SDS PODCAST EPISODE 961: DISTRIBUTED ARTIFICIAL SUPERINTELLIGENCE, WITH DR. VIJOY PANDEY

Jon Krohn:	00:00:00	70,000 years ago, a single evolutionary leap transformed scattered tribes into a civilization building species. Humans. Today we're aiming to enable an analogous leap, but this time in machines. Welcome to episode number 961 of the SuperDataScience podcast. I'm your host, Jon Krohn. For this extraordinary episode, I asked Dr. Vijoy Pandey to return to the show to explain how as head of Cisco's elite incubation engine out shift, he's helping bring together the currently isolated genius of today's individual AI agents to enable distributed artificial super intelligence. Here, Vijoy eloquently explain how distributed artificial super intelligence will advance human health, reverse climate change, and much more. Enjoy. This episode of SuperDataScience is made possible by Dell Intel Fabi and Cisco. Vijoy, welcome back to the SuperDataScience Podcast. We had so much fun with you on the show last time in episode number 941 that I had to bring you back right away. Vijoy, how you doing?
Vijoy Pandey:	00:01:06	Doing great, Jon. And so happy to be back again.
Jon Krohn:	00:01:09	Excellent. Yes. And in our last conversation, you painted a bold vision of the future. I was so excited by it. It was a world where humans and AI agents form collaborative societies to tackle humanity's biggest challenges. It's the kind of techno optimism that I am fueled by, but it was nice to hear somebody else who has that vision and you added so much concreteness to it. So again, people can go back and check out episode 941 for all of the skinny, all of the picture that you painted back in that episode. But some of the key things that we're going to build on in this episode are that you introduced us to the concept of multi-agent systems and the open source platform agency, which enables agents to work together seamlessly. But you also explained how today's agents are isolated geniuses brilliant, but unable to share their knowledge or collaborate effectively.

00:02:02 And you challenged us to imagine a future where agents don't just solve problems individually, but scale their intelligence collectively. And so this is part of why we had to have you back on so soon because that left us with a tantalizing question, how do we move from isolated agents to a world of distributed super intelligence? And so you're back to talk about that today. Today's episode is focused on distributed artificial super intelligence and I can't wait to hear all about it. So yeah, I don't know if you want to kind of give us a high level view of distributed artificial super intelligence right off the bat.

Vijoy Pandey: 00:02:44 So here's the way we think about this. If you think about, I mean we talked about human agent societies moving forward and we talked about agents coworking with us, not in hierarchies but in flat organizations. And I strongly that all scientific discovery, all human work, all technical work, all services work in the future, all social interactions are going to happen within these multi-agent human societies. And the next breakthroughs are going to happen through artificial super intelligence, which is achieved not through these big, large, isolated geniuses by like you said, but by collective intelligence, distributed intelligence and sprawling societies and civilizations. And the way to enable that is to build a cognitive fabric that allows all of these multi-agent and humans to come together and collaborate and solve for those big problems. So that's what we are looking for and that's what we are working towards and it's a big long journey, but we are excited to be on that trait.

Jon Krohn: 00:03:51 Yes. So the term super intelligence, artificial super intelligence is kind of a loaded term for a lot of people where some people kind of feel like it's so wooly that it's almost useless to bring up. But I actually think we can define it quite concrete concretely,

Vijoy Pandey: 00:04:12 Right? So like you said, I mean there are many, many definitions of super intelligence and as an industry we've gone from defining intelligence and AI to a GI and then a GI was not a well-defined term, so we went to a SI, artificial super intelligence. Some people are thinking that's also not a well-defined term. People are talking about powerful ai. Like you said, some folks are in the camp that we have some level of super intelligence already. We've passed the tutoring test for example. Is that a bar that we set ourselves for super intelligence maybe. But what is that feeling that there is something that is still missing? I mean, we know we all use AI today and it's super powerful. In some use cases it's pretty bad in certain use cases. So there is certainly something missing. So the way I think about them is what others have been talking about in terms of economic viability as well as technical viability.

00:05:11 So if I were to summarize the economic viability is do we have a system of autonomous agents, autonomous AI that can perform a hundred percent of what a human can do without human intervention? And whether it's a niche task or it's a broad task, I mean that's a for debate, but no human intervention. Can a system of AI agents come together and autonomously that task to a hundred percent? So that is one definition, that's the economic definition. Some folks like OpenAI are talking about this. DeepMind also has a similar definition. And then there's the technology definition, which is can you have an objective driven reasoning engine that can invent that is beyond its training data? So can a system autonomously generate novel ideas, novel discoveries, invent something that is beyond anything it's ever seen and that passes the rigor test? Of course you can hallucinate and discover stuff and again without human intervention. So I think these are the two definitions that folks have been rallying behind. So whether it's young lacoon, whether it's Demis OpenAI, I mean all of these, if you were to summarize

these were the two buckets, whether you call it A-G-I-A-S, I call it some other name. But these are the two sort of definitions that we are rallying behind and we want to look at it from the economic perspective as well as the technology perspective. And we want to achieve that through a scale out mechanism. And that's the goal that we're going after.

Jon Krohn: 00:06:50 Great definition there. I like how you also mentioned how there are people talking about things other than a GI, other kinds of terms like powerful ai, which is Dario am I think kind of popularized that term and we've done an episode on that kind of definition as well. So I looked this up. We have this previous episode on five levels of a GI that Google DeepMind kind of founded. That's episode 748. And then we've also got 832 if you're interested in hearing more about this kind of idea of powerful AI as a kind of alternative way of framing a GI or artificial super intelligence. But regardless, I think now we've probably talked about it both of us enough that our audience, that our listeners have some idea of what we're talking about basically machines that, as you said, can be tackling all the kinds of tasks that a human could be tackling and doing it at a level that is equal to or superior than what we're doing.

