtual realms, there is this longing at the same time, for a kind of tactile human connection. People love going out to restaurants again. I mean, the students love being in person in class again. There's a reason for that.

SM
William Little, thank you for talking with me on With Good Reason.

WL
Thank you.

SM
William Little is a professor of Media Studies at the University of

Virginia. This is With Good Reason. We'll be right back. Welcome back to With Good Reason at Virginia Humanities. Technology advances at the speed of our own biases, even in the medical field. Daphne Yao is a Professor of Computer Science at Virginia Tech. And she's working to help correct some of the bias in AI and medicine. She and her team have developed an AI fairness technique, where they're inputting datasets of underrepresented patient populations to improve the accuracy of things like whether or not a patient will die in the ER within 24 hours, or the likelihood a patient will survive with cancer five years. Daphne, you and your team have found a solution for reducing bias in AI, which you believe will reduce potentially life threatening prediction mistakes for minorities. How so?

DY
Yes. So this is, I used to study cybersecurity, not digital health. You know, health is really the biggest challenge that our humankind is still facing. With all those computational power and abilities, how can we make healthcare management better? And so, so that's really how I get into this. How can I contribute to with my computing knowledge?

SM
You and your team have found a solution for reducing prediction mistakes for minority populations. Help me understand how this happens. What sort of predictions are being made, and how are they not factoring in the needs of minority patients?

DY
So the predictions are, you know, for example, things like what is the likelihood of a patient to surviving cancer in five years? What is the likelihood of a patient die in ER in 24 hours? And so those kinds of decisions, and the prognosis, are very useful information, and in machine learning could be a great tool, a great friend, or a great assistant. Machine learning can do this reasonably well, because it has this amazing ability to process massive amount of past data. If you give them millions of records in the past with the information on this patient, with this kind of blood pressure, has died in the ER in 24 hours, and this patient with this record, this size of a tumor. And so machine learning people have trained models to do this. And the model just need to output a decision given a patient's record. And but the problem, the fairness problem, comes in when we look at in the past that we have been training one machine learning model for the entire population, for the entire nation, regardless of your race, your gender, your age. But in reality, the patient's data that goes into teaching the computer how to make predictions are not well balanced.

SM
Why do you think that happens? Because I would imagine if the people feeding the massive amounts of data into the machine so that they are smarter and can more accurately predict are doing a sample of the

entire nation, you would have all kinds of minority groups, as well as a majority white population, presumably, factoring into that machine learning. Is that not happening?

DY
The limitation of machine learning models, the computer is that if they don't see a lot of data then they don't do as good a job. And so it was told again, and again, this is the type of patient, a white patient, a certain age range is what you have to learn, and you have to learn it really well. And for minority patients, it happens less often. And so the computer may not be able to capture the differences between different groups.

SM
If you get the prediction wrong of the medical outcome for a patient, how does that affect the care they receive?

DY
This is a great question. In the emergency room prediction, where we want to predict whether or not this patient may likely be dead within 24 hours, the wrong prediction may have two consequences: if you say accidentally predict death, but then the patient is not going to die, this patient will survive, then then you are over cautious, which is not that bad. And the other type of mistake machine learning model can make is more serious, where you miss on mortality prediction, you predict that the patient will survive, but then the patient may develop, say, sepsis very quickly, where blood pressure drop very quickly and then you miss that. And so that kind of prediction would be very fatal. And so what we are work, we have tested two different data sets: one year the year data set, the other is a cancer five year of survivability prognosis prediction.

SM
Give me an example of a case, or set of cases, where you know that the standard models have not accurately represented the outcomes for patients in either minority groups, or young versus old.

DY
Just to use a one typical machine learning model. In that case, that you will have a wider variation of this accuracy, this prediction of mortality accuracy, the whole group would have 61% of accuracy. Out of say 100 people, 100 patients who passed within 24 hours of emergency room visit, the model can only accurately predict 61 of them. And so the other 39 patients will not be predicted accurately. They will die in the hospital but the machine learner will not be able to predict that. And so this is the whole group accuracy. But then if you look at the Black patients, the Black patients accuracy is only 50%, which is a big difference from the earlier 61% accuracy, 61 versus 50. And so the problem is that if you don't say tell me machine learning model, what is your accuracy for Black, or for Hispanic, for Asian, for

people under 30 years old? The machine learning model will just tell you, oh, we're predicting 61%.

SM
So very simply, is your team finding, let's say, people under 30, and massively inputting this data into the AI mechanism so that it is much more accurate for reporting health outcomes for young people, and doing the same for Asians, Latinos and African Americans.

DY
Exactly. That's exactly what we have find that to be effective, to be able to improve the prediction. And as one way of improving it we say, okay, the minority patients, we don't have a lot of their records. And so maybe we should create more of those records in the data set, and then train machine learning models, train the computer with those enriched data set. people have been doing this type of enrichment in the past. However, in the past, the enrichment is not differentiated. And people can reach all the disease group patients of Black, white, young patients, as long as they are, they belong to this disease group, that rare events group, we are creating more of those cases. But our finding shows that no, no, that is not very effective. You have to, only in reaching the subgroup or the sub population, that you're particularly interested in, its substantial improvement over the one model fits our prediction.

