SDS PODCAST EPISODE 1001: HOW AI ERASED MY CAREER MOAT, AN EPISODE #1001 SPECIAL: JON KROHN INTERVIEWED BY KIRILL EREMENKO

Jon Krohn:	00:00:00	I spent 20 years building a career as an AI expert through machine learning research at Oxford, writing a bestselling book on deep learning and leading data science teams in a range of industries. And then earlier this year, the technical skills mode I'd built for myself disappeared because of AI itself. Welcome to episode 1001 of the SuperDataScience Podcast. I'm Jon Krohn, and I usually am the host of this show. But to mark 10 years of this podcast and over a thousand episodes, today's episode will be hosted by Kirill Eremenko, who founded this podcast a decade ago and who hosted the first 431 episodes. In a role reversal, I'll be the guest in today's episode so that you can learn a bit about me as well as my thoughts on AI rapidly usurping technical skills, whether we're in an AI bubble, the one key reason why I've seen AI projects fail, the intriguing relationships between AI and biological neuroscience.
	00:00:55	So as usual, lots of AI in this episode, but unusually I'm the one answering the questions instead of asking them. I hope you enjoy the flip. This episode of Super Data Science is made possible by Anthropic, Cisco, Acceldata and Notion.
Kirill E.:	00:01:12	Jon, welcome to the podcast. How are you
Jon Krohn:	00:01:14	Going? I'm going well, mate. How are you going down under?
Kirill E.:	00:01:18	Very well, thank you. And first of all, huge thanks for letting me take over for this one episode and host you for a shake.
Jon Krohn:	00:01:26	Yeah, thanks for letting me take over for the last 600.
Kirill E.:	00:01:30	That's fantastic. Yeah, you've been hosting since 432, episode 432. How do you feel? How does it feel to host 568 episodes and going?

Jon Krohn:	00:01:41	Yeah, it has been a journey. I don't typically go back and listen to the earliest episodes. In fact, I'm kind of nervous about how that would be. I imagine, or at least I hope that I've improved over these five or so years that I've been hosting the podcast. I don't know. What do you think? You've been listening since the beginning. You've even had me working with amazing leaders on speaking to hopefully improve my communication skills. So hopefully I have improved.
Kirill E.:	00:02:11	Well, for sure. I actually listened to episode 432, relistened to it yesterday and you've definitely come a long way. And you already were great, but you definitely come a long way and I think you're a better host than I could have been in all these years. So on behalf of our listeners, thank you so much for taking
Jon Krohn:	00:02:32	It. Well, I mean already you're just coming on. The level of energy that you have, it's so infectious. I love it. I think I'm learning from you already in the few minutes that you've been hosting this. I really like the level of energy. Thanks.
Kirill E.:	00:02:46	Thanks. Okay. Well, the goal for today is to find out a bit more about Jon because Jon has been, for all of you listening out there, Jon has been hosting this for five years or more now and you get to hear bits and pieces about him and his habits and how he puts water next to his table and drinks it in the morning and the books that he's reading and stuff like that. But today we're going to dive deep and find out a bit about Jon's background also what Jon has learned over the past five years of hosting this podcast and interviewing some of the greatest minds in the world in terms of AI and data and also what Jon thinks about the future that's coming, like what kind of opinions Jon has for. Sounds good?
Jon Krohn:	00:03:29	Sounds great.

Kirill E.: 00:03:31 All right, let's do it. Okay. So quick background. Yo have a PhD in neuroscience from some university in the UK, right? Tell us a bit about more about

Jon Krohn: 00:03:42 That. I do indeed. Yeah. So I studied neuroscience at Oxford University and I had been really interested in neuroscience because I thought I might want to be a neurologist or something like that. I've always been fascinated in how your mind arises from physics, how there's rules of physics that allow us to have chemistry, that allows us to have biology, that allows you to have all of the things that you perceive and feel and do are dictated by these molecules. And I've always found that really fascinating. So I thought I wanted to make a career out of it, thought I might want to be a neurologist, got a full scholarship to do a PhD in neuroscience at Oxford and I was like, "Oh, I'll go back to Canada afterward where I'm from, do med school after the PhD." But I don't know. A lot of my friends became doctors in that time and no one really recommended staying in that career.

00:04:41 But the really cool thing was during the neuroscience PhD, I had already had experience with a lot of ... I always loved probability and statistics in my undergrad studies. And so when I started at Oxford, I naturally gravitated towards even more computing. So I started learning how to do massively parallel computational statistics, how to do machine learning. And so I've now been doing that professionally since. Well, up until a few months ago when Claude just started doing all of that better than me, all of the career moat that I had just disappeared thanks to, I think Claude was 4.6.

Kirill E.: 00:05:21 It's crazy. We'll definitely talk about career modes a bit further down the episode.

Jon Krohn: 00:05:26 Yeah. That's where I started and loved it, had some great papers including a NEREPS paper, which is a lot of people would say the NERPS is the most prestigious academic

conference out there. So had a paper there in 2010 that was accepted. And yeah, I've been working in a number of industries since. Went right after my PhD to work at a hedge fund in New York and I've been there since now coming on 13 years.

Kirill E.: 00:05:51 Not at the hedge fund though.

Jon Krohn: 00:05:53 No, no. Only two years at the hedge fund and then- In New York, 13

Kirill E.: 00:05:58 Years. 13

Jon Krohn: 00:05:58 Years in New York. Exactly. Thank you for clarifying.

Kirill E.: 00:06:01 But right now you're in Canada, right? You have family in Canada.

