AI and You

Transcript

Guest: Kordel France

Episode 72

First Aired: Monday, November 1, 2021

Hi, and welcome to episode 72! Today's guest is Kordel France, an AI engineer and founder of Seekar Technologies, an AI startup. Going over the transcript, I lost count of the number of times I said "Wow," because of how many unexpected topics we got into. I had no idea Kordel was going to start talking about autonomous tractors, and I know… what's the word… *nothing* about AI in farming. So that got my attention immediately. And as improbable as it might sound, that topic segued into AI in defense, although not some kind of weaponized tractor, but more… conventional… missiles. And also, Seekar contributed two products free of charge to physicians to help ease demand on medical staffing and screen for the virus faster. They are currently building the first clinical AI tool used to advise neuropsychologists in diagnosing mental disorders.

We are going to cover a lot of territory in this interview. Let's see if you are as surprised as I was by some of it. Here we go.

All right, Kordel, welcome to *AI and You*.

Peter, thank you for having me.

Sure. Now, tell us something about how you got into this field, because it's so new, just about everyone in AI was doing something else before that. What brought you into this field? And was that accidental? Was it happenstance? Was it deliberate? Give us the story.

Sure, so it was more or less deliberate for me. It actually goes back to when I was a toddler. So I grew up on a farm. And on the farm, my father had some autonomous driving software for these tractors, these farm implements. And being able to see these really large machines steer themselves without having to have human interaction as a kid was pretty profound for me and was, at the time, something straight out of a sci-fi movie, sci-fi novel. And that really kind of catapulted my interest into robotics and autonomous systems. And my father always told me, "If you want to go into that, you need to really excel in mathematics," which I did. And I've always really found joy in finding order within chaos and math obviously helps a lot with that. And pattern recognition, obviously, goes hand in hand with that. And so I continued to study that path throughout high school, college, and ended up starting a career in engineering and autonomous systems. And from there just kind of progressively got more and more heavily involved with artificial intelligence. Before I started Seekar with my co-founder, I was actually in autonomous systems for a very large defense company within the United States. So able to see some pretty profound things on how autonomy can be used in defense, and the limits of that, and where we should draw the lines, and where there's some advantage in being able to use autonomy. And kind of where AI should, and maybe should not be used in that regard. Because

in my opinion, there's a difference between autonomy and AI. But that kind of brings us to the present day, where all this has kind of added up to the point where I ended up starting Seekar Technologies, which is the AI company that I currently lead and work for.

Sure. Now, you got my attention right off there, you are the first guest on the show to have experience of autonomous tractors, so I want to know about that. When did you see those? What are they capable of doing? What are their limitations? How long have they been used? How much do they cost, though? What is the cost-benefit ratio or the business case? Well… you get it, I want to know.

Sure. So what's interesting is it started with self-steering. So basically, these tractors can go through the field and steer themselves very straight, to the degree of like tenths of an inch. So they can maintain tenths of an inch of accuracy throughout this entire field, which is pretty profound, right? You really need straight lines in order to harvest some of these crops and plant them correctly. And so autonomous steering helps with that. And this was back 20, 25 years ago. So I mean, it was a long time ago, which is pretty interesting, because right now, we're just seeing autonomous steering coming into consumer vehicles to the point where you can actually use it and purchase it. Granted, autonomous vehicles are a little bit more of a complicated user environment than tractors or farming is, but the cost-benefit is profound. It's obviously come down every year, being able to purchase an autonomous package for agriculture, just because computers are getting cheaper, algorithms are getting smarter, and we're finding better ways to scale. But you'd be hard-pressed to really find a farm these days, at least in the United States where there isn't some sort of an autonomy package on at least part of their agricultural fleet. And it started with self-steering and now they're getting to the point where you can actually leave the entire cab of the tractor and it will basically till the field, it will plow the field on its own, turn itself around, and have very minimal human interaction, which is very, very interesting, because, in my opinion, that kind of helps pave part of the way for self-driving cars. I feel like they've kind of built on top of what they learned in agriculture because it's easier to prove in an agricultural setting.

Wow, did not know this. I'm wondering, can you tell me how they stay registered to a tenth of an inch? GPS isn't that good.

Yes, so actually there's a base station. So a satellite basically relays information to a base station. And we have to go set about this five-foot base station up in a corner of the field, which helps directly communicate that GPS signal to the tractor that's driving. So really, you have a satellite that's getting information to that base station, that base station is relaying information to a tractor, which is only a few 100 meters away, or maybe at most a mile or a mile and a half away, so it's a lot closer. So that response time, that control loop is a little bit more minute, which allows you to get those fractions of an inch. But it's pretty profound that it's able to retain that for such a long period of time, because a lot of the times, these tractors can be going for six to 10 to eight hours at a time.