00:07:48 And there are interesting things where it seems kind of obvious that if we were able to unlock an algorithm or set of algorithms, maybe distributed algorithms to be able to achieve human intelligence that is, sorry, capabilities that are equal to human intelligence across all the breadth that humans are capable of, it seems like it would be potentially instantaneous that you would then have super intelligence across all of these capabilities because if you had distributed systems learning off of each other, teaching each other how they could be doing things better, it seems like that would proceed very rapidly. Awesome. So now that we have this definition of artificial

super intelligence under our belts, Vijoy, tell us about this idea of distributed artificial super intelligence or as I've seen you abbreviating it, DASI.

Vijoy Pandey: 00:08:44 Right. So if you think about human intelligence, because one thing that we have in common moving forward is the comparison bar for all of us is human intelligence. And whether you think about artificial intelligence, A-G-I-A-S-I, powerful ai, whichever definition you might have, we just talked about two definitions. One, which is economic in nature, one is technical in nature. There seems to be always this comparison metric, which is let's compare it against humans because that's the best thing that we know when it comes to intelligence. So if you think about humans and how we evolved and how in our intelligence evolved, we actually evolved intelligence across two axes. So the first 300,000, 400,000 years, human intelligence was actually scaling up vertically. So we were getting smarter and smarter. We were inventing tools, we were inventing processes, but it was limited because we weren't communicating that intelligence.

00:09:44 So the communication was a big missing piece. And so what ended up happening was whatever we invented and whatever processes that we came up with and the dangers that we were aware of and how we reacted to those dangers and the way we stitched our clothes together or whatever we wore, I mean I'm not an expert there, but whatever we did was limited to the lifetime of either that individual or that process. And so we became more and more intelligent, but it was very, very limited. And so we were scaling intelligence vertically. But as we all know, every system including intelligence can be scaled onto axis vertically as well as horizontally. And so there was this big evolutionary jump around 70,000 years ago, it's called the cognitive evolution in humans where we discovered language not just sounds, just paintings and patterns, but language. How do you convey meaning

semantics between people and then how do you convey that across humans, but across tribes and across generations?

00:10:57 So what happened when that cognitive evolution happened was we invented three things. The first thing being shared intent. So as a human society, we started sharing a common intent. Let's go and build this not just in the lifetime of me as a person, but in the lifetime of this tribe or this group of people. So shared intent and coordination as a result. The second thing that we invented was shared knowledge. So this is what's colloquially known as standing on the shoulders of giants. So I build a knowledge base, then you add to it, then you add to it, or you modify it and you keep doing that. And that is cumulative human knowledge. So we invented shared knowledge. And then the third thing because of the first two is now we could do shared innovation. So innovation itself itself wasn't a singular pursuit or an individual pursuit, but it was a shared pursuit.

00:11:59 So that's what happened when language got invented and semantics got invented and you started scaling horizontally because now you are inventing as a collective instead of as an individual. And so what we are seeing, and the big thesis here is that so far in intelligence, in artificial intelligence, in artificial super intelligence, we've been building bigger and bigger individual geniuses. And the framework and the infrastructure to do collective intelligence, to do distributed intelligence has been missing and that's what we want to go after. So distributed intelligence to us is to enable to build a framework that allows for shared intent, shared knowledge, and shared innovation to happen in this multi-agent human society.

Jon Krohn: 00:12:50 Nice. And so if I can make a bit of an analogy here, you can correct me where I get this wrong, but it sounds like

prior to 70,000 years ago when some pre homo sapiens species, or was it homo sapiens? I don't even, yeah, you can correct me on that point as well when I give you a chance to speak again. But this human species develops language and it allows for the first time us to go from learning things inside of one brain, the neural weights, the neural connections inside of one skull, and now be able to share knowledge. And then in recent millennia that sharing of knowledge horizontally across many brains has been accelerated dramatically through first writing and then through the printing press and then digital storage, the internet GitHub. And so we're able to share information more and more between human brains. And so it sounds like what you are up to at shift with distributed artificial super intelligence is allowing that same kind of change in machines.

Vijoy Pandey: 00:14:13 Exactly. I mean, I think, I don't remember what my book recommendation was, Jon in the last episode, but if it wasn't Sapiens by UL Harari, then let me just state that should be the book recommendation to go ahead and read. Because one of the big elements there that he talks about is how humans or homo sapiens took over the world. And that was through the power of the narrative. That was through the power of fiction and the power of a shared fiction. So whether it's religion, whether it's capitalism and money, whatever it is, the stories that we build and the way we rally behind those stories and the way we act to them and modify them is actually our biggest secret and the biggest invention. But that's like more of a anthropological statement. But if you think about in scientific terms and in computer science terms, what we really need to do is enable agents to again, share intent, share knowledge and share renovation.

00:15:23 And the way to think about that is let's look at agents today where they are today. We've built these really powerful entities, and let's take a concrete example. We

built a multi-agent framework, a multi-agent application for a customer for, and in that case we had two agents. One agent was doing network configuration and validation. Another agent was making sure that it's compliant, it's secure, and it scales properly and it's within budget. The problem is that today both of these agents are individual geniuses and they need, but they need to not only talk to each other but understand each other, share intent, share knowledge, and then innovate together to build out that network. That framework does not exist. So how do we take the same solutions or the same design patterns that happened with humans 70,000 years ago that changed the cognitive evolution to happen that made the cognitive evolution to happen within humans and make that same evolution happen within agents. So can we look at building about an agentic cognitive evolution? And that's the problem statement that we want to go after.

Jon Krohn: 00:16:43 That is wild. It is big. And I'm not surprised that you're a big fan of Sapiens given these kinds of interests that you have. It is one of my favorite books, certainly one of my favorite nonfiction books and I've bought a lot of copies of it for people. I don't think it was your recommendation at the end of your previous episode on the show because I'd probably remember that. I don't remember what it was. I'll try to look it up at some point while we're recording here, but I don't think it was that book. I'm pretty sure I'd remember. So let's dig into this sharing intent concept a bit more. So you've talked about these three core capabilities that are required for distributed artificial super intelligence just as these three things were required for humans to be able to horizontally share information and achieve things together.

00:17:31 So sharing intent, let's focus on that one first. We'll talk about sharing knowledge and sharing innovation next. So how is intent different from knowledge, for example? So I

guess the key thing is with intent, I'm guessing it's related to being able to figure out what some agents, whether that's a biological one or an artificial one, is planning on doing. And so that's kind of different from knowledge where knowledge is kind of like a set of facts maybe. And so yeah, so intent is it's really this intentionality, this ability to predict what an agent is trying to do.

Vijoy Pandey: 00:18:15 Let's break it down. So let's take that example that I just talked about. So intent, all of us are individuals as individuals have intentionality around a task that we want to do. So let's say prehistoric times, since we're talking about that, we wanted to capture that hill. Now that is a shared intent while we capture that hill or the village on the hill, we are all subject matter experts. So some might be warriors, some might be medics, there might be a king or a leader, there might be a treasurer. We only have so many coins left or so much grain left. So there are all of these subject matter experts with their own intents and their own optimization functions. The shared intent is can we capture that village on the top of the hill within budget with the least amount of casualties in the shortest amount of time, as an example.

00:19:19 So that is a shared intent because even though all of us might have different individual optimization functions and intents, a shared intent allows us to solve a broad problem together keeping the various individual intents under balance. And so in the concrete example that I talked about earlier, these two agents, so one agent is there to configure the network and make the best network possible. Since it's coming from out shift and Cisco, we might not really dig into the details of compliance and cost and time because that's not what this agent is supposed to do. The job of this agent, this individual genius is to be really, really good at building out a network with all of the devices that it has in hand, that scales well, performs well and so on. And then there's this

other agent in this case that is looking at compliance, that is looking at security, but is also looking at budget and on time delivery. So the shared intent is can you build out this new network, which is awesome. Performant scales, well carries all kinds of traffic. What is on time is under budget and is compliant and secure. That is a shared intent. You might have to make trade-offs on one side or the other to make that shared intent happen because individually they cannot solve it because they lack that common goal. So it's a common goal. So that's shared intent coming to shared knowledge because you picked that and what's the difference there, the shared knowledge.

Jon Krohn: 00:20:55 Well, quickly something that I'd love to just dig into on shared intent there for one quick second is it sounds like in the distributed artificial super intelligence framework that you've come up with, semantic protocols seem to be the kind of practical implementation of that shared intent. So maybe you could kind of tell us maybe even in the context of shift example, and I love that you're giving us that concrete example, it makes it so much easier to understand what we're talking about here. Yeah, so maybe in the example you could go into these semantic protocols so that we can understand how that particular solution allows sharing intent in agents. And if you happen to be able to get more on the capturing the town on the hill as part of the analogy to help us understand semantic protocols, I'm really enjoying that one.

Vijoy Pandey: 00:21:47 Right? So if you think about them example, because that's more concrete, let's go there first. The out shift agent, which is the network configuration agent, has been built on technologies that are specific to shift. So we've built an agent using the cloud technologies, the models that are available to us. It actually mostly aligns with what we see in terms of regulations and laws in the US because what we are most familiar with, it looks at design

patterns that are pretty broad because as Cisco and as out shift, we've looked at network design across many, many, many customers. So that's our worldview and that's out shift agent's worldview. It's like saying you and I, you have an Android phone and I have an iPhone and we can send a text to each other. So you can see a green bubble and a blue bubble barely. You can send text to each other, but you're speaking Japanese and I'm speaking English.

00:22:52 So what's in those texts we have no idea around. So we can connect to each other, but we can't communicate, we can't collaborate. There is no shared intent, there's no communication of information, there's no semantic exchange of information. There's no meaning behind this connection. And so the agent world today is full of connectivity, but devoid of meaning is the way I look at this. And what we are trying to do as a first step to enable shared intent is to bring about meaningful communication. And so when these two agents coming back to our concrete examples, when these two agents talk to each other between out and Swisscom, not only can they sit behind two models, two cloud technologies and connect to each other and exchange messages to a protocol like Gateway or MCP or whatever, but then talk about the layer above this and actually think about the meaning.

00:23:57 Start communicating semantic meaning messages that actually convey meaning and why is that important? Why that's important is we actually want to start looking at a common understanding. So when we talk about shared intent, the first thing to get to is grounding and a common understanding what is the problem that we want solve and do we all agree on the same terminology, the problem space, the same objectives, the same goals? So that's step one. Step two is discovery. So maybe while we start looking at the objectives, we realize there's missing

information. So can we start looking at and thinking about pieces of information and knowledge that are missing resolution of conflicting information? So we might say, I might say X, you might say Y, but there's now a common goal. Can we now resolve those conflicting pieces of information so that we are going after a common goal coordination?

00:24:57 That's the big thing because again, I might have to give up some things. You might have to give up some things to actually achieve that common goal. So coordination and then negotiations that come with it, and then finally evaluation of those goals. So all of these things can only happen at the semantic layer. It does not happen to just message passing and the protocols of that layer. So the semantic layer that we are building out is actually a layer or a set of protocols that allow all of these things to happen, whether it's grounding, whether it's discovery, whether it's resolution, coordination, negotiation and evaluation. And so those are the sets of things that we want to do to then enable shared intent to happen.

Jon Krohn: 00:25:43 That is super fascinating. And maybe this is going a bit too into the nitty gritty or maybe this is not even a question that makes sense and you're going to be able to refine this for me, but when you're describing these kinds of semantic protocols that allow for clear shared intent regardless of whether the algorithm is outputting in English or Japanese or S Swiss German or whatever, so it sounds like this might be some, how is this information captured? Tell us about what is the substrate, what is the information in these semantic protocols that's being shared between?

Vijoy Pandey: 00:26:26 So if you think about this, agents are human-like at machine speeds and scale, but also humans do not operate on machine speed and scale. And so is and only if situation here where agents can operate like humans, but

they can also operate like machines. So they can do API calls, but they can also do an lp. Humans, I mean, I don't think I can do API calls at speed, right? Maybe I can communicate at speed, but so there is definitely a space and a mode of communication that agents can perform that humans cannot. So I would answer your question in three different ways. The first being that I truly believe that all innovation all work that's going to happen tomorrow, this is where we started, is going to happen through multi-agent human societies. So agents and humans operating at the same fabric, same level, not in hierarchies but at the same level.

00:27:35 And so that is a goal that we're going after. If that was the case, then the lowest common denominator is natural language. And so we have to solve for the natural language problem because there are going to be humans and agents in the mix for a long, long, long time to come. So that's the lowest common denominator. Then you go one step further up and say, okay, let's say, and let's think about the A SI definition that we started with, which is can agents solve for either discovery or economically viable work a hundred percent of the time without human intervention? And so in that case, there is no human in the loop by definition. And so what does that look like? I would say even in that case, the lowest common denominator is still language because you will get agents and you will get components that are not from the same vendor, they're not from the same framework and so on and so forth.

00:28:37 So the lowest common denominator does become a language, and I'll put an asterisk there and come to that in a second. And then the third level is let's say we all did kumbaya and we agreed this thing is really inefficient and it's like putting neural links in our brains and that's the fastest way to communicate with each other. And that's your scary nightmare, Jon. But let's say we go there, then

absolutely, yes, the vector space is the way to communicate because the amount of information that you would have in vector space is actually reduced dramatically, significantly when you convert that and you tokenize that and actually make it go towards language. So I would say if everything was standardized and there was no capitalism in mind and people all agreed to things, then I would say vector space is the way to communicate. But in reality, I would say we would fall to number two or number three because that is how things will evolve over time.

00:29:48 And so that's the answer I would give. And I would go back to that asterisk for a second and say that we are talking about language here, but language models are not the only models that we are going to tackle in the future. So they're going to be world models. And one of the things that actually we've just read recently that came out of the Fair Labs and a bunch of other places is when you think about world models, tokenization is actually a big detriment. So vector space is the only way to communicate. So I would say that is something that we need to solve for anyway, but for now we are going after tokenization and language and that's what we're trying to solve for in distributor support hazards.

Jon Krohn: 00:30:34 I'm so glad you dug into that asterisk there at the end because it's amazing how this whole conversation, the entire more than 30 minutes that we've been recording so far this episode, any example that you've been talking about, I imagine with language tokens as the inputs and outputs. And it is so interesting how that is such a small set of the kind of meaning that can be conveyed, the kinds of problems that can be solved, the kinds of understandings of the world that we would need to solve, all of the kinds of problems that human intelligence, human intelligence can solve. And so yeah, it's a really interesting point in history where it seems where we've

been able to do things like take all of the language on the internet, and that's been a relatively easy thing to feed through large language models that are increasingly massive and therefore have all of these unexpected properties emerge out of them intellectual abilities.

00:31:38 But there's still years and years and years, probably decades of work ahead of us to even be capturing the data that we need for world models 3D navigation of the world. And yeah, it's a completely different kind of information storage. And yeah, I love how at the end of your asterisk there, you touched on how because if you take a multimodal system, a multimodal AI system, it has to have, if it does both machine vision and natural language processing, it has to have embeddings a vector space where that meaning converges. And that's not something that you can't just look at that and get a snapshot and understand it as a human looking at it. But we might not be far off from a machine that can kind of explain that to us. And so maybe bringing this all back to shared intent and semantic protocols, it sounds to me like the semantic protocols that the distributed artificial super intelligence framework that you've developed at out shift, it allows us to, regardless of what is being transferred, what specific information is being transferred between agents or humans in this multi-agent human collaborative system, whether it's natural language tokens or pixels or vector embeddings, whatever it is, the semantic protocols are a way of ensuring that the proper meaning is understood between multiple players in the system.

Vijoy Pandey: 00:33:17 The way I would phrase it is that we are starting with the language side of the equation. And so what we are starting with is the tokenized output. We are starting with natural language because that's where the world is today. And the other thing I would say is it's not apples to apples, but you can represent anything that you see or experience you can represent through language, whether

it's natural language like English or whether it's a mathematical language or the language of physics and equations. But most of the things that you see around you, you can represent in some form of language. It's an approximation. It's not apples to apples, but you can get pretty far along by just broadening your view of what a language is beyond natural language to the languages of science and physics and mathematics and so on. So where we are starting out with is language, because that's also the world where Cisco plays in, where we are the glue in enterprises between heterogeneous systems in brownfield environments, and that's where most of the world is going to be. We are not going to appear in a world where it's a very purist one model just from one vendor. I mean, that's not where we play. So for us to tackle this heterogeneous, distributed brownfield environment is where the bigger problems are in enterprises.

Jon Krohn: 00:34:47 Awesome. I love that answer. Thank you for clarifying everything for me. I think I now have a great picture of this first pillar of the distributed artificial super intelligence framework. Before I get onto the second pillar, I have a really quick random social side question for you. Vijoy, have you come across or maybe even watched a new TV series called Pluribus?

Vijoy Pandey: 00:35:07 Please do not give that away. I've been waiting to watch it. It's in my queue, and I've been like, everybody has been telling me you've got to watch it. I was waiting for all the episodes to drop and now they've dropped. So I'm probably going to binge watch it right after this episode.

Jon Krohn: 00:35:22 Nice. Yeah, I don't want to give anything away. I really enjoyed having the story unfold for me without knowing what was going to happen at all. So I won't spoil anything for you at all, but I think you're really going to enjoy it. And I can't wait to record a third episode with you. So that can be focused exclusively on pluribus. There's ways

that I'm not going to go into it all. I don't want to spoil anything for you, but there's things about our conversation that we've just been having, semantic protocols, these kinds of things that I think you'll, yeah, you might pull out of, you might see, you might understand what I'm saying now. Now you're

Vijoy Pandey: 00:35:58 Watch now making really difficult, Jon.

Jon Krohn: 00:36:01 Alright, so onward to the second pillar of your distributed artificial super intelligence framework. So you call this cognitive memory fabric, and this pillar is all around sharing knowledge. Like 20 minutes ago I interrupted you as you were about to start talking about sharing knowledge to distinguish it against sharing intent. So now the floor is yours. Go for it.

Vijoy Pandey: 00:36:26 Yeah, so intent, I mean intent can be shared through protocols and semantic protocols and so on. We just talked about that, where you need to figure out grounding, you need to figure out missing information, evaluation and so on and so forth. But that is for the task, the objective that as a multi-agent human society you want to achieve, you need to share in intent, but there is knowledge that is being built over time as well. So remember when we started back and compared this genetic cognitive evolution to what happened with humans, the big thing there was can we build up a knowledge base, the corpus of human knowledge as we like to call it, the standing on the shoulder of giants as we like to call it, what does that look like for heterogeneity? So what does that look like for many agents, for many vendors inclusive of what humans contribute in that framework.

00:37:30 And so the way we are thinking about that is to build out what we're calling a cognitive memory fabric. And the reason to pick those terms is so it is cognitive in nature,

so it's not data, it's not information, it's actually higher order. So that's number one. Number two is it's memory because it's persistent. So you are keeping it around for a while. You're summarizing things and you are extracting knowledge out of it. So that's number two. And number three, it's a fabric. So it's distributed in nature because by the very fact that we are looking at heterogeneous environments, there are things that we want to share and place in this fabric. And then there are things that we want to keep private. So as an example, I mean you and I are having a conversation, this is a bad example, but this is getting recorded, that this is going to get published on the internet, but let's say we were sitting in a coffee house conversation, that conversation is private.

00:38:36 So everything that we talked about in terms of shared intent applies because we are grounding ourselves, we are exchanging, we are discovering things, we are resolving conflict and so on and so forth. But there is no shared knowledge. I mean the shared knowledge is only, it's not really shared. I mean, you've imparted knowledge to me and I have imparted knowledge to you. So we've become better individual geniuses. In some ways we've scaled ourselves vertically, but we have not published anything that anybody else can benefit from or including us. I mean 10 years later, we might forget a lot about that conversation. So we've solved the shared intent part, we've solved the semantic communication part, but we have not solved the shared knowledge part so that the bigger agent, human society can actually go and do something with it. And so to us, the way to think about shared knowledge is so that we want to move the ball forward when it comes to the knowledge base of this shared multi-agent human society. And the way we think about building that out is we are looking at all kinds of memory. So we are looking at working memory, we are looking at knowledge, we are looking at ontologies, we are looking at knowledge graphs, and it's going to be a

combination of all of those things shared between multiple entities with strict privacy and access control mechanisms layered on top of that.

Jon Krohn: 00:40:13 Nice. I like it. This makes a lot of sense to me. It kind of reminds me how in our previous conversation when you were in episode 9 4 1 with me, you talked about how having different kinds of systems like knowledge graphs paired with probabilistic predictions, so this kind of blend of very firm, very rigid knowledge with probabilistic language, that those two kinds of systems paired together are what can make the most powerful AI systems of today. And so it sounds like perhaps with this now new term, new pillar, cognitive memory fabric that we're talking about in this episode, it sounds like it captures that idea from the previous episode, right?

Vijoy Pandey: 00:41:01 That's exactly where this came from. So going back to again, the example that we've been running with the concrete example, which is you have two agents and they're trying to build the network for Swisscom. And this idea behind the cognitive memory fabric actually came because of the work that we did with Swisscom where we realized that probabilistic engines like agents need to work hand in hand with deterministic software and a good solid formal representation of the world that exists. And so in this case, again, in the case of Swisscom, this was the network digital twin. And so what we realized was if you just throw agents into the mix and say, go and build out Swiss comm's next network, that's a recipe for disaster. You can go the other way and say, let's just build digital twins and write deterministic software. And then you're not exploring the entire massive space that generative AI and genetic software does allow you to do.

00:42:10 So you are writing deterministic software, which is very narrow in scope and space. So the combination of the two which says, let's build a formal model of the world around

us, in this case, a network digital twin, and then layer on these agents that can go and explore and just go at it and then be creative and do that shared innovation that we talked about. That combination is super powerful. And so we had some learnings from there. There were like three or four use cases that we went after and we realized we need to generalize this layer, this layer of representing formally what a word looks like so that other agents can act on it over time and contribute to it is what was missing. And so that's what the cognitive memory fabric is, which allows you to not just use it as a working memory of sorts, but synthesize it, summarize it, evolve it, redact, add to it, and that is what we are building out in a generalized manner. And that's what, yeah, it exactly came from this use case that we talked about.

Jon Krohn: 00:43:21 Cool. I'm kind of surprised that I nailed it so much on the head if, because, so lots of other points in this conversation or I guess any conversation I'm in, I'm like, ah, I think I have this really relevant point you brought up before and you're like, ah, well actually, so I'm delighted to have got one there. So we've now talked about the first two of three pillars for distributed artificial super intelligence. So the first one was semantic protocols for sharing intent. The second one, which we just were talking about was cognitive memory fabric for sharing knowledge. The third pillar is sharing innovation. And I understand you have something called cognitive engines as the formal name for your solution as part of this distributed artificial super intelligence framework.

Vijoy Pandey: 00:44:05 This brings up a quote from Raj Reddy who's a Turing award scientist computer scientist. He came up with this quote, and I heard it recently being shared by SAT NA on one podcast, I think it was Lex or wa, I can't remember which one it was, but he brought this back, which is artificial intelligence in general, AI can either be guardian, angel or a cognitive accelerator. And that quote resonated

with me so strongly because we were thinking about cognitive engines, and then Satya brings up this quote from Jedi. I'm like, yep, that is exactly what this cognitive engine layer is trying to do, which is we've done shared intent through semantic protocols. We've done shared knowledge through a cognitive memory fabric. What do we do with it? The whole aim is now going back and thinking about how do we do distributed super intelligence? How do we do economically viable work or new scientific discovery a hundred percent through agents, through a horizontal as well as a vertical scaling mechanism.

00:45:25 And so that is a problem that we want to go after. And so while we get words, let's say solving for something brand new in this, again, going with a running example of let's build a net new network which nobody has seen before. How can we build a net new network? Let's say it's a low earth satellite network that they've never built before. That is new innovation that is net new that none of them have been trained upon, but they need to get towards building that through shared intent, shared knowledge, and shared innovation. So this is where the shared innovation piece comes through. None of them have seen it. They want to invent something new. The way you invent that is through making sure that you have cognitive engines that are helping you along the way. And there are two types, like Jedi said, there are accelerators and then there are guardrails.

00:46:21 And so the accelerators allow you to do things faster. So I want to lean on the simplest thing could be, like we said, it could be a calculator. I want to lean on something that does certain task faster than any human or agent in this case can do. So we are building out engines that can accelerate certain aspects of a workflow. So that's an accelerator guards is making sure that it's not going haywire because you're doing something new, you're

inventing something new. You might not be within the boundaries of what is allowed. So how do you ensure that it's again, going back to security or compliance or it does not violate the laws of robotics. I'm just making that up. But there are guardrails that allow you to work in spaces that don't disrupt the laws of physics. So it could be as simple as, you know what this a world model and I know what jello feels like and that's encoded, I know what a wall feels like and that's encoded and this asset test that all robotics people go through, can you sort blue berries, right?

00:47:37 So I know what a blueberry feels like and that is encoded. And so these are guardrails that need to exist so that the models don't go beyond those guardrails. And so I think those are the two classes of engines that we want to build and then let these two agents of ours go and innovate. So you have the knowledge base, you have a way to communicate intent and coordinate, and now you have these two engines that you can rely upon sets of engines to accelerate work innovation and be guarded. So you're not doing something stupid. And across all of this, now you have shared innovation, not just go ahead and have fun and give me the solution to the problem I'm looking for. Sorry, yes. The solution to the problem I'm looking for.

Jon Krohn: 00:48:25 Yeah. Yeah. This is so fascinating and there's kind of even this energy in the way that you're describing it, I feel like I now need to try roll forward with your brain and kind of get your vision of how having something like this distributed artificial super intelligence can change the trajectory of what machines can be doing. We're kind of starting to talk about it there with maybe one agent has encoded what a blueberry feels like, and so it's able to be able to delicately pull blueberries and this sharing that information with other agents that have other expertises could allow for innovation to happen. So maybe there's

another agent that knows how to use solar power to move a vehicle around a field. And so those two together can say, oh, wow, we should be making a blueberry, picking a solar powered blueberry picking device. And so that gets me thinking about the same kind of vision for the future that you talked about a little bit in the previous episode. So in episode 9 4 1, you talked about how drug discovery could potentially be accelerated by agents working together. So you could have LLMs planning the research, you could have protein folding models, exploring shapes, nucleotide configurations that lead to new kinds of proteins, robots running wet lab experiments, compliance agents monitoring everything. So with decentralized artificial super intelligence with that framework, how does that kind of pipeline that you were describing in your previous appearance on this show evolve?

Vijoy Pandey: 00:50:11 Right. So that's a super, super awesome example there because it brings in capabilities that cross vertical boundaries. So let's think about direct discovery where you have maybe a large language model that allows you to first plan for what you're trying to achieve. So this is the goal building exercise and iterating exercise. So you want to build or create a drug that does X. So when you build a plan for that, you actually want to define the problem, the goal really carefully and iterate with a large language model to then break down that goal into discrete tasks. That's step number one. And you figure out which parts of the task goes to which kinds of agents moving forward. That's the first step. Then from that, you might go into an agent that is based on an alpha fold like scientific foundation model that is doing the protein folding exercise and exploring that space in a pretty comprehensive manner from that alpha fold is going to give you a set of possible outcomes that look interesting and you want to test that against these compliance agents, these effectiveness agents, these manufacturability agents and so on.

00:51:33 And so that's what you do next time. That's an iterative curve itself that you'll have to go through for a long time till you narrow the space further down and get to, yeah, maybe these are the 10 options that I want to take and do some wet lab experiments, actual rubber meets the road kind of situation here. And then you get into physical AI or embodied ai, and you have a bunch of these robotic labs doing these wet lab experiments to make sure that it passes master in the physical world before you get into maybe human trials or even animal trials for that data. And so this entire pipeline is going across multiple kinds of models from various vendors who are exports in each one of these things, and they need to get connected and talk to each other, and that's where the world is today.

00:52:26 So just being able to pass messages is where the world is today, and that's what we are doing with agency and HWAY and MCP and all of the work that's happening today. But then all of these entities need to then share common intent, which is let's build or create a drug for this specific purpose with this price point in mind, with this population in mind, with these risk characteristics in mind. So that is the goal that needs to be defined, and that is the shared intent. And so you need a set of sematic protocols to make sure that you iterate on that, make sure you discover information, resolve conflict, negotiate, coordinate, and get to a common goal. So that needs to happen. Then we need shared knowledge because as you're doing these things, as you're iterating between the compliance agent, the manufacturability agent, and the protein folding agent, they need to share knowledge and look at what's worked, what's not.

00:53:33 The compliance agent even needs to pull from the past and say, you know what? Even if this looks like it's going to work, you're going to get stopped when you're looking for approval from FDA. So there needs to be shared knowledge and you need to add to it, redact, modify, and

then finally you need these guardrails and cognitive engines that are doing guardrails and acceleration so that you can do these experiments faster, but within the boundaries of what's possible. And so that's the shared innovation that needs to happen above the shared intent and the shared knowledge layer. So that is how I see the drug discovery of the future as going to happen, as going to make me a healthier person and make me live for 200 years and let me play the guitars that I still have to pick up. And I still don't know how to play.

Jon Krohn: 00:54:29 You don't know how to play the guitars behind you.

Vijoy Pandey: 00:54:32 This is always, it's like super int intelligence, Jon. It's like you'll always be working towards that goal.

Jon Krohn: 00:54:39 I see, I see. I got you. Well, they're occluded a little bit, but they still look beautiful. I'm sure they'll sound nice once you have the time to play them. When artificial super intelligence is handling everything from drug discovery, climate change issues, healthcare education, it's unreal, the number of areas that these artificially super intelligent systems when they're distributed, when they can share intent, knowledge, innovation effectively, what they're going to be able to do for us in the coming decades. It's staggering, and I'm really excited for it. Thank you for playing such a big part in this. Something that we talked about in the preceding episode that we haven't in this one is that you are doing this from the position of leading out shift, which is an internal incubator at Cisco. Probably most of our listeners have heard of the company Cisco. For listeners who might not remember from your preceding episode on the show, can you remind us of out shift's mission and where does distributed artificial super intelligence fit into the broader work you're doing there, including really mind blowing things like quantum networking that you mentioned in the preceding episode?

Vijoy Pandey: 00:55:52 So out shift is Cisco's internal incubation engine. So we are a group of people that are looking at horizon three, horizon two problems in spaces such as artificial super intelligence as well as quantum computing and spotting the next trends and paradigm shifts that are going to happen in these spaces and making sure that Cisco is well positioned in those spaces to build businesses with those paradigm shifts. And the way we are thinking about those things is we put the Cisco on these problem spaces, and you mentioned this in the guitar playing narrative that Jon, but to me, these two problem statements of super intelligence and quantum computing are problems that we as an industry have worked on for decades in the past, and I believe we'll be working on them for decades to come. These are shifting goals and these are hard problems. And so we've be working on both of these problems for a long time to come, but as Cisco, we have a pretty specific take on both of these swim lanes.

00:57:00 And the take is that we are a networking and security company primarily and then an observability company. And so we are in the business of enabling distributed systems. So the way we look at any problem statement is whereas the world is trying to build these massively vertical, big bad systems, or we like to call it lone wolves, we are here to empower the pack. So while you are building the lone wolf, we are here to build our distributed scale, horizontal scale, scale out systems and empower the pack. So when we look at super intelligence, and we've talked about this in this episode, we are looking at distributed super intelligence and the same thing that we are doing in quantum computing where as people build out these quantum machines which are bigger and more effective and can solve larger problems and have more and more qubits, we are here to build a quantum network that connects all of those quantum nodes together in a distributed quantum computing system so that you can solve problems, larger problems

faster. So that's our entire gig when it comes to both super intelligence and quantum, and it's our job to spot the next paradigm shift and enable it to Cisco for Cisco.

Jon Krohn: 00:58:21 Really cool Vijoy, for listeners out there who are inspired by the vision that you've been talking about today, which is about as inspiring as it gets, what can they be doing? I understand for example that you have a white paper that's just come out, it's called Scaling Super Intelligence, a cognitive architecture for distributed artificial super intelligence. Maybe tell us about that white paper where people could find it. I mean we'll have it in the show notes, but maybe there's kind of the white paper. What are other things that people can be doing to engage with out shift agency or be involved in distributed artificial super intelligence?

Vijoy Pandey: 00:58:58 So that white paper actually talks through a lot of the conversation that we just had today in this episode, and it talks through the parallels to human intelligence and how that has evolved as well. It talks about the details behind the three layers we just mentioned. So reading that white paper would be step one. Definitely go and check that out. I would say if you are a data scientist, a developer, or just plain not happy with the state of things when it comes to energetic AI and wished that agents could do something better and bigger and bolder and not be in this state where they are today, come and join that vision and come and join that mission because we are all in the same boat. We're trying to take agents and make them solve really hard and sometimes really boring problems. And the reason they're boring is that we've talked about them time and again, and we have not had the tools or the capabilities to solve them in new and different ways.

01:00:02 And I think we are at the precipice of having those tools at our disposal. So let's come together and build those use cases out and bring us those use cases. Let's solve

those boring hard problems through these new tools that we have and look at that paper and help us solve this in a holistic way. Because yes, we are trying to solve the identity problem through agency. We are try to solve the communication problem and messaging problem through eight way and MCP, the observability problem and the discovery problem. Those are all great problems, but we are looking at the weeds here. We need to take a step back, let's say tomorrow we solve for all of those problems. We will not be solving for this vision that we have. We will not get to a place where agents can come together and solve for this definition of super intelligence because the bigger problems are the semantic problems, and we need to start with those problems right now. And the way to start with those is read the paper, but bring us the use cases and join us in this journey.

Jon Krohn: 01:01:08 Awesome. That sounds great. And because the agency is a open source Linux Foundation project, people can also be getting involved through that route as well.

Vijoy Pandey: 01:01:18 That is a great way to get involved. You can go to agency.org. That's A-G-N-T-C y.org agency is a Linux Foundation project. We are also core members and launch members of a two a project from Google, which is also a Linux Foundation project. And then just this November, we became a launch partner for the Egen AI Foundation where MCP exists, and that's also a Linux Foundation project. So there's a theme here, but we are launch partners for all of those three because we believe all those three projects are complimentary to each other. They bring things to the table, they need to get glued together to solve for that communication layer problem that we just talked about. And then yes, above this we are building out the semantic layer and there's knowledge layer to make everything work better. So Jon, yeah, one thing I just wanted to state there is that I think I misspoke that Agent AI foundation where we are a launch

partner as well as where MCP currently resides. It was formed in December and not in November, so I just wanted to clarify that.

Jon Krohn: 01:02:29 Awesome. So many things for us to dig our teeth into after the episode. We already have your book recommendation from earlier in the episode. It's Sapiens by Yuval Noah Harari. So we can skip right to my final question that I ask all my guests, which is Vijoy. Oy, you are brilliant. You are such a joy to listen to. I loved our previous episode. I loved this episode with you. I can't wait to do another one. I'm sure our audience can't as well. But in the meantime, how can they get your thoughts? Where can they follow you on social media?

Vijoy Pandey: 01:02:59 I'm the most active on LinkedIn, which is LinkedIn slash in dash bjo. So we'll put the link in there. So I'm the most active there. I do post on X, but it's actually getting very noisy, so it's hard to extract signal from the noise. But you can also go to vja point a.com, which is a personal website. But the best place to get all of this information is actually to head towards out shift.com, which is where it's not just me because I'm just the spokesperson, but it's the team that is actually innovating, the team that is doing work. The team behind all of the things that we've talked about. They are providing more details and more flavor to what's happening. So go to our oxford.com because that's where you'll find a lot of the stuff that we've just touched upon and a lot more to come.

Jon Krohn: 01:03:49 I love it. Congrats also on getting the LinkedIn URL. That's just your first name. That is quite a catch. Nicely done. Alright. It is been so awesome having you on the show, and as I just said, I can't wait for the next time. Thank you so much for taking the time with us from your busy schedule.

Vijoy Pandey:	01:04:05	This was amazing, Jon, as always, fun conversation. Thank you for having me.
Jon Krohn:	01:04:12	Always great to have Vijoy Pandy on the show. In this mind blowing episode, he covered how distributed artificial super intelligence is. The idea that artificial super intelligence will emerge not from building ever larger individual models, but from enabling AI agents to collaborate much like humans scaled intelligence through language. 70,000 years ago, he talked about how that cognitive revolution gave humans three capabilities, shared intent, shared knowledge, and shared innovation. The same three pillars VIIs team is building for AI agents. On that note, his team at Out shift, Cisco's internal incubation arm is building this framework while contributing to open source projects like Agency A two A and MCP. So you can get involved too. As always, you can also get all the show notes including the transcript for this episode, the video recording, any materials mentioned on the show, the URLs for Vijoy yous social media profiles, as well as my own at superdatascience.com/961.
	01:05:08	And now it's my turn to give thanks to everyone on the SuperDataScience 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 another amazing episode for us today for enabling that super team to create this free podcast for you. We are deeply grateful to our sponsors. You can support this show by checking out our sponsors links, which are in the show notes. And if you yourself would ever like to sponsor the episode, sponsor an episode. You can get the details on how by making your way to Jon cron.com/podcast. Otherwise, share the episode with somebody who'd like to be blown away by distributed artificial super intelligence.

Review the podcast on your whatever podcasting platform you use to listen or on YouTube, subscribe obviously if you're not already a subscriber. But most importantly, we just hope you'll keep on listening. I'm so grateful to have you listening and I 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.