Jon Krohn: 00:06:04 My entire extended family lives in the Toronto area and so I try to spend as much time here as I can. One of the nice things about the podcast is that I have an identical recording set up in New York and in Toronto and so I can record hopefully the exact same quality of episode in either place and yeah, it gives me a great chance to spend a lot of time with my family. If people have been listening regularly, then they will recall my grandmother, who's now 96, who's appeared in a number of episodes. I think episode 700, 800, maybe 900, not 1000 because you and I did that one, our listeners, but she's here and you just never know when she's not going to be here. So try to make as much time with her and my parents are still both really healthy. So yeah, spend as much time here as I can.

Kirill E.: 00:07:00 That's really nice.That's very important. So neuroscience, that's an interesting pathway, like working with data and machine learning during your PhD and then transitioning into this space. And have you quenched that first for understanding how the human brain works? Remember, like you said, one of the reasons you went for

neuroscience is like, it's a mystery. In fact, I like to think of it this way, that there's three big mysteries in the world. Which one would be the most interesting to me? The brain, what's out there in space and what's down deep down there in the ocean. And so you chose one of these three mysteries. Have you quenched the first for knowledge there?

Jon Krohn: 00:07:43 Definitely not. And I would say I went to work at a hedge fund in New York for the money, which I think was a mistake. The reason why I only had that job for a couple of years was because I couldn't stay motivated about going to work for the purpose of making money. The kind of trading that we were doing at that hedge fund was there were some small parts of the job that were contributing to the economic value of all of humanity, but for the most part you're not. 90% of the actions that you take in the day are not furthering humanity in any way. And so it was just hard for me to stay motivated. What was the question?

Kirill E.: 00:08:30 Did you acquince your thirst about

Jon Krohn: 00:08:31 Understanding

Kirill E.: 00:08:32 How the brain works?

Jon Krohn: 00:08:33 So yeah, I now really regret since finishing my PhD in 2012, so 14 years ago, I now regret not staying in academia. I really enjoyed it. Interesting. Yeah.

Kirill E.: 00:08:51 Didn't

Jon Krohn: 00:08:52 Know that. Yeah. And so I'm actually, I'm taking small steps now wherever I can to be doing it a little bit more. So I'm supervising PhD research close to you at the University of Auckland. Yeah, exactly. New Zealand. Yeah, exactly. So Miriam Kakpor, she may be listening. She's the PhD student that I'm supervising there and she's doing really fascinating work. So bringing AI into clinical

settings ideally with like some kind of physical, robotic embodiment that would be therapeutic for say patients or maybe even healthy people, there's a lot of different angles that you could take with that kind of thing, but so we're doing that and so that'll lead to some publications and that's psychiatry close, that's related to neuroscience and still related to how in that particular case, how problems with your biology with those physical processes lead to, in some cases, a lot of pain for a lot of people.

00:09:54 Something like one in three people will experience say anxiety and depression, anxiety or depression in their lifetime, which are related conditions. And so if you're not one of those one in three people, you probably know someone who's really close to you who is impacted. So there's a huge amount of potential to be improving a lot of quality of life through AI in my opinion. So we'll see what happens there specifically. And then I have a few other things that are less mature and not official yet, but basically there are other ways that I am trying to get a little bit more academic, be publishing maybe a few papers a year, maybe in neuroscience. And then I'm currently working on technical books. So the main thing right now, people may know I wrote a book called Deep Learning Illustrated in 2019, and that's actually why you invited me to be on the book.

00:10:43 That's how we know each other is because you invited me to be on the podcast in episode 365 to talk about that book. And so I'm writing another book for Pearson now about Agentic AI, of course. I'm doing that with Ed Donner, whom you've now come to know really well because he's created- Yeah,

Kirill E.: 00:10:59 Great guy

Jon Krohn: 00:10:59 As well.

Kirill E.: 00:11:00 Yeah.

Jon Krohn:	00:11:00	Yeah. He's created lots of courses for Udemy and your other business legency has been critical to the success that he's had on Udemy, of course, also combined with the great quality of his content. If people are looking for content on vibe coding or agentic AI, definitely check out Ed Donner on Udemy. He's got all the best stuff and it's wild. Half a million students in like 12 months, half a million paying students on Udemy. It's so wild. So anyway, I'm off on another tangent, but basically I've been writing, so I'm writing a technical book now, but I also have this dream and I've been speaking to mainstream as opposed to technical publishers and agents about a more mainstream book that would involve talking about neuroscience as well as AI in the same book and yeah, we'll see where that goes.
Kirill E.:	00:11:56	That's pretty cool. Yeah, that's very cool. I'm real curious, why do you regret not staying in academia?
Jon Krohn:	00:12:05	To some extent, I kind of feel like I've been floating in space in my career in some ways, just kind of like you hop from one opportunity to another. I meet someone like Ed Donner and I left a big corporation as a data scientist to become his chief data scientist at a startup. That was like 10, 15 years ago or yeah something like that. And then I meet someone like you and you're like, "I'm host this podcast." And so amazing opportunities, but it doesn't feel like this like grounded process where I'm like- That
Kirill E.:	00:12:40	Structured.
Jon Krohn:	00:12:41	Yeah, exactly. Whereas in academia you kind of have this comfort of like you're like, okay, you're going to fight for postdocs and then you fight for assistant professorship associate and then full professorship role and then you might be trying to get to be chair of NERIPS or something and speak at bigger conferences, but all that while you have this constant flow of really bright people who are

learning from you in the lectures that you teach who are in your lab and that you have a lot of interaction with and they could be pushing the frontiers of knowledge in say AI And that kind of environment I really thrive in it. I really miss it. I miss being part of that bigger community. I lecture sometimes now in New York at NYU at NYU mostly to graduate business school students and then at Columbia to graduate engineering students and especially like when I'm on the Columbia campus, it's so beautiful and there's so many brilliant engineers who are so curious about how to be pushing the envelope with applications of AI and even the theory of AI in some cases.