Wow. And so then you got into defense, and you've got some things to say about autonomy in there. Can you tell us something about the applications of autonomy in defense that you had involvement with?

Sure. So to the extent of what I can say, I was involved with autonomy within projectiles, essentially. So pretty large projectiles and missiles. But I was able to be exposed to a lot of different things that allowed me to see how autonomy can really help benefit a soldier, right? Because soldiers are in very, very high-stress environments, and they have to make very large decisions in very small amounts of time. And autonomy can help in that regard by providing them situational awareness and help kind of reconciling some of the things that they maybe shouldn't have to worry about. But there's the other camp where autonomy starts to turn into artificial intelligence and artificial intelligence weaponry, and that's where people start to get a little bit squirmy and there's a lot more fear in that because a robot shouldn't be in charge of determining such a drastic decision on if someone should live or die. And that's where we need to be able to kind of define the line and draw the line. Granted, I didn't work on any of those projects where we were really explicitly using AI in a sense of defense, just because those projects in the Department of Defense itself are pretty minimal, to begin with. But autonomy itself is what I was primarily involved with, and able to see a lot of what, even not-- If you don't even use AI, you can still get some pretty sophisticated functionality on software and robotics and everything that still kind of maintains defense ethics. And you had a gentleman on the podcast a couple weeks ago that talked about the Geneva Conventions. You won't even begin to touch that really with at least the systems I was working on, but you really start to bring those into question with autonomous systems and AI in general in defense. And so yeah, that's kind of an interesting debate right now.

And you're talking about Tony Gillespie there who has written a book about systems engineering for autonomy. And you're also reminding me of another guest we had a few weeks before that, Peter Asaro, who is an activist with respect to lethal autonomous weapons and parsed that quite finely for us.

Yes.

And then that brought you to start up Seekar. So you've got this background here in autonomy that started when you were a toddler on the farm, that's practically a superhero origin story right there. And now it's taken you through your involvement in the Defense Department, and you've started Seekar. And just before we leave that defense work, I have had people with clearances tell me that the stuff that we see now is technology they were working on 20 years ago. Is that an assessment you'd agree with?

It is, yes, and that's part of the reason why I warranted us to starting Seekar is because there's a lot of the technologies that are being put forth on the field today are antiquated in defense.

Right. So you talked about motivation for starting Seekar? What is that?

So my partner and I both came from Department of Defense and we noticed a trend, kind of what you were just alluding to. There's a lot of technology that is being pushed forth today in defense, and really a couple other industries such as healthcare, were actually developed many years ago and that's for good reason, right? There's they have to be proven, right? You don't want to push something new that doesn't have a track record of success into the hands of a soldier because the consequences are dire if the technology is still at high risk. So it needs to have a very large track record before a defense company would consider using it. However, by the time that that technology does get that track record, it could be antiquated. And we see that a lot with some of these technologies that are being pushed in defense, that are using technologies that we're at their prime 10 to 15 years ago and they're a little behind the times. So that's kind of what started or warranted us to start is we kind of want to help change that. And we wanted to work our way into defense, but we ended up starting in hunting to start with and then moving into medical and now defense. But we really want to try to construct a way to allow soldiers and allow people to take advantage of these really high technological capabilities of artificial intelligence in industries that they're not being applied to. And what I mean by that is there's a lot of investment of time and a lot of investment of capital in industries such as self-driving cars, and smart home assistants in regards to artificial intelligence.

You said you started out in -- it sounded like hunting, which probably I misheard.

No, that's actually correct. So our goal initially when we started a Seekar was actually to try to bring forth these capabilities in defense and to try to be an AI player in defense. But we didn't want to start out by doing that because again, we know the consequences of us being wrong are pretty severe. So we started with hunting and we actually got quite a bit of work developing smart scope technology for hunters and being able to verify, "We're going to look at the field of view of your scope before a shot is fired to ensure that you're not firing upon an endangered species or there isn't something that you shouldn't be shooting at in your field of view." So it's kind of this extra layer of security that you're using artificial intelligence and computer vision with to try to add additional security for hunters that we were hoping would pose the gateway into defense. And we actually got a lot of work in that regard, developing that smart script technology and some extra technology for the hunting industry.