SM
Are you finding there's excitement once people learn about what your efforts are? Are there people in the field who realize what you're doing and are championing this effort?

DY
Yes, yeah, we got, I would say very, very supportive reviews, which is very encouraging. And I've been sharing our findings with folks in our profession, you know, both computing, and also in medicine, they are just amazed that this is such a intuitive approach. And people haven't really thought about it. This AI fairness problem is relatively new, you know, with the rise of AI five, six years ago, people start to realize that is AI cannot solve all problem. We have to tell AI we have to give, provide guidance to AI.

SM
Well Daphne Yao, thank you so much for sharing your insights on With Good Reason.

DY
Thank you so much for having me.

SM
Daphne Yao is a professor of Computer Science at Virginia Tech. A brain trust at Google recently developed something called the

AlphaFold 2. It's a game changer. It allows scientists to actually see in 3D the structure of molecules they're trying to target to create new medicines. Amarda Shehu is a Professor of Computer Science at George Mason University. She says this structure prediction miracle is really making a difference for the future of science and medicine. Amarda, tell me about this new development in Google, called the AlphaFold 2, which allows scientists to actually see what they're trying to target when they're creating medicine. I've read that it's been called the structure prediction miracle. Is it miraculous?

AS
It's pretty close to that. It's an amazing, a powerful tool in the hands of molecular biologists. We didn't have such a tool for decades, five decades of research. Many scientists like me came out of the labs, writing our dissertations on how to write up code to figure out what these three dimensional shapes of proteins were, just knowing their chemical composition, what were the atoms that comprise them. But this is a leap. This is certainly a leap.

SM
And the leap is what? What can you actually see through AlphaFold 2?

AS
Right, for many, many human disorders, the protein molecules end up often being the target of pharmaceutical companies because more often than not, they are the reason why something is going wrong inside our bodies. We do need to know what they look like in three dimensions, so we can fit it with something with design. So having a view, a model, of the three dimensional structure, is really key for our ability to advance drug discovery, and more generally, treatment of human disorders. So that is the reason why it's so exciting that now we have these tools, such as AlphaFold 2, advancing our ability to predict, to compute, these three dimensional structures of proteins, basically, with a push of a button.

SM
You know, I looked at one of these online, and seeing the structure, the picture of it in sort of 3d form, seemed so simple. And yet, being able to illustrate it could not have been simple, right?

AS
Well, they are, they're very interesting structures, the structures of proteins. That's also what fascinated me when I started as a PhD student. They have so much exquisite complexity to them, because there are a lot of voids inside, there are a lot of crevices on them, that we target or leverage in order to stick, to bind to them small molecules, drug candidates that we design in pharmaceutical companies.

SM
Why did this come from Google? You'd think that something like this,

this 3d picture of a protein that has really been a game changer, you'd think it would have come from Stanford's science lab or another medical facility around the country.

AS
Okay, so now you've made a really, sort of a painful (INAUDIBLE) for academics. So, on October 20, we were quite shocked, I would say, shell shocked, really, as academics, when we learned that the DeepMind team, so DeepMind is sort of like an academic lab, but it is funded by Google, it's part of really, I should say, the Alphabet, the parent company of Google, and it's a lot of researchers there with guidance from Demis Hassabis, and others higher up in the hierarchy, that really go after complex, challenging computational problems. They were first interested in games, such as chess and the Chinese game Go, and then they cast their eyes towards these sort of decade long problems in molecular biology, and protein structure prediction, I suppose, really piques their curiosity.

Every two years, we gather as part of a competition, we call it CAST, Critical Assessment of Structure Prediction, where all the academic labs get together and they effectively compare notes. How well does your method do? Is mine better than yours? Which one is the most accurate? As a way of understanding what is the state of the art and how much have we advanced. And in 2020, the DeepMind team won that competition. They debuted their system that they called AlphaFold 2, and they blew everybody out of the water. It wasn't even much of a competition. They performed much better than everybody on all sort of categories of difficulty, which is a way through which we organize different protein targets.

SM
Is Google's DeepMind a lab some place where all these minds are working together side by side? Or is it various people paid by Google all over the country or maybe even all over the world?

AS
Well, I can speak to the physical collocation, but they are very coherent in their activities. They have an objective, and they collaborate with one another. But I wanted to address your point, why this team, why this particular team. I want to make sure that I give credit to academic labs and academic researchers. There were a lot of us over the years that accumulated a lot of knowledge on what works, and what doesn't work. But what the DeepMind team was able to do is, it was able to take effectively many imperfect methods and make them more powerful by having virtually unlimited computational resources, financial resources, and human resources. So in an academic lab, we are very constrained for how much money we can throw at a problem, right? And how many students we can pull together to work on a problem. And also how much computational time and how many computational resources. We cannot afford to do a lot of prototyping.

So this team was able, because they have access to very unique computational equipment, they were able to prototype, to try many different ideas, to scratch out those that couldn't work very quickly, and come with a design that finally worked. So in a sense it was, what they did was take an academic labs research or many academic lab research, on steroids, they were able to do it in a short amount of time because of this unprecedented access to resources, human resources, financial resources and computational resources. But they really built their ideas over decades of research by academic labs.

SM
Does it give you an idea of what you would like to create? Do you want to sort of replicate that winning model of how to make these leaps and bounds through funding and pulling various programs together?

AS
Right, so you're asking all the difficult questions. This is a conversation that us in academia, now are having, right? Can we, and the question is somewhat a little bit more pessimistic than what you framed. It's more in the sense of, can we even compete, right? We sometimes put it in those terms. Can we compete with teams such as DeepMind, or can we compete with teams now at Meta that are casting their eyes on molecular biology problems? I am not particularly concerned with competing. I'm excited about what they're able to do. I don't believe that we can replicate what they do, because in academic labs, for funding, we go to the federal funding agencies, right? We ask for them to fund our proposals, our grant proposals. And typically, it's not a lot of money, right?

We're also limited because we have an educational mission, we have to educate our students. We cannot just 100% of us focus our time on research or pool our students. They're not employees, right? They are students, we're educating them. So we have, our mission is a little bit more nuanced, a little bit richer. But I am excited about the research and development that is happening at these big tech companies, because they are allowing us to really push in certain problems. But they're also inspiring, through shoots, or sub shoots, or threads of research that previously we would not have been able to do. I can speak for myself, my lab now has a much richer research agenda. Because now that I have access, for instance, to structures, and I have access to models, computational models, by DeepMind, the limit is just my creativity in terms of what I can do with it, and what kinds of deep biological questions or health problems I can start investigating. So it's really a jumpstart, I would say for many academic labs.

SM
You know, you are known as a researcher unafraid to tackle exceptionally challenging problems. You've really embraced them. And also, as someone who is a fantastic teacher. I want to diverge from AI

for just a moment and read something that you wrote in your application to the State Council of Higher Education of Virginia, where you were awarded an Outstanding Teacher award. And this is what you wrote, it's so moved me. "She stood in my office, hair covered in her scarf and posture slumping. In halting word she told me that I didn't know her, but she had been in the Ph. D. program for three years. She had struggled to fit in, though her grades were good. And she'd bounced from lab to lab. Would I take her in? I would be her last chance. It was the fall of 2017. As I listened to her a dozen diverging answers formed in my head, but the only one that came out in a heartbeat. She was startled as if my words were the beginning of a cruel joke. To assure her I followed up with, 'I have the perfect project. Let me give you some reading to get us started. Would you like to meet in a few days?'" How did she fare after that encounter?

AS
So you're talking about my student Nasrin Akhter, and she's, she's not alone in that experience, in the experience of feeling a little bit lost and not finding sort of enough passion or enough guidance, not treated very well. She is actually, and I say "is" though she graduated, she still is one of the most productive students in my lab. She went on to publish many journal papers, many conference papers, she won the Outstanding Graduate Student Award from the computer science department. And she's now teaching her own students. She's now an assistant professor at SUNY Buffalo. And I believe this is her third year. She did a graduate, she defended her PhD dissertation just at the beginning of the pandemic, but that didn't slow her down. So she is, she's now hopefully accepting her own students in her lab. But that is a wonderful experience, I would say. And while when you read it, it might seem that I did a favor to Nasrin, it was not a favor to Nasrin. I like to think she did a favor to me, because it really, it enriched my life, it enriched my lab, not only in terms of intellectual ideas, but in terms of learning, working with different people and really broadening my understanding of who can succeed, and who can push forward, and you know, who can come and enrich my life?

SM
You know, often, young women who are scientists have a hard time finding the mentors who can, who they can model themselves after, as well as see someone succeed in that regard. Your own mother, however, was a neurologist and doctor in her community.

AS
Yes. And she was very busy. She was one of the few that had two specialties. She was also a psychiatrist and a neurologist. So I learned very early on from her how to be effective, how to multitask. She would often take me because there, when she found herself without having childcare, she would take me to the hospital and leave me in the room with the nurses.

SM
Well, Amarda Shehu, thank you for sharing your insights with me on With Good Reason.

AS
Thank you. It's been a pleasure.

SM
Amarda Shehu is a professor at the College of Engineering and Computing at George Mason University. She has been named an outstanding faculty member by the State Council of Higher Education for Virginia. With Good Reason is produced by Virginia Humanities, which acknowledges the Monacan nation, the original people of the land and waters of our home in Charlottesville, Virginia. Our production team is Allison Quantz, Matt Darroch, Lauren Francis and Jamal Millner. Cassandra Deering and Aviva Casto are our interns. Special thanks to Jenny Taylor for booking assistance. For the podcast, go to with good reason radio dot org. I'm Sarah McConnell. Thanks for listening.