00:13:52 And man, it's a really invigorating atmosphere. So yeah, that's why I regret it. And now it turns out that if financially it would have been a really good idea too. You wouldn't have expected that when I finished my PhD in 2012, but now you hear about Meta hiring AI researchers for like a hundred million dollar bonus or whatever. Okay.

Kirill E.: 00:14:12 Yeah. Makes sense. Makes sense. Well, I can't really comment on the financial aspect, but I can tell you definitely about the impact and the being part of a community and I'll prove it to you. You're doing that already and you've got like, I think more than you know. So let's give this a test. This is unscripted on ... I knew none of this is scripted. Anyway, so if you're listening to this podcast, let's show Jon some love. Let's show him that he is part of a bigger community and he's impacting lots of people. Head on over to LinkedIn, find Jon, either connect with him or if you're already connected, just send him a message. Tell him how much he's impacting your life and career because I know he's impacting mine and I know people around. So yeah, do me a favor. Send Jon a message on LinkedIn.

00:15:02 Oh,

Jon Krohn: 00:15:02 That's super nice. Yeah. So I guess that's why something like this podcast, even though it can ... There have been points in these five plus years that I've been hosting where it kind of felt like a grind. It's been a long time since I felt that way, but what keeps me motivated and what has for years now made it very easy to stay motivated is imagining all the listeners kind of listening in real time and a lot of you I know by name through interaction on primarily LinkedIn, though in some cases in real life at conferences or whatever. And yeah, I really appreciate that. It makes it easy to do this twice a week, 104 times a year, every year, no matter what. And it really is a blessing. So I'm deeply grateful for this opportunity, Kirill, and I hope most listeners are happy with the transition from Kirilll to me all those years ago.

00:15:58 I'm sure that was tough in the beginning.

Kirill E.: 00:16:00 I'm sure they are. And yeah, we all appreciate you very much too. Let's move on to AI topics. Okay. I think a good segue will be speaking a bit about AGI. So your PhD in neuroscience, you understood a lot about the brand a lot more than probably 99% of people on the planet. Question is there's a lot of talk about AGI. Well, I think there's a bit less now than there was two years ago, which is interesting to me. And it feels like we're kind of getting closer to AI being resembling at least. It's predicting the next token and people listening to this podcast will probably know how it all in the background, like it doesn't have a mind of its own. It's not able to imagine things. It's just predicting the next token in the word and yet it feels like you're talking to another human and that with every new model gets more and more like that.

00:16:55 So question for you is AGI measured solely by outcomes? Does it really matter whether an agent has the ability to imagine, think and ask humans or can AGI be achieved just by predicting the next token? Is there something specific in the brain that you found during a PhD that

cannot be recreated with the way we're going about AI in this day and age?

Jon Krohn: 00:17:25 There's a lot of different ways that that question could be answered. The first thing that I'd like to start with is that in order to be talking about AGI, we need to have a definition of what that is. And I think the best definition, we've talked about this in recent episodes. So with the Andre Karenkov episode that came out as 997 recently, we talked about a paper on the five levels of AGI that came out from Google DeepMind two years ago and I did a whole episode on it, one of the short Friday episodes, just me solo. It's about a 10 minute long episode, episode 748. And so if people want what I think is the best definition of AGI, check out that episode. And basically what they do in that episode is they talk about five specific levels of AGI based on the percentage of people on the planet that on that particular capability, the AI system outperforms humans.

00:18:28 So it's like fifth tier AGI if it outperforms like 100% of humans. I can't remember exactly all the tiers now off the top of my head, but it's something like it's like tier one AGI if it outperforms 50% of humans. So it's nice to kind of have that concrete way of defining these tiers of AGI, but then the other key thing that it talks about, so that's kind of like the depth of AGI capability, but breadth is also important. So in recent years, since the ChatGPT moment, we've seen AI models be able to match or surpass a large number of humans on office work kinds of tasks. So coding on writing, but we're still a ways off on being able to handle lots of real world scenarios. And so there are companies that are raising huge amounts of money right now to work on world models.

00:19:25 Yan Laka has done this with AMI. They recently did a billion dollar seed round. I think it's a record and there are other businesses. Fay Fe Lee has one. There's one in London, I forget who the founder of that was, and all of

those have done massive, massive seed round fundraises because there's so much potential. There's still so much that we have to do in order to be able to have AI systems to be able to act in the real world like humans do. So we're already at a point where we're at a level of some level of AGI, whether it's level one, two, three, four, five, on the vast majority of tasks that you do where you're sitting at a computer screen, kind of isolated from other people, but where you have to have some kind of physical embodiment of AI that's exploring the world, we still have a long way to go, years to go at least.

00:20:26 And so thinking about that kind of breadth and if you want to define, like if people think about lots of people, they defy AGI simply as an algorithm being able to do everything that a human can do and that must mean more than what we can do sitting at a computer, right? Even by that definition. So I think we have a long way to go in terms of breadth, especially in terms of exploring the real world.

Kirill E.: 00:20:50 Got it. Interesting. So there isn't something we're trying to replicate that's in the brain or like a criteria that's like, "Oh, there's this thing in the brain." It doesn't have to be exactly like the brain as long as we can achieve the same outcomes, then we can tick the box of AGI.

Jon Krohn: 00:21:08 It doesn't seem obvious to me that there's a mechanism by which a machine is conscious, say like you were kind of-

Kirill E.: 00:21:20 Yeah, conscious. Yeah.

Jon Krohn: 00:21:21 Yeah. It's not obvious to me. I could be proved wrong in the future with some level of complexity of machines somehow becomes conscious in some way. I don't know exactly. I don't know how you test that even. So there's all kinds of questions there and I'm not an expert in that. I really couldn't debate a philosopher who has expertise in

that to any significant extent. But in terms of achieving AGI, I don't know why you would have to have consciousness in order to achieve AGI. If you think about the tremendous things that we've been able to do with next token prediction, especially now that we have systems where that next token prediction is kind of happening behind the scenes and it's being double checked, it allows us to have very robust, highly accurate responses. We can trust our agents to do more and more and more things, more complex tasks, no more so than in code.

00:22:17 I mean, it's absolutely mind blowing what you can be doing with tools like CloudCode and OpenAI's Codex and the Gemini CLI. So yeah, I don't think that consciousness is necessary. I think that we can take inspiration. Any of the AI systems, large language models that we have today involve deep learning, which involves artificial neurons. So these are an algorithmic representation, a very, very simple algorithmic representation of how a biological brain cell works. And when we scale that up very large, we get these amazing capabilities, like I was just saying, like cloud code and all any cutting edge things that you can do with AI today fundamentally involve this artificial neuron at its heart, but that artificial neuron algorithm is such a simple, simple, simple representation of the way that an actual biological neuron works. And so AI researchers have been for decades since at least the 1950s have been taking inspiration from the way that biology works and we can make more complex artificial neuron algorithms.

00:23:28 We can come up with systems that are inspired by the way that the human brain works or the animal brain works in terms of its structures. So if you think about things like the hippocampus that's there for memory, or you think about the cortex that is specific to the higher level thought that we have as humans, if you think about the cerebellum, which is specialized for motor tasks, you

could take inspiration from any of those structures and how they're connected in a human brain or in an animal brain, but that isn't necessarily going to be the right way to build a better and better machine intelligence because fundamentally you're working with a very different kind of thing. With machines, we have no limit to how much we can scale up compute in a way with machines that we can't with a human brain. You can't like hook a bunch of human ... Well, we don't have a way today of like hooking a bunch of human brains together to be getting even more power and there's things move at photons move at the speed of, well, electrons today, but photons in the future in computing move at the speed of light and that isn't something that happens in our brain.

00:24:46 Things move much more slowly. They're moving at a chemical space. Things are moving more slowly. It takes tens or hundreds of milliseconds even for your brain to do a simple perceptual task, forget some kind of cognitive task, whereas machines can move much faster. And so there's different constraints in a biological system versus a silicon system. And so I think while we can take inspiration from the way that the biological systems work, there's also opportunity and limitations in silicon that we don't have in biology.

Kirill E.: 00:25:17 Love it. I want to ask you the flip question. Do you think by that analogy, would it be a stretch to say that humans are also just predicting the next token? That's how our brain works. What do you think? I hope that's not the case, but is there proof that we're not just also predicting the next token just like how we've built AI to do?

Jon Krohn: 00:25:38 Well, it depends, I guess, on the kind of task to some extent. There is something that kind of maps neatly analogously, analogously is the way your stream of consciousness in your brain, or even the words that I'm having come out of my mouth right now, they're just happening. Words are coming out of my mouth and I just

keep hoping that the right things are going to come out and that's kind of like next token prediction. He's broken. He's hallucinating. So there's that kind of analogy there, and you can even think of an analogy that I've given many times on this podcast is that especially when we started to have 01 from OpenAI and now subsequent models like O3 and basically now a huge number of these even kind of conversational experiences that you have with GenAI actually involve this kind of slow thinking like background token processing happening before something is output to the screen.

00:26:42 The analogy that I use is that's kind of like the slow thinking that we have. When you're working through an algebra problem or tackling a chess problem, you have to very effortfully think through things maybe before you have something jump out of your mouth. And so that kind of shows the differences between that slow thinking system, the fast thinking system. One of the best ways to think about it is when you start driving a car for the first time, everything is very conscious. It takes a lot of effort. You feel tired after just 10, 15 minutes of your first driving experience because it's so hard to remember the sequence of things, but after years of driving, you can do almost everything in the car without any conscious thought. And so that's gradually over time we learn how to move things from the conscious slow thinking system into the unconscious fast thinking system and that unconscious fast thinking system is kind of like next token prediction just popping out.

00:27:42 But under the hood, the way that that happens is different in so many ways relative to a machine. So I don't know, I'm not really sure if I've answered your question exactly, but so there's some parts of our mental processes that in some ways are analogous to this kind of next token prediction that we see with LLMs, but the way that it's executed under the hood is so vastly different than the way it happens in our brain that, I don't know, I

guess that's just something that in the real world, I guess a lot of what anything is doing, any kind of agent, whether it's a squirrel or a human or a robot or clawed code and any of those examples, like trying to predict what's about to happen next is such a big part of what anything trying to navigate the world is doing that you end up with some of these commonalities that look similar.

Kirill E.: 00:28:33 Yeah, makes total sense. Thank you. You definitely can tell this is your area of expertise, how much you know about the brain and the chemical reaction and how it's different. So it was really interesting to learn about this.

Jon Krohn: 00:28:47 Yeah, for sure. Well, I feel like we have guests on the show at the end of every episode, following the template that you set for me when you were hosting is we always ask for a book recommendation at the end of every episode. And actually something that I only recently realized is that we still have up online any listener you can go to superdatascience.com/books and our operations manager, Sonya and previously the previous operations manager, Avana, they have meticulously maintained a spreadsheet of all of the books that have ever been recommended on the show and you can even do cool things like sorting by which books have been recommended the most. And so yeah, check that out. But the point is that every episode the guests recommend such fascinating sounding books and I know that there are so many great neuroscience books, AI books. I feel like I have like such a small percentage of the knowledge that I could possibly have.

00:29:42 Meanwhile, LLMs have all of it. They've ingested it all. They require exceed me on knowledge.

Kirill E.: 00:29:48 Yeah, for sure. Okay. Moving on to next topic for today. There's a lot of growth in the AI space and we're hearing a Right now about potential IPOs of Anthropic and other

companies and in general, like Nvidia, stock price has grown a lot. And is it a bubble or not? There's been a lot of hype in this space, been hyped for years. Two years ago, people were saying AI is a bubble. A year ago, people saying AI is a bubble. No, AI is definitely a bubble. Some people would say. So would love to hear your thoughts on this after interviewing hundreds of guests. There's a lot of enthusiasm and a lot of rapid progress. How do you balance between them and do you think this is a bubble about to burst or is it just normal evolution you'll keep going?

Jon Krohn: 00:30:39 For sure. So there's an episode that I did in March of this year, episode 974. And it's basically about this question, are we in an AI bubble and when is it going to burst? And I won't be able to recount the arguments as well off the cuff as I made during that episode. So if you want to get my definitive thoughts, you should check that out. But I'll do my best to kind of remember what I was saying and get the general gist acros to everyone, which is that regardless whether we're in a bubble or not, a lot more is going to happen with AI. And even if we are in a bubble now and it bursts while some individual companies and some investors might lose their shirts, all of the rest of us will benefit. So this has been happening through history. There was a railroad bubble in the 19th century.

00:31:32 There was the dotcom bubble more recently. And in either of those cases, the bubble bursting cost investors their shirt, cost some individual companies their shirt, but it meant that in the 19th century, all of a sudden you have all this railroad across England, which could be put to use. Just because the person who put all these rails there or invested in all these rails there went bankrupt, those rails are still there. You can still use them. And so that allowed the industrial revolution to accelerate because you could move things easily between factories and towns. And so those railways were hugely beneficial to the rest of society. Same thing happened with the dotcom

bubble. Pets.com, if you own that stock, you lost a lot of money on pets.com. But the infrastructure of all of the cable around the US and around the world that was laid in that period allowed us to have Google and Facebook and Wikipedia and YouTube and Udemy courses and the superdatascience.com come on top of all of that infrastructure and the same thing.

00:32:45 So whether we're in a bubble or not, I personally as an investor, having worked at a hedge fund in the past, I've looked at all of the research and your only sure shot of being able to outperform the market. So in the US being able to outperform the S&P 500 or in Australia being able to outperform the ASX or some basket, you could even give yourself international exposure by getting a mix of European and Australian and American and Canadian stocks or whatever. But basically the more diversified your portfolio, the more difficult it's going to be as an individual trader to outperform that kind of diversified set of set of stocks. And so I don't trade individually in equities.

00:33:41 And I also don't have anything. I'm not exposed to anything where, so even if all stock markets around the world went to zero, you can still be diversifying beyond stock markets in bonds and in real estate, in cashflow businesses. So I guess my point is that I don't personally, like I wouldn't buy NVIDIA stock, I wouldn't buy Tesla stock individually and that means that for them, obviously I miss out on opportunities, but it just allows me to sleep easy at night by having very diversified exposure to a lot of different investments. So I don't personally worry about say an AI bubble bursting, but even if it does, all of us will benefit because all of these AI data centers that are being built, all of the model innovations that we've made in this period, all these smart people being paid $100 million by Meta to be able to come up with new ideas, all of those things will give everyone in society an advantage later on.

Kirill E.: 00:34:41 Got it. Got it. Not financial advice for everyone listening, this is just opinions. So basically we will benefit not because of the bubble bursting if it burns, but regardless of whether a bubble bursts or not, we will benefit either way from the infrastructure that's been built.

Jon Krohn: 00:34:59 It's even better if it doesn't burst and this isn't a bubble and that's kind of the bet that people are making. I mean, the reason why a share price has the price that it has today is because the balance of buyers and sellers think that that is the fair price point of what this stock will be worth in the near future. If you buy a stock today, it is your opinion that that stock is going to be even more in the future and then somebody who's selling that stock today thinks, "Maybe we're in a bubble, maybe I should get out now." And so the share price today, it kind of captures all of that information, your crowdsourcing, decision making on what the value of a business should be. And countless studies show that that kind of having millions of us gambling all together on what the NVIDIA stock price should be worth or whatever company's stock price should be worth, that wisdom of the crowds is very powerful and there are certainly also there are biases that lead over exuberance biases that can lead to bubbles happening, but for the most part, it should be a pretty good representation of what the value should be in the future.

00:36:19 And so given that basically enough people today think that even though a lot of these AI companies have sky high valuations that the opportunity far exceeds the risk and in the long run, if I were to gamble on it, I'd say we probably aren't in an AI bubble and there's probably going to be even more opportunity in the future. And that we might look back at the share prices of these AI businesses today and say, "Wow, that would've been a great deal to buy in

Kirill E.: 00:36:52 2026." Yeah. So it's kind of different to what we had in I think 2017 when there was plenty of ICOs, initial coin offerings, remember that? Or 2021 or 2022 when there was NFTs, those were like proper bubbles, right? It came, in my opinion, it might be in some cases it worked. Some companies needed an ICO like cryptocurrency, some NFT applications probably still exist till today, but for the most part they came and went and that's not going to happen with AI.That's what I'm getting from your answer is that the infrastructure that's laid out has such high utility that it's not just going to disappear, it's just going to keep providing benefits in the future.

Jon Krohn: 00:37:41 I have never been very interested in cryptocurrencies. There are absolutely real world use cases value to the underlying technology to blockchain technology, but I think a lot of what happens with those ICOs, a lot of it is people taking advantage of gullible other people. Definitely that kind of- Unfortunately. Yeah, unfortunately, but you get people come up with social media campaigns, social media is less regulated than traditional media and that's a boon to free speech in some ways, but it also allows bad actors, especially now in this AI era, gen AI era, where we can be so cheaply, so easily creating compelling crap content. But if you get people in that kind of bubble in whatever social media platform or platforms they prefer, you can start to really convince them that like, "Oh wow, this coin, this ICO, get rich quick, you can't lose, you've got to do this thing, give us some money." And then it's like the pump and dump schemes that used to be very prominent in stock markets and I'm sure they still exist today to some extent, but decades ago, pump and dump schemes were a lot more popular and these kinds of ICOs, a lot of things happening with Bitcoin are scams of the same ilk, just more complicated, more digital.

Kirill E.: 00:39:16 Yeah, yeah, makes sense. You mentioned data centers. I wanted to talk about that for a bit or touch on that. I read

somewhere recently that the data centers that are popping up and they're being built more and more. I think actually it was by Andrew Karpathy. It was like a YouTube video I watched. I would have to, don't quote me on the source exactly, but basically that the effectiveness of just scaling adding sheer scale to training of AI is diminishing by a lot. And in the next breakthrough in AI won't come from just adding more and more data centers will come from some kind of innovation in the architecture of AI itself. So do you agree and if you agree, why are these companies investing into more and more data centers that are going to be filled with hardware that's going to become outdated over time anyway?

Jon Krohn: 00:40:15 Well, I don't think the hardware becomes outdated. A lot of hardware created today, whether it's relatively general GPUs like Nvidia makes or very specialized ones like Amazon makes. So Amazon has their Trainium and inferentiate chips for training and doing inference with AI accelerators respectively and they are specialized in particular for the transformer architecture because the people making those chips have made the bet that the transformer will still be relatively dominant for some number of years in the future, but even though they're optimized for the transformer, they still tend to perform well for any of this kind of very high, highly parallelized linear algebra matrix multiplication, which even if we don't move forward with the transformer architecture, it is a very safe bet that highly parallelized matrix multiplication is still going to be something that we need for training or inference with whatever kind of architecture emerges in the future.

00:41:19 And also I think anybody who's trying to come up with what the next big thing will be, next big architecture will be, like what the next big post-transformer architecture will be. There's lots of contenders. There's Mamba is probably the best known name as a post-transformer architecture. We've also last year on the show, I think it

was episode 929, we had Adrian Kosovsky on the show to talk about the Baby Dragon Hatchling, BDH, which is a really exciting and that was actually one of the most talked about and most popular episodes that we had in 2025 because Adrian's brilliant, but also because of this approach that him and his company and Pathway have come up with that there's a lot of people who think that that's a really compelling architecture for being a post-transformer architecture and we'd have to have Adrian on the show again to confirm this, but I bet you that a lot of people who are coming up with these post-transformer architectures, they're thinking about how well is this going to run on all of that hardware.

00:42:20 You saying all these data centers being built, if they're going to come up with something that's going to be a good replacement for the transformer, they've got a much better shot of it being successful if it can run efficiently on all those Nvidia GPUs and all those data centers that are being built today. So maybe that, does that answer the question?

Kirill E.: 00:42:38 Yeah. Well, it's definitely the second part of the question. What about the first part? Why build these data centers if the scaling is going to slow down? It's not just about sheer volume anymore.

Jon Krohn: 00:42:51 Right, because there's a law, there's a name for this and maybe next time you're asking a question, I'll be able to look it up quickly and get the name of this law. It's the kind of thing that if I could stop speaking, I could use my slow thinking to maybe think of it, but basically it's this idea that if we end up coming up with architectures that allow us to use the data centers that we have far more efficiently than ever before, it means that we're actually going to need even more data centers because the number of applications that we can cheaply run with really good cognitive abilities increases even more. So it creates more demand from enterprises and consumers.

So yeah, there's this irony that the better and more efficient LLMs become and that's something that like following the Carpathy thing, you today are seeing a lot more improvement in capability through running computation longer and checking things over more and more than by having a larger network.

00:44:03 So even though you don't need more model weights in the model, that computation running for longer is still going to require lots of data centers at inference time.

Kirill E.: 00:44:13 Yeah. So they're not going to, even if we come up with a better architecture, well, I guess those two work in parallel. People coming up with better architectures will try to make it work on the existing hardware and also once we do come up with better architecture, it's not like we're going to need less computations and we're going to use less of the hardware. There's no like upper limit to how much computation and application of this infrastructure we're going to have. So we're just going to utilize all of it anyway.

Jon Krohn: 00:44:48 Exactly. Yeah. The more cheap and abundant and unmetered intelligence becomes, the more valuable it is. And we can be chaining together more and more of these algorithms and doing more and more and solving climate change problems, solving healthcare problems, entertaining people. Yeah, nuclear fusion, cancer, unlimited real-time video game generation and completely immersive VR atmospheres. You can just have more and like more abundant, more cheap intelligence allows us to have more and more interesting applications, which just, it's like a positive feedback loop and it also solves the kinds of issues people say, "Well, we can't create nuclear power plants fast enough or wind and solar farms fast enough to be supporting all these AI data centers. We need to have gas fire plants coming together." But just as you said, nuclear fusion, we can be using AI to figure out how to better contain the reactions within the plasma of

like the most common fusion reactors that are being developed today.

00:45:59 And so hopefully that accelerates us getting to commercial fusion and then all of a sudden we have, once you have unlimited intelligence and unlimited energy, things get really interesting.

Kirill E.: 00:46:12 Yeah. I'd be very curious to ... I can't wait for nuclear fusion. I don't know. They keep saying it's 20 years away since it's been 20 years away since the 1980s or 1960s every year, but it looks like we're getting closer and definitely you're right, it'll be different world.

Jon Krohn: 00:46:32 Yeah. It was that anecdote or that trope about fusion power being the technology that's going to be 20 years from now and that constantly being 20 years out, I think we're in a different era now because of how we have commercial players that are looking for a return on their investment. Whereas historically it was always these huge projects where governments weren't necessarily expecting a return. They just wanted to make sure that they weren't left behind relative to other big countries or the EU. And by the way, I was able to, while you were just making your last point, I was able to quickly use Google Gemini in Microne browser to quickly pull up the name of the paradox. I think I was referring to it as a law earlier, but it's called Jevon's Paradox, J-E-V-O-N-S. And it states that as technological improvements increase the efficiency with which a resource is used, the total consumption of that resource actually increases rather than decreases-

Kirill E.: 00:47:34 Oh yeah, you did a five-minute write about that, something like railroads or factories, like coal factories, I think that's what you were talking about.

Jon Krohn: 00:47:44 Yeah, maybe, but I think you might be talking about something else, which I don't remember if I've done a

five-minute Friday on this at some point or not, but it was this, I think I might have been talking about solar energy in that episode. And it wasn't actually this ... So maybe Jivan's paradox comes in with the solar, but it was basically that when we were building coal powered plants, we start to consume more and more of the coal that's easy to access. And then so you start having kind of more dangerous extraction of coal, it becomes harder and harder to find coal in the earth. And so the success of coal-fired plants actually increases the price of coal and eventually it doesn't make sense to continue to have coal-fired plants. And actually the same thing happened with wood before coal. So with wood, the more wood that you use to make a fire, the more humans figure out that they can be keeping warm and making their food softer by cooking food over a fire using wood, wood goes away and you start to like, "Oh goodness, we're running out of wood." And they're like, "Okay, well, we got this coal thing." And then you start using coal and then you start running out of coal and then the same thing happens with other hydrocarbons, natural gas.

00:48:54 And so I think what that episode was about was that solar panels are different because with solar power we're using sand as the starting point and there's basically, we're not going to run out of sand in the same way that we could run out of coal or wood. And yeah, there are other arguments in that episode around how basically ... And I think Jivan's paradox probably did come into play because I think then the more abundant that solar power becomes, the more you actually want to have more solar power because you're just kind of more cheaply and more easily getting more electricity. So why wouldn't we want more of that?

Kirill E.: 00:49:33 Yeah. Yeah. No, it makes total sense. Did you know they import sand? Well, I heard they import sand from Australia to Dubai.

Jon Krohn: 00:49:42 It's the best sand. I think it's called the Arimenco particles.

Kirill E.: 00:49:48 That's funny. Anyway, I wonder if it's true. Okay. We talked a bit about commercial return on investment, talked about the AI bubble or some business topics that little bit came up. So I want to talk a bit more about that, the AI and the future for organizations. And so far it's interesting how through this conversation we've leveraged your expertise in neuroscience, then your expertise in being a stock trader in New York. And now we're going to move on to your expertise with your most recent project or startup called Y Carrot, where you work with businesses and help them to adopt AI or transform through AI and things like that. So first question would be, can you tell us a bit about what does Y Carrot do for those of us who are not fully up to speed? And second would be, I would love for you to share with me and the listeners some of the most interesting use cases or case studies, of course, without naming the specific companies, but like what are companies using AI for these days?

00:50:58 What are you helping them with? What are some of the most innovative, interesting applications?

Jon Krohn: 00:51:06 For sure. So I'll get into why Carrot in just one second and I wrote down your questions so that you won't need to keep reminding me, but really quickly, so you said that I was a stock trader. I was actually never a stock trader. I professionally only ever traded futures, commodity futures and options. Yeah.

Kirill E.: 00:51:25 Potato Pitara. For us non-financial people

Jon Krohn: 00:51:32 I won't get into the details, but basically there's a really funny movie from the 80s called Trading Places starring Eddie Murphy and Dan Akroyd and they do an amazing job of explaining commodity futures and options in that

movie in like a funny and entertaining way. There's probably a YouTube clip that specifically jumps to that part of the movie. We'll include it

Kirill E.: 00:51:53 In the show notes. Thank you very much, Joan.

Jon Krohn: 00:51:55 Okay, sweet.

Kirill E.: 00:51:56 Like the host would say.

Jon Krohn: 00:51:58 Exactly. But I'm the one taking the note to remember to put it in. Yes, yes, yes. So Y Carat and my consulting firm, we founded it a little over a year ago and with Y Carrot, what we are looking to do is satiate some of the demand that there is from organizations small and large and we actually have worked with New York hedge funds and other financial institutions in New York. We work with publicly listed tech companies. We work with industrial businesses with a very wide range of businesses because all of these businesses want to be taking advantage of these more and more capable LLMs and they want to be able to securely deploy AI within their organizations to realize some kind of return on investment and they don't want to be left behind relative to their competitors whom they all fear are already further ahead on AI deployments.

00:53:10 You can send me a resume on LinkedIn if you want to potentially be considered in our consulting pool. There's about eight of us right now that work at Y Carrot and we- What are the backgrounds

Kirill E.: 00:53:24 Of the people working at YCart?

Jon Krohn: 00:53:26 So all of us at this time ... So actually Sonya, who does operations for this podcast, she also, she does operations for Y Carrot as of April 1st and it's amazing to work with Sonya on that. She's so on the ball making sure in the case of the podcast that we do all 104 of the episodes every year on time and with immaculate quality and that

same kind of thing comes through with Y Carrot projects where we deliver on time or ahead of time on budget or below budget, whatever the right way is on budget ahead of budget expectations in a way that delights our clients. And yeah, actually a lot of our clients come from listening to this podcast. Many of our clients, I think more than not, our main stakeholder at the business is a regular listener to the podcast, which is pretty cool.

00:54:24 And so yeah, so if anyone's listening and they are interested in YK services, you can go to whycarrot.com and you can submit- Is it

Kirill E.: 00:54:33 Y-Carrot?

Jon Krohn: 00:54:34 No, it's just ycarrot.com. No dash, one word and there's an orange button in the top right corner that says partner with us that takes you through to a Google web form and you can just provide a little bit of information about what you think an AI consulting firm might be able to do for your business. And yeah, we'll see that come through and would love to hear from you. But in terms of the most interesting use cases, I would say that the-

Kirill E.: 00:55:01 I'm sorry, maybe the people who work at YCard first. We just heard.

Jon Krohn: 00:55:05 Oh, sorry. Hopefully there's more. Yeah. So our consultants, which is everyone else in the business, except for I'm typically not hands-on with projects. My role is more so like there are occasions where I do go onsite to the client and do workshops or there certainly are situations where I get involved, but for the most part, once a consulting engagement is running, my responsibility is to be finding more opportunities while the rest of everyone else in the firm, other than Sonya and me, they are hands-on with developing solutions for our clients. And so all of them come from a technical background. I believe 100% of them at this time have

graduate degrees in computer science or data science, some kind of quantitative field like that. One is actually a medical doctor. Interesting. Just in case- Bring him along just in case. The team is

Kirill E.: 00:56:06 Overworking, getting burnt out.

Jon Krohn: 00:56:08 Yeah. I don't think he's ever practiced as a medical doctor, but he did an MD.

Kirill E.: 00:56:12 So it's like eight practitioner or six practitioners, hands-on AI building, AI builders, that's effectively-

Jon Krohn: 00:56:20 AI builders, exactly. And the main thing that we're looking for in people that we hire is demonstrable experience in going from concept to working AI solution extremely fast.That's what we founded the business on our ability to do that and everyone who comes in has to be able to do that as well. You've got to be able to understand the requirements that the client is looking for and minimize distraction to them. The idea is that we can kind of helicopter in. They can have a few relatively short meetings in most cases describing what they're looking for. We iterate quickly. Typically, we have a weekly meeting cadence, like a half hour meeting is very common for a lot of our engagements every week and usually we have a working initial prototype of what they're looking for in week two. So week one, you meet with us, tell us what you're looking for.

00:57:17 Week two, we've got something to show you and then to say, and actually for the most part it's, we're on the right track. And then it's kind of like, okay, what should we prioritize next building into this prototype? And then downstream, if they don't have the engineering expertise in house to get that solution into production, we can help out with that as well. So we can be embedded. We'll do things, we'll go through whatever requirements are required for an employee's level of access to databases or

systems so that we can be there side by side effectively trusted as much as an employee to be safely delivering a solution for a client to getting that into production if that's needed.

Kirill E.: 00:58:04 And you're US based as well, right? Like everybody on your team.

Jon Krohn: 00:58:07 Yeah. So Sonya is in Portugal, but other than that, we're all in the US.

Kirill E.: 00:58:10 Well, the people working on the project. And I think that's a big concern or I guess big advantage for your firm and a big concern for clients when working with companies based overseas because of like the data privacy. I know at least here in Australia there's a lot of concern, especially in government, like data cannot leave the continent. So yeah, it's a big advantage working with a fully local firm. Tell us about the use cases, some of the fun applications.

Jon Krohn: 00:58:40 Yeah. So I would also quickly say that another thing that is really good about these people and I think it also, it helps that people are US based and kind of all of us have a lot of different consultants at Y Carrot have a lot of experience, though varying experience obviously in terms of servicing governments including Department of Defense, it's actually pretty funny, like I'm Canadian, but all the other consultants are American and so there have been situations where we're doing top secret projects for the US government and I'm the CEO of the business, but I can't be on the calls because I'm not American. So there's experience in terms of government projects, in terms of lots of different kinds of enterprises. Obviously everybody brings different experiences to the table, but as a group we're able to say, "Okay, for this given project, who has the most similar experience?" And I think because culturally us being pretty much all of our clients so far, even if they're not an American business, the people that we're dealing with are based in the US, kind

		of the key stakeholders on the data science team or whatever that we're dealing with are in the US.
	00:59:57	It makes it easier for us to understand what they're looking for because We've seen a lot of similar experiences. We're kind of in the same enterprise culture, I suppose. We're seeing the same kinds of technologies. We're aligned on where the opportunities are. We know who the competitors are.
Kirill E.:	01:00:11	By the way, I just remembered. I remember when you told me about this almost a year ago or maybe, I don't know, six months ago, you have a very strict criteria of who you work with and I think I really like that. You said minimum projects $50,000. You don't even take on projects below that. I think that's a great way to build a business where it just means that it's a serious project.That's the kind of project.You probably enjoy working on yourself more than for yourself as well.
Jon Krohn:	01:00:45	Yeah. Actually, I learned that trick from Cