SDS PODCAST EPISODE 967: AI FOR THE PHYSICAL WORLD, WITH SAMSARA'S PRAVEEN MURUGESAN

Jon Krohn:	00:00	What happens when a truck driver falls asleep at the wheel and there's no sell signal for miles? AI on the edge, designed for physical systems is the answer. Welcome to the SuperDataScience podcast. I'm your host, Jon Krohn. Today's excellent episode features Praveen Murugesan, VP of engineering at Samsara, a physical operations platform that processes over 20 trillion data points for the world's leading organizations across construction, transportation, manufacturing, retail, logistics, and the public sector. Hear all about how some SARA engineers make AI impactful across such a broad range of physical applications in this practical and excellent episode. Enjoy.
	00:41	This episode of SuperDataScience is made possible by Dell, Intel, Acceldata, and the Open Data Science Conference.
	00:49	Praveen, welcome to the SuperDataScience Podcast. Great to have you on the show. Where are you calling in from today?
Praveen M.:	00:54	Hey, Jon. Thanks for having me. I'm calling in from Amsterdam in the Netherlands.
Jon Krohn:	00:59	Yes. Appropriately in Amsterdam, a big center of shipping for today's episode, because we're talking all about physical systems. You're one of the world's leaders in providing software for physical systems. And so we're going to talk a lot about your company, Samsara. But right before we get into that, I want to talk about the data that you have to work with because I was blown away as we were doing research for your episode on the volumes of data that you have. So as VP of engineering, at Samsara, you're responsible for overseeing teams that apply AI to interpret massive volumes of video, vehicle and sensor data, 20 trillion data points covering 90 billion miles annually. So tell us about that huge volume of data

and what the challenges are, what the exciting things are about dealing with things on that scale.

Praveen M.: 01:52 No, it's a great question. I think it's funny that actually Sanjit founder recently mentioned that, hey, we probably are collecting more real world data than anybody else. And I think maybe Uber was the only one that's probably collecting a little bit more.

Jon Krohn: 02:09 And you actually, you previously worked at Uber.

Praveen M.: 02:12 I did, actually. It was a pretty good journey, like seven plus years. I think for me, my time at Uber actually made me want to keep working on more physical world problems. It's one of the reasons I kind of stuck to the space. And Samsara particularly is very interesting for the same reason that you mentioned, which is like the volume of data is just so much, which meant for me, it's like a kid in a candy store where you could actually build a lot of really cool applications with all the volume of data that's like nobody else has. The other part that was quite interesting was just the physical world operations aspect of it. So maybe to answer your question, going back, so some sort of people who are not familiar with is a business where we try to build solutions or like software and hardware solutions for our customers who operate in the physical world.

03:12 When you say the physical world, you can think about an example would be like people that are working in transportation or people working in construction or like field services, like even like the plumber who shows up at your door. When you have issues, you probably work with a plumbing company who have a lot of plumbers working in an area who they're going to dispatch one person. So we build software for all of these type of solutions. And for all these companies, when you think about it, there's a common themed problem in the sense of they all care

about safety as a theme, both of the people that work there and also of like the assets or vehicles that they operate, which usually can run in the thousands. They all care about efficiency, like how do I operate really efficiently? And a lot of these businesses operate with very low margins, so this is very crucial for them.

04:10 And then the world is becoming more and more sustainable today. So many of these companies think about sustainability as a goal and like how can they move in that direction? So that's effectively the type of solutions that we try to offer in terms of with this data. Now, maybe like I'll give you a couple of examples and we can even talk a little bit more of, maybe deep dive a couple of them. But when you think about safety and we cater to people who are like driving in the real world, we actually sell sensors, which in this case could be like a physical camera that you can install in your vehicle, which has AI running on the edge in this sensor where it's able to give you quick feedback around driving behaviors. For instance, you might be driving out on the road for a long time, it can detect trousiness and actually tell you like, "Hey, you should pull over for instance."

Jon Krohn: 05:04 That's like a camera in the cab of a truck or something like watching the driver or how do you tell if a driver's tired?

Praveen M.: 05:10 So it's actually like the camera is actually in the truck. So the thing that we sell is actually like safety cameras. So we actually have the technology where you can install the safety camera in the cab of the truck and it is actually like watching and learning from the behaviors of people's behavior as they're driving. And we run AI models on the edge, like obviously vision models that actually can detect for these behaviors. It basically does like vision modeling and then also some degree of temporal reasoning on the edge. And a part of the fun problem for us is like, how do

you run these models with the constraints of what comes together when you're running it in like constrained hardware.

Jon Krohn: 05:57 Yeah!

Praveen M.: 05:58 Usually happens on the edge.

Jon Krohn: 05:59 That sounds really interesting. So if you don't mind me digging into this use case a little bit more, obviously you can't tell us things that are your proprietary secret sauce, but to the extent that you can, can you tell us a bit more about this edge compute? What kinds of compute do you use on the edge to be able to do? It sounds like you'd be doing machine vision processing, which can be pretty complex, pretty computationally expensive.

Praveen M.: 06:24 Yeah. So we actually do, at a high level, we basically do a lot of our training in our backend for these models. It's primarily like inference on the edge. So for training, we have our own machine learning platform, which we built in- house. It's like a Kubernetes native stack and we also leverage Ray and it supports training, deployment and like feature management, et cetera. And once we do that, we actually have a system that can automatically manage the various versions of models that we have and can deploy onto the edge and the inference happens in the edge. So obviously these models, a lot of them are like proprietary that we built, especially based on use cases. We are also starting to experiment with everybody in terms of tiny models from like off the shelf, tiny LLMs, et cetera, for doing like voice translation, et cetera, on the edge as well.

07:23 So that's something new that we are doing. But when you really think about a lot of the vision models, it's like proprietary models that we build internally and it caters

to different use cases that we offer, which primarily like in the mobile camera was like around safety.

Jon Krohn: 07:38 Right. Really cool. Yeah. There's a lot of different ways that you could be setting up your machine vision algorithms to be specific to your use case, to be compute efficient on the edge there. And then so I know that you're not doing much training necessarily on the edge, of course. It's mostly inference on the edge, but even that's not something that I have any experience with, having cameras, doing machine vision in real time. Is that something that's pretty lightweight and easy today? Is there a lot of hardware that can do that kind of inference in real time that's not very expensive or rare?

Praveen M.: 08:12 It depends upon how you balance the constraints, right? So I think for us, really, there is hardware that you have to then operate with the hardware limitations, especially when it comes down to the inference models. That's usually what we try to do. It's not like we put in very expensive hardware out on the road, because we also have to think about the overall bomb cost of the hardware that we actually leverage. It's a balance, right? But ultimately at the end of the day, what leads our decision making is more around what type of use cases that we want to operate with and what is the compute that we need and how do we do it smartly. So I think a lot of the edge computational reasons are more along the use case itself. When you think about it, if I have to do an alert around predicting whether you are using a mobile phone, which is a common thing everybody's guilty of.

09:06 So that is a use case that we need to be always operating on the edge, primarily because you have latency considerations and potentially connectivity considerations, etc.

Jon Krohn:	09:18	For sure. I can imagine when people are driving, they can be driving a truck anywhere and you're not always going to have cell phone signal access.
Praveen M.:	09:26	Yeah, that's totally true. It's funny actually, look, one of the customers I met recently, they were in this business of actually laying out electric poles. They're basically operating, it's a utilities business. Their job is to bring electricity places where they don't have electricity, so which means they also are operating in an environment where they don't have roads. So we do hear that quite a lot and a lot of our customers, because of the nature of customers we have, you're operating in these environments which are most definitely like connectivity is not a guarantee. So we by default build with that intent that like, connectivity is not a guarantee. And how do you also build these solutions from a lens of can they gracefully degrade and operate in that environment where it could be off the grid for let's say like a week and come back. And that's fairly common as well.
Jon Krohn:	10:23	Right. Yeah. I was basically only thinking about intermittent for a few hours loss of access to a server, a remote server. But yeah, you could potentially be a week or more in that scenario you just described where somebody's setting up the electrical infrastructure, there's no roads, no cell phone towers. That is a really interesting use case to have to be building for. And I guess a lot of your solutions, you need to be ready for that kind of situation. Cool. Well, so we've gone into a bit of detail now. I kind of dug deep on this specific video camera one. What are other kinds of applications that you work on at Samsara?
Praveen M.:	11:01	Yeah. The other one you could think about is like efficiency. And when you think efficiency, the two class of problems that we could maybe talk about, one is like fuel efficiency. Fuel is usually one of the biggest spends when

for people who have like a lot of vehicles. I consistently meet customers who spend north of like a hundred million dollars a year just on fuel. That's like a very common thing. So for them, they would talk about like fuel savings in many different dimensions. And like a common one is on like fuel efficiency savings by finding which of your drivers are actually like not driving really well, are there like mechanical issues on your vehicle? Which one should I prioritize and how do I drive better fuel efficiency, which could quickly ladder up to like millions of dollars of savings for people. And the other use case on the space of fuel, which we've started recently working on, it's an interesting one, is on fuel theft.

12:04 We actually realized that, I think there was a study that I read recently where it was like $130 billion problem in the world where fuel is actually stolen or like siphoned off vehicles, et cetera, like fuel-

Jon Krohn: 12:19 $130 billion annually, I assume. That's correct.

Praveen M.: 12:22 That's

Jon Krohn: 12:22 Correct. Being stolen from vehicles, I had not thought of that leakage.

Praveen M.: 12:27 Yeah, I had not either. Actually, it was quite interesting. So what happened was we'd heard about this from some specific customers in specific regions where the cost of labor was very low and fuel cost is sort of similar globally. So there we actually saw more of a spike of this behavior. So one of the challenges all our customers in that space were telling us who operate in that region was like, "Hey, could you build a good solution for this? " And a lot of the solutions out in the field required you have to go buy external sensors, additional sensors, and actually install them into your vehicles.

13:07 So the challenge for us really was like, "Hey, how do you scale this solution for all types, classes of vehicles and that can operate globally?" Because the problem that you inevitably encounter is these sensors can be very noisy and the variants can be different based upon the vehicle you're driving. So there's no one single easy solution here. And we originally explored this problem by doing some very simple or heuristic based approaches like a rolling median as an example to smooth out data, to find variants in fuel levels and look for scenarios. But those started immediately, like there were diminishing returns, I would say, where the lossiness or the noise was just too much and effectively that was not a problem that we could solve. So we inevitably had to move down the path of, okay, let's actually look at different techniques here in terms of how do we apply machine learning algorithms.

14:06 Actually we do more like a statistical model here which can actually smooth out the variants and like effectively improve the hit rate when we find the anomalies in fuel. Interestingly, we originally did it for that one region and then we decided, okay, we have all this data and customers globally, let's actually scale this up and see what's happening. And what we realized is it's not a very regional problem, it's a global problem. So that was quite interesting for us as a learning, just taking the data and like trying to apply it for like a broader variance of customers.

Jon Krohn: 14:45 So basically you came up with some kind of model, like something like a, you don't have to tell me exactly, but I'm imagining it in my head as like a logistic regression model to detect to have some probability of this kind of adverse event happening. So you have all of these factors feeding into the model and you were able to get kind of labeled training data in one region. And then when you apply that globally, you found that we're seeing this all over the world.

Praveen M.:	15:14	Yeah. The main part of the problem is the variance in terms of sensor readings. So think about this as the part of the problem is we could have, like let's say you drive like a Toyota Corolla and I drive, let's say like Master or something.
Jon Krohn:	15:32	Yeah. I'm driving a Corolla and you're driving a Ferrari.
Praveen M.:	15:35	It's okay. That would be great. But the general idea is these senses are not equal. So when fuel level is reported by these sensors, there's a lot of variance and there's a lot of noise. You also have scenarios like when you try to think about this from a lens of big trucks, they usually have dual tanks, like meaning there's not just like one tank. The sensor reading is only from one tank and there's actually like fuel flowing between these two tanks, right? So the fundamental crux of this problem is about for you to do really high quality precision for like these detections, you have to have ways of eliminating the noise. And so we primarily do a lot more like statistical models, I would say, and like tune from there. That's the approach we use.
Jon Krohn:	16:31	Really cool. Fascinating application areas. And it gives us a good sense of how Samsara is applying AI across so many physical problems. You mentioned construction, transportation. We also have, from our research, we pulled out warehousing, manufacturing, retail, logistics, and some public sector applications as well. So it's a huge, vast problem surface area. And the idea with Samsara is to be helping your customers operate smarter across all of those verticals with an open platform that's obviously built to scale, dealing with the trillions of data points that your company is regularly processing. So yeah, so really exciting place to work, I imagine. So let's talk a little bit more about the edge compute. I found that pretty interesting. And so when you're doing this kind of edge computing and you're helping say commercial

drivers operate more effectively, digging a bit more maybe into that kind of cab camera example, how do you decide when you're going to be doing on device compute on the edge versus sending something back for processing on a server remotely?

17:49 Because I do understand, in the previous example that we were talking about, it sounded like we were kind of doing all edge compute, but based on our research, it sounds like sometimes there are some things that need to go back and be processed on a remote server. And so how do you handle these trade offs between latency and reliability and quality?

Praveen M.: 18:06 Yeah. I think it really comes down for us in terms of the constraints or constraints for the product problem, right? When you really think about the problem of like keeping people safe in the real world, you want to be able to ensure it's operating in a reliable environment without latency constraints or connectivity constraints. And so mostly whatever we do in safety, a lot of those solutions usually operate directly off the edge, right? But we do a lot of these problems which don't have those constraints. An example would be like all the fuel solutions that I mentioned, we have all of them operating directly in the cloud. There's like no latency constraints in any of those. We also work on a lot of problems around like routing, for example, like plan creation. It's less of a machine learning problem, more of an operations research problem, you could say where you're like trying to basically do like a traveling sales plan and like generate routes for your customers in terms of what's the most efficient way for them to fulfill.

19:21 So all of that happens more like in the cloud. The one other thing we are doing a lot, like most other companies, which I didn't mention is like, we're starting to leverage LLMs, it's becoming very easy. We have like a AI gateway,

like pretty much everybody at us now. And now that kind of opens it up where any product engineer could just directly start building or like building with these capabilities, especially like common solutions like summarization, data extraction and all of that into the product. So we have a cloud dashboard, which is increasingly becoming more AI first, not from a lens of AI first for AI sick, but really like how do you eliminate work when you see what people need to do and where could you get that leverage? So we're moving in that direction quite a bit, like how do you understand the use case so we can actually like pull in more of like automations and like agent take AI with LLMs.

Jon Krohn: 20:20 Right. Yeah. Let's talk about that a bit more. I mean, you said how having an AI gateway, most firms are doing that, but actually that specific term, AI gateway, I'm not really familiar with that. So yeah, tell us a bit more about this. It sounds like it's kind of, is this a way for engineers, scientists to have access to kind of like maybe pre-vetted models that you feel are safe? Is that kind of the idea behind the gateway?

Praveen M.: 20:44 That is correct. I think it's think about it as there are like multiple different models that we leverage and for multiple LLMs available today. Now, there could be like the best use case for the different use cases catered to different usage patterns of LLMs in itself. Now, how do you have like one single gateway, which you can like a central endpoint which is available within like the Samsara ecosystem, whereas a product engineer could actually like leverage that endpoint and based on their specific use case, define the constraints to say, "Hey, this is what I'm trying to do and like leverage this particular model for me. " Rather than like for building their entirely new facade or think about like the layer or the facade for every individual model that exists out there. So we use like OpenAI or like the Anthropic models or like even

Gemini and there's also like, we kind of built this gateway with a little bit of a layer of intelligence where you could be smart about things like cost and also make sure we are keeping up with our own internal rigor around being enterprise ready in terms of security, compliance, whatnot.

22:04 So all of that goes away from a data scientist or engineer's responsibility and they could actually delegate that out to another team and they just get it out of the boxes. So that's like a platform capability that we have.

Jon Krohn: 22:18 Nice. So the gateway kind of allows you to, it allows your engineers, your scientists, your developers to be able to take into account capability versus cost considerations right out of the box.

Praveen M.: 22:29 That's correct. Yeah. And we also design, it also gives you feedback loops where like for you to, as you almost think about as a training ground where before I'm ready to deploy this across all the customers and get it really into production, you could even get like very quick feedback loops from like, not just like try it out, but also like understand cost profiles, but also understand quality and promote to production. So it kind of gives you like this ecosystem which you could use like an example solution which I thought was pretty cool that somebody built was like a product engineer in the org, like built this solution where like one of the common customer problems that we had heard was, "Hey, we want to identify when our drivers are misusing the vehicle assigned to them." Meaning lots of companies have policies. When they assign you a vehicle, you're supposed to use it for specific needs.

23:28 Now the only way for you to do that historically was like you have to see through all the content of video footage or like randomly pick stuff or core sample for like individual,

right? Now you can technically delegate that problem to like a video LLM model and then which can automatically, like with the constraints you offer, it can actually like flag things. So we actually built that like a product engineer who's not done a lot of data science in the past was just able to pick it up and build it. Now that's the capabilities that it's unlocking, so which means like a lot of engineers, the lines have blood quite a bit in terms of like these different engineering stereotypes used to exist.

Jon Krohn: 24:11 Yeah. Yeah. Such a great example where you're basically, you're saying that historically you might have needed an AI expert, somebody with a PhD in machine vision to come into your company and be able to have this algorithm that could automatically label your vision training data for you. But now any engineer in the business and probably even non-technical people, they could vibe code a solution in a lot of cases to say, "I'd love to be able to label these data." And I can imagine a scenario where you could vibe code that kind of solution very quickly.

24:50 Yeah, really cool time to be working on AI for sure. And yeah, I'm sure we have a lot of listeners out there who have spent many years developing their chops on AI, but I'm sure we also have listeners out there that are just able to power through now. And so they can listen to an episode like this and hear ideas about like, "Oh yeah, that's so easy. I can be coming up with these ideas and applying them at scale, even though I'm relatively new to the space." A question that I have for you that this conversation is leaning into perfectly is with the huge amounts of data that you have, 20 trillion data points, and I'm sure that is growing exponentially every year, with 20 trillion data points at your disposal and these kinds of exponential advances in AI, whether it's this ability to vibe code, this ability to have an algorithm that

can automatically annotate vision data for you, how do you define your physical AI roadmap?

25:51 What kinds of internal decision filters help you help executives or help people who report to you distinguish between big long-term platform bets versus something that is an impressive demo, but might not have great operational payoff?

Praveen M.: 26:12 So I think it's a good question. In fact, actually we have a lot of that. I think at Samsara it's like, I would say like a lot of engineers generally, we run hackathons and so on. So there's like lots of demo culture in terms of like building things, which we think is great in a way. How do you filter for the noise when you think about it, right? So the main point in any company or any ecosystem still that holds true is you, it's not about the technology, it's really about getting like the outcomes and driving impact for your customers. So one of the things that we always try to drive our roadmap around is what is the business problem or what's the problem that the customer is actually facing and how does the solution that we offer create like either operational leverage for the customer or effectively is it moving the needle from a lens of like, if you put this as a product, will there be a willingness to pay?

27:19 It's usually a very good filter, I would say, whether you're chasing randomized ideas or like something that's actually concretely solving pain. So those usually are good filters. So we always started it from a lens of a customer and that's how the company operates. We rely on this notion which is very ingrained in the business called customer feedback loop. Anything we build needs to be shared with like a few customers and we get some feedback and then we it right from there. I generally think about artificial intelligence or like any of these new technologies that's like incredible from a lens of like doing

and trying new things, but ultimately all the companies that win out there are the ones who put the customer at the heart of the problem and trying to give them something which actually creates leverage. And that's the approach we try to take as a filter whenever we look at problems to invest back in.

Jon Krohn: 28:17 I like that. So it's the idea of this loop all the way through to the customer to ensure that the kinds of projects that your team, that the engineering team at Samsara are investing their time in are likely to have a good payoff to be valuable to your customers downstream. Cool. So yeah, so we've talked about now LLMs and those and the cool capabilities that all of us have access to now. Something else, another emerging technology that you've talked about in a blog post we caught is talking about quantum computing. So I wanted to get into that a little bit. In an article you noted, and we'll have a link to this article, it's in something called Digenomica. And so we'll link to that so people can read it in full, but you note that connected operations platforms already do many things that transport re-imagines and that quantum computing may continue to enable this by crunching quote unimaginable amounts of data.

29:18 So you talk about something that would today take a traditional compute platform, 10 septillion years can be done in minutes with quantum. So yeah, tell us about what quantum might be able to do in your space around routing, predictive maintenance, scheduling, and yeah, go ahead.

Praveen M.: 29:41 Yeah. I think when you really think about it, like a very practical example is like routing and scheduling, you could say. I'll give you an example, which is like pretty much everyone who studied computer science has gone through, like the traveling salesman problem. It's an NP hard problem, which means like you Now, just scales

exponentially in terms of computational needs for you to solve it. And anybody who's trying to solve that problem effectively uses optimization techniques. So in a simple sense to say it's like a traveling salesman problem where you're basically saying for routing, a vehicle has to visit 20 stops. Find me the most optimal way for it to visit the 20 stops. I do believe it takes, there are a little over, I think a lot more than trillion. I think maybe quintendal, I think is the next one, 10 trillion, like number of combinations that exist.

30:36 So I think over time algorithms, we've gotten good at approximation algorithms and then reducing the search space of that solution space and then coming up with answers. But I think getting to optimal solutions for the space of like route planning is like a great opportunity with quantum. Now, we just talked about traveling salesmen in terms of the one individual route problem. Now move that along to think about, hey, I have to make a 10,000 stops. I have 15 vehicles. Find me the most optimal. Yes, you can do that. It's also a very classic engineering problem called like the vehicle routing problem. And again, solutions today uses approximations, not like precision answers. Now, let's extend that two steps further. Now we've talked about stops. Now, between stops, what happens is routes. Think about it from a lens of like, I need to go from stop A to stop B, I need to parse a bunch of segments to get there and I have options within those segments.

31:44 Now the problem again multiplied to become something bigger. And let's add one more variance into it. Every time, whenever we all drive, we know that we actually have real-time traffic lights. That's real-time traffic, very real thing. Now, lots of what we are doing is like these approximation models are like just taking some predictive examples and then not really using closer to real-time data. And then they're just trying to model, okay, this is

what I think likely will happen. So with technologies, when computes sort of becomes really easily accessible and like the limits goes off, your ability to do things with much more high precision is possible. Obviously with things like routing, you can get like really hypersize and a lot more efficiency. The gains can be achieved. There's always the argument we can make of what is good enough, which I think there's like a point of diminishing returns at some point, but the possibilities are endless and the different variables you can react to are pretty much endless.

32:55 So I think what I generally think will happen is like today, a lot of the technologies we have help you make assisted decisions. And I think that's kind of great for you to get into a world where truly self-operating systems, I don't think we are anywhere close because like with these type of problems or like the compute capabilities plus the level of access to data to do make real time decisions are like limiting factors in that scenario.

Jon Krohn: 33:27 Yeah. Fascinating. It's going to be interesting to see all the application areas that quantum ends up making a big difference in. We've had in the past year, a couple of times, we've had Dell's global chief AI officer, John Rose on the show, and Quantum is one of the spaces that he's most excited about. He makes an analogy toward the way that you've been able to see for decades that super powerful AI capabilities were coming around now. And that now you can see the same kind of thing coming years from now with Quantum.

Praveen M.: 34:02 It's super exciting. Yeah, it's definitely like the pace of innovation in technology is just incredible and like super excited to even see what it'd be like in the next couple of years. So

Jon Krohn: 34:14 Yeah. Yeah, exactly. Really exciting indeed. So one other application area that is exciting for me that we haven't talked about yet is when we were preparing the research for your episode, your team sent to me this video, which I'll put in the show notes on commercial navigation, which is an area that had never occurred to me is so important. And so let me explain this. So when I need to go somewhere, typically I'll use Google Maps. There are other obviously navigation apps out there for kind of regular users like myself. So Apple Maps or some people might use Waze. I think a lot of professional drivers, Uber drivers are using apps like Waze, but all of those kinds of solutions don't take into account commercial considerations. So for example, in this video that I'll include in the show notes, right off the bat, it shows how there's this big problem in trucking globally where trucks will be directed to go under a bridge that is too low for them to go under and the truck gets destroyed.

35:28 And yes, a really dramatic footage is shown in this video. And so it sounds like Samsara has come up with a solution, which is a commercial navigation tool. So a tool that is useful for avoiding bridges that are too low, and there's probably other kinds of scenarios where a commercial navigation tool is essential as well.

Praveen M.: 35:50 It's actually interesting that you picked on the low bridge problem, right? So we actually learned about that problem through a customer. A customer in UK actually came out and told us like, "Hey, we actually get these issues a lot where one of our trucks hits a low bridge. They have a lot of them in the UK particularly." And it gets into this process where they now have to pay a lot of fine. They also have to, it's almost like a strike against in their record. There's a lot of implications. So we didn't know about that until this customer told us and they said, "Can you do something about this? " So we did two things actually on commercial navigation, which is like,

one of the things is like, we were starting to work on a product where we've heard this from customers through and through that like, "Hey, the navigation that I have, which works great for passenger vehicles is not something that I can use for commercial needs." And constantly the manifestation of that came in the direction of like most of our customers would end up paying fines.

37:01 They would get up into situations like exactly what you had talked about in terms of low bridge scenarios, et cetera. And they wanted to make sure they're compliant. I think a lot of, for them, like the driving need is also like compliance and safety of their drivers. So that's a big reason. So we built this product called like commercial navigation, which is exactly built for that purpose. So one of the key differentiations here is the navigation is aware of like your particular context, which means what type of vehicle are you driving? And then it can change the mapping context to suit essentially what you are driving. So if you are actually driving a truck with like, let's say like a 20 plus foot trailer, you will automatically be sent down paths which are like compliant to where you can drive rather than like down roads because there's just a path that exists, which is an important piece that we built.

38:00 And then we also built this capability where every organization usually has like different policies. Some people have actually told us like, "Hey, driver cannot make a right hand turn ever." So that's like a no-go. So we've also built in capabilities where you can customize- Do you mean

Jon Krohn: 38:17 Left-hand turn? They can't make a left-hand turn ever.

Praveen M.: 38:19 Oh, I think left-hand turn maybe. Yes. Yeah.

Jon Krohn: 38:22 Because

Praveen M.:	38:23	That's the- I think it was actually a left-hand turn.
Jon Krohn:	38:25	UK, but it's in the UK. It's in the UK, right? So then it is right-hand
Praveen M.:	38:30	Turn. I think you might be right. Yeah. But I've heard this from primarily companies that operate in more of the domains of school bus operations, for example. So that's one that I'd heard about this in. And then the mobility services, which is like where they are like paratransit type services, right? So we actually have some customers in those domains who actually ask for specific ... You could only do this one side turn based on region geography. And then we also want to enforce things like you always have to arrive at the right side of the road. So one of the unique things that we've tried to do is take foundational mapping capabilities, apply the vehicle context and individual constraints that are needed for our customers to kind of build this end-to-end experience that's like super compliant, take safety at the heart of it. And that's what we've tried to build.
Jon Krohn:	39:31	Really cool. I love hearing these use cases and to take a lot of the examples that you've provided, the use cases, Samsara use cases over this episode, it seems like kind of the general idea here is to be taking something that previously might've been manual and making it automated, making it algorithmic. And so things that previously might've been done in a spreadsheet are now done algorithmically and automatically. And so things like fuel costs, traffic patterns, bridge height can all be taken into account automatically. What are the kinds of organizational bottlenecks or hidden costs that you've seen when companies try to swap out human heuristics for algorithmic planning?
Praveen M.:	40:21	Well, that's a good question. I actually came across one very recently, just like a couple of months ago, right? So I

know I described this problem around the traveling sales plan plus the VRP for like route planning. So we actually built this product on route planning. Let's let our customers ... The goal we went out with is like, we have to get the most optimal routes for our customers and something, a problem that gets an engineer excited. Okay, this is like an engineering algorithmic problem. Let's actually build something great. Now we built that product and then we got it out in front of some of our customers and many of these organizations have a role called us like a route planner who's actually like responsible to make sure these routes are made and ready to go for like dispatch the next day for the various routes that their fleet has to go deliver to.

41:12 So as a part of our, like what I mentioned in terms of feedback loop with customers, like we went and showed them like, "Hey, these are the orders you need to deliver tomorrow. We created the most optimal routes. Here you go. " And the person goes, "Wait, I can't recognize any of these routes." And they actually said something that was quite interesting. They were like, "We have this driver, Sam, who's actually need to go do this delivery to this particular grocer and they expect them at like 10:00 AM tomorrow morning." Now that's like human context that the route planner always had and there was a lot more like human to human engagement reason for why they wanted it that way. So they wanted to preserve some of the relationships that they had between driver to the actual customer of theirs. And we just lost all of that because we were just thinking of this from like, it's an engineering problem, let's build the most optimal routes.

42:11 So what we decided to do was like build an iteration from there of this product where we realized, hey, we have to meet the customers where they are. And one of the challenges is going to be, it's not like everybody's looking for just like complete pure automation on the other side

to get to like something that would ultimately be like replace the human, get the system to do everything. Whereas what effectively people wanted was like in an assistive experience, how do you get me operational leverage, but then don't sort of destroy what I've already built here in terms of like cultural value, the human connection, et cetera. So what we ended up doing was, again, took the team technology that we built, but then adapted it to say like, we'll actually, instead of giving you magical routes, we'll actually take in your current routes and we will work to iterate from there.

43:04 We'll give you suggestions based upon your current inefficiencies we identify. So we will bring you along in this journey in terms of the transformation. So you're getting to efficiency, but then we are also learning from you in terms of the human constraints that exist in these problems, right? So this is a very interesting learning for me, which also kind of resonated where I was like, okay, I always wonder of this world where like, should we have route planners? What would we do to actually take pretty much most of the work and automate it? But then it gave me perspective to say like, what we really need to be building today is like assistive intelligence, which people can leverage because there's the human component that's incredibly valuable for people in business and like running successful companies.

Jon Krohn: 43:54 Great example. I'm glad that you were able to dig into that. I thought it might be a really tricky question, but you were right there prepared with an excellent answer. And so speaking of preparation, you've had a huge amount of success in your career leading high growth engineering teams all over the world, Chicago, Amsterdam, Berlin, Belgrade, Bangalore, and so on. And Samsara seems like an amazing place to work. Fast Company listed it as one of their most innovative companies in the world. The Fortune Future 50 ranked

Samsara as number seven in a recent ranking there. And so that's like super exciting. Double digit percentage growth in terms of revenue year over year for many years. So if people who are listening want to work with you, want to work at Samsara, are you doing any hiring and what do you look for in the people that you do hire?

Praveen M.: 44:57 We definitely are hiring. As you mentioned, the business is growing pretty fast. And I think the main key part we generally look for are like, we want people who actually show curiosity is one of the main things that we look for in terms of like really ... This is a very fascinating and interesting space. So we think anybody who's a builder needs to really get fascinated by the problem space in terms of what physical operations could mean, like want to meet customers and we enable that and learn from how operations work. So I think that's a very important piece to building great products, like making sure the builders really understand the audience and the problems. So curiosity is like one of the key ingredients that we really look for. And then effectively, it's going to be agility, I would say. The other one where this is a high growth startup and like startups in general, I would say like it's a scale up right now.

45:57 We could say it's a public company, we're continuing to scale up. We still like to operate like a bit of like a high growth startup, which means like a lot of adaptability and agility is a part of it, right? So not everything's figured out, not everything's like super bureaucratic. Do you actually love an environment where you're enabled and you can move fast and show curiosity and actually like build with focus from there on and get solutions out the door? That's usually what we look for.

Jon Krohn: 46:28 I love it. Moving the needle on projects, really making it happen. Exactly. It does sound like one of the key things. And Praveen, this has been a fascinating episode. In

general, you're doing outstanding, really exciting work. How can people be following you after this episode to get more of your thoughts or to get more on what's going on at Samsara?

Praveen M.: 46:54 You can actually follow me on LinkedIn. I should be pretty easy to find. I'm not super active on socials, I would say. So yeah, so LinkedIn-

Jon Krohn: 47:04 Too much work

Praveen M.: 47:04 To do. Yeah. I think it does get hard when you're primarily operating every day. So LinkedIn's a good way to follow, keep in touch. I do often talk to people, especially people reach out in terms of like interesting problems that we are working on and want to talk about like exploring new ideas, solutions. So definitely welcome that to reach out.

Jon Krohn: 47:30 Awesome. And then I ask all of my guests for a book recommendation and I already know what yours is. I can't wait to hear more about it.

Praveen M.: 47:38 Yeah. I know we were just talking about it before we started. The book actually, one of the books I read during the holidays was this book called Do Androids Dream of Electric Sheep. It's actually like a fairly old book. I think it came out in the, I think 50s or 60s. I don't remember. So this is actually a book that explores the idea of like a world where Androids and humans coexist and what does it actually mean to be human versus being an Android? It was a very fascinating read. I actually really enjoyed it and it's actually very appropriate for the times we live in with the pace of velocity with respect to the changes that's happening. It's not a very heavy read. It's actually a story novel. Recently I've been reading more like novels actually, but this one I actually enjoyed it and I thought

like something that everybody, this particular audience would enjoy it.

Jon Krohn: 48:41 Yeah, I love it. So it's 1968 to Android's Dream of Electric Sheep. It's by Philip K. Dick, pretty well known American science fiction writer. I haven't read it myself, but I've got to. And I'm just really quickly looking at a little bit of information on it online as you've been speaking and I didn't know that the book is the basis for the film Blade Runner.

Praveen M.: 49:03 Oh yeah. I think the ... Yes, I think it's been a long time since I watched Blade Runner. I think they take some concepts from Blade Runner, but I don't think reading the book is special. So I would say read the book. Yes. Exactly. Yes.

Jon Krohn: 49:21 Yeah. And I think it's not necessarily the same plot. It's just kind of like an inspiration. But yeah, I didn't know that. Really cool. Thank you for that great recommendation, Praveen. And it's been so much fun having you on the show. We've learned so much about the kinds of problems that you're tackling at Samsara with AI and yeah, hopefully we can get you on the show again in the not too distant future to get more updates on the exciting applications of AI that you're applying on such massive scale.

Praveen M.: 49:49 Yeah. Really enjoyed the conversation, Jon. Thank you.

Jon Krohn: 49:54 Fascinating episode with Praveen Murugesan today in it. He covered how Samsara processes 20 trillion data points covering 90 billion miles annually, how in- cab safety cameras run vision models and temporal reasoning directly on edge hardware to detect drowsy driving and phone use, even when connectivity is unavailable for a week or more. He surprised me by letting me know that fuel theft is a $130 billion annually global problem, and

that Samsara uses statistical models to detect anomalies despite noisy and inconsistent sensor readings across different vehicle types. He described that an AI gateway serves as a centralized endpoint that lets any product engineer leverage multiple LLMs while automatically handling cost optimization, security, and compliance. He talked about how commercial navigation differs from consumer GPS and how the traveling salesman problem with even a small number of stops has a mind boggling number of combinations, but how quantum computing could soon overcome this to enable real-time precision routing that accounts for live traffic data.

50:58 As always, you can get all the show notes, including the transcript for this episode, the video, recording, any materials mentioned on the show, the URLs for Praveen's social media profiles, as well as my own at superdatascience.com/967. All right, that's it. Thanks to everyone on the Super Data Science podcast team, our podcast manager, Sonja Brajovic, media editor, Mario Pombo, our partnerships team Natalie Ziajski, our researcher, Serg Masís writer, Dr. Zara Karschay, and our founder Kirill Eremenko.Thanks to all of them for producing a super episode for us today and for enabling that super team to create this free podcast for you. We are completely dependent on you and on our sponsors. So you can support this show by checking out our sponsor's links, which are in the show notes. And if you'd ever like to sponsor an episode yourself, you can get the details on how by making your way to jonkrohn.com/podcast.

51:45 Otherwise, share, review, subscribe, but most importantly, just keep on tuning in. I'm so grateful to have you listening and hope I can continue to make episodes you love for years and years to come. Till next time, keep on rocking it out there and I'm looking forward to enjoying another round of the SuperDataScience Podcast with you very soon.