Let me get a picture of how this looks because that's pretty riveting there. You're saying you've got a scope that is smart -- What does it do if it decides there's an endangered species in the view? Does it lock out the gun or does it just put up a warning sign?

So we were building it to lock out the gun is the intent, but what we did, to begin with, was just put a red box around you know, "Hey, don't fire because look, this is here." So it was really just kind of a visual representation of what you should and shouldn't do. But our goal was to build it actually into the mechanics of the firearm so that we could help prevent some shots that shouldn't be fired. And we didn't get that far with these particular people or with these particular customers, just because the effort was cut short due to some funding efforts on their end with COVID-19, but that was the intent the entire time was to build it into the gun.

This is yielding all kinds of things I did not know existed. This is fascinating. By the way, going back to my earlier comment, what I was saying was that my friend with the clearance was saying that the tech that they had developed 20 years ago under clearance was what is emerging today. I imagine this is rather industry and field sensitive, but it sounds like you were looking at it the other way, that it had to be proven out before it would be adopted in defense.

Yes, so that's a good point, there are two different aspects of that. So defense actually is a big incubator for innovation as far as radar technology, stealth technology, they're the inventors of that. We have these technologies today because of them. And they're the ones that are the front runners for that. What I mean as far as some of these technologies that we enjoy as consumers, particularly with AI, like target recognition and object detection, is something that at least isn't as prominent in Department of Defense, at least from our fields that we were in at the time, and from our scope. And the reason was, a few reasons. One is because AI is very hard to explain. And if an algorithm can't explain its decision in a similar manner as a soldier can, then it's not going to be employed, right? But a lot of the efforts that are being taken to use AI, really, they're kind of trying to over-promise. So they're trying to say, "Hey, we're going to try and do a 200% increase in X, or we're going to try to yield a 300% increase in Y," or whatever. And really, you want to try to do incremental increase and try to augment soldiers' capabilities instead of replace capabilities. We're taking the aspect of providing additional situational awareness, rather than removing a soldier from the picture altogether.

Right. So it's not hard for me to speculate that what you're developing with the hunting rifle could be adapted to smart missiles so that if they saw that there was a child in the area they were about to hit, they diverted, perhaps faster than the remote controller could react. Any comment?

Yes. So that was the intent. We didn't want the debut of our technology to be used on these very expensive defense products with very cataclysmic consequences if they're wrong. So we started with something such as target practice where you have a very benign environment where you're just shooting at a target, and we're just trying to do target recognition in that regard, like a paper target. And then moving into animal conservation, where we're actually identifying different animals within the field of view, within the rifle scope, and then, hopefully, elevating more and more beyond that. So that was the intent is to try to provide these advanced recognition capabilities to defense but starting at a much smaller scale to prove the technology out and gain a track record behind it.

So we've got applications coming out of Seekar for hunting, defense, and medical technologies. What ties all these together? What's the value proposition or the focus of your company?

So the focus for our company is primarily computer vision. So basically, anything with image recognition, image processing, signal processing within images. We've done a lot with sound processing as well. But we've touched multiple industries just because we built the platform to be able to be repurposed for-- You know, we didn't want to have to rebuild an entire image recognition system for healthcare versus defense versus agriculture, whatever. We really wanted to be able to build the trunk of the tree that can be basically minorly branched out to

accommodate different industries and different practices. And that was the intent all along is to build a computer vision platform that's easily adaptable.

But your bio says you're building a clinical AI tool to advise neuropsychologists in diagnosing mental disorders. Is that machine vision?

It is, yes. So the COVID-19 pandemic brought up a lot of good opportunities for Seekar in regards to us being able to develop. So we had some requests, some opportunities from some doctors to develop a radiology screening tool that basically looked at X-ray images and determined whether or not a chest X-ray was indicative of COVID-19 or if it was pneumonia, or emphysema, or something else. So we basically took that target recognition technology that we'd built for hunting originally, repurposed it, trained it very heavily in a loop of some doctors on some radiology X-rays to identify what patterns were particular to COVID-19. And then we donated this product entirely to the cause of COVID-19 just to try to relieve demands of physicians. So this product is called COVID-AI, it's available on the App Store. And it's an app where you can basically import an image of a chest X-ray or take a picture of it, and it will tell you whether or not there are signs of COVID-19 in it. And what this allowed physicians to do was to basically triage patients a little easier and help them kind of direct care as needed to help relieve the demand of some of these radiology technicians. That branched into the current product you're speaking about, which is a neuropsychology product. And this is actually using an app as well, where you have a smartphone or a tablet where the camera is on the patient, and we're not recording data, we're just actually monitoring-- And this is part of the patient's consent, obviously. You have to have their consent and they're willing to participate in this trial. But the camera is recording what the patient says and monitoring their body, their facial emotions, what they're saying, how they're moving, if they have any idiosyncrasies about their movement. And what we're looking for with this is we're monitoring the patient to try to find earlier symptoms of neurodegenerative diseases such as Alzheimer's, Parkinson's, etc. Because a lot of these diseases once you find them, they're too late. And there are a lot of studies out there that you can read about that help back this, but there are symptoms within people's voice that are indicative of certain neurodegenerative diseases you can catch early, symptoms you can find in people's hands when they start shaking, in their skeletal tracking, in their body tracking. We're trying to take all of this together and map it to help augment a neuropsychologist's capability so that they can take this back with a report and say, "Okay, I've interviewed this patient, and these are the symptoms that I think that they've displayed. What did the AI capture?" And the AI is going to make a prediction to the physician and the physician is going to go back and say, "No, I don't think that was it," and it's going to train the AI. Or the physician is going to say, "Okay, I agree with this. This AI made an accurate assessment," in which case it validates it, and the AI gets smarter and smarter and smarter. But it goes one step further within this interview with the patient in that it's actually transcribing everything the patient is saying, and everything that the doctor is saying. So as I'm speaking, it's basically doing a movie script-like style, a document of everything I'm saying, everything the doctor is saying, so when she goes back and reviews this later, it's easy for her to recall. Because when we were approached with this problem, the physicians were saying, "A lot of times, we don't review the patients until two weeks later because we're in such high demand right now, and because of mental issues with COVID-19,

and people have been quarantined." And so they're just under such high demand. And this allowed them to recall faster so that they didn't forget any key details about the patient's interview. It's something that we're working on right now to try to train out. It will need a clinical trial. But we're working with these physicians right now to try to build it out so that it can actually hopefully detect a lot of these neurodegenerative diseases sooner.

I see. And the instances, the examples you described, would have called neurobiological, and the description that I read said neuropsychology and mental disorders so that made me go to schizophrenia, depression, narcissism. Does it get into that? It's easier to see how it would work and be useful for the other things, but how far do you go in the other direction?

Yes, so we actually started with depression. Schizophrenia is another one of the things that we do look for, but we started with depression and PTSD. So being able to identify key biomarkers for those particular conditions. And I'm not claiming to be a doctor at all. I'm advised by some very, very smart doctors so I'm going to explain it the best way that I've been explained, but they seem to be a little more adept at or it's easier to identify conditions such as depression and PTSD. But we had to start somewhere by training AI. So we built a foundation of training these two conditions, making it accurate in being able to identify these two conditions before we move on to something more complex such as schizophrenia, Alzheimer's, etc. But that's where we started and it's actually helped augment their capabilities a lot more being able to say, find different things within their voice because one thing that's key to what we're doing here is being able to kind of find out how you emotionally react to certain things. So if I go in for an interview and I may have PTSD, and a doctor asks me about a certain event, and I seem to emotionally react in a way that might indicate that a topic should be maybe delved a little bit deeper into, the AI is going to pick that up and say, "They seem to have some disengaged eye contact here, some volatility in their voice, etc., so let's look at this particular topic a bit more.

Right. Where do you get the training data for that?

We get them from the neuropsychologists and the patients that consent to participating with this type of study.

So that's supervised learning and so you're getting data that's already been labeled.

Correct. Yes, correct.

I think this is amazing to think about how the same system that can look through a rifle scope and tell you if a penguin has wandered into the field of view can be adapted to spot COVID on X-rays. And I think that's just a testimonial to how AI can be so general purpose. And so just so I'm clear on the application of that, it's an app you can download on the iPhone, you point the iPhone camera at an X-ray, and it says COVID, not COVID. Is that how it works?

Yes, that's correct. And what's interesting is we actually put that app through a clinical trial. So it was validated against X-rays that were validated against a board certified radiologist. And it proved to be I think it was 95.4% accurate in confirming what the board-certified radiologist confirmed was the condition present in the X-ray. So it was AI that actually went through a

legitimate clinical trial during the pandemic. But yes, it's downloadable in the App Store, it's called COVID-AI.

What is the average accuracy of radiologists?

The average for identifying COVID-19? I'm not sure what the average accuracy is. I'm sure it's probably much higher now because when COVID-19 first came out, it seemed to be a little bit less-- Nobody had seen it before, so it was a little less prominent in radiologists' minds. But I'm sure it's much higher now. I don't know the exact metrics on that, though.

So machine vision is what ties these things together.

That's the end of part 1 of the interview, we're breaking this one up for the usual reasons. I wasn't expecting to do that, but it was just getting so fascinating I had to keep going.

By the way, we had a personal encounter with AI here at the show; a day after the episode on Disinformation was published, I got an email that YouTube had pulled it due to a violation of their terms of service on statements about COVID-19, like we were promoting disinformation and conspiracy theories. If you listened to the episode, you know we were *denouncing* disinformation and conspiracy theories, so this was obviously one of those automated decisions that was looking at the content at a pretty primitive lexical level and not even looking for negations let alone doing sentiment analysis. Fortunately there was an appeal button, so I hit that and told them that they were doing their community a disservice by pulling this episode since it precisely aligned with their goals in creating that standard. And wonder of wonders, the very next day they emailed me to say that they'd made a mistake and restored the episode. Yay!

Another way that this show has personally encountered AI is that our post production sound engineer, Lee, who does an impeccable job in making us sound as good as humanly possible, so much better than those first few episodes that I hacked out myself on Audacity, told me that the industry has moved on from bandpass gating and that our noise reduction is now being done by AI. Alright! The rise of the robots marches on.

In today's news ripped from the headlines about AI, Github together with OpenAI have created CoPilot, which does for computer source code what GPT-3 does for language; in other words, just as GPT-3 can take a prompt like the beginning of a poem and then write more of the poem in the same style, you can type the beginning of a function into CoPilot and it can suggest what comes next. The underlying technology, called OpenAI Codex, is a transformer, the same as GPT-3, only it's trained on the public codebase in Github. If you think about how much code is stored on Github, it only makes sense that you could do something useful with it with AI just as GPT-3 has done with the text on the Internet.

Because comments are part of that code base, CoPilot can even write some code from you just typing what you want, because if it's seen something like that before in a comment, it'll have an idea what comes next. Sometimes that's even a whole function, albeit a small one. In my book Crisis of Control, I suggested that in 2027 a technology colloquially called CopyCats would be invented that allowed the computers to write software from verbal requirements. I though that was excessively optimistic, and it probably still is, but I also didn't think something with the capabilities of CoPilot would come along this soon.

I had a hard time believing at first that this could be useful, because, well, poetry is one thing – you want random creative off-the-wall responses, but computers don't do too well with random inputs, and I expected that the most this could do would be to create syntactically correct code that produced nonsense outputs. But people who have used it are generally enthusiastic, some of them hugely so. Alex Polozov, a senior researcher at Microsoft Research, tweeted, "Not exaggerating, Copilot will be in top-3 tech developments of 2020s."

Now, whether the code it come up with will work is not always certain. It's a lot like Google Mail suggesting what you might want to type next. A lot of that is predictable, like if you type the word "ulterior," it would make sense to suggest the next word be "motive," because when have you ever seen the word "ulterior" not followed by "motive" or "motives"? But the big issue swirling around CoPilot concerns licensing, because a lot of the code in the public Github repositories is protected by the Gnu Public License, which means you don't have to pay for it, but you can only use it if you propagate the license with your code, you include the source code, and you allow people to copy the source code for free. The charge levied against CoPilot is that because its inputs include GPLed code, it may be laundering that code into applications where the GPL license should apply. As an analogy, if you put the beginning of the lyrics of a well-known song into GPT-3, it might complete it either in a way that's not been done before and therefore qualifies as original, or it might complete it with the precise text of the original song, because it has quite likely seen those lyrics and realizes that they form the most common completion of what you put in. But now, what it's delivered to you is a copyright work, and you may not know that. You authors know that if you put even a sentence of the lyrics of any song in your book, you have to pay a license fee or you will be in trouble.

So you may end up with code that incorporates chunks of copyright code, although Guthub says this doesn't happen often. However, Armin Ronacher posted an example of where CoPilot had filled in the code for a fast inverse square root function by copying it from the source code of Quake III, including a sweary comment, and then preceded it with a comment claiming it was copyrighted by someone completely different. You can just hear the lawyers sharpening their pencils in the background.

Of course, there are also concerns that a system like this, if it gets capable enough that people start trusting it like today's self-driving cars, will one day run into a wall. Programmers should test their code, of course; but what if they get CoPilot to write the tests as well?

Next week we'll finish the interview with Kordel France, when, amazingly, we get to autonomous cars, and explainability, and artificial general intelligence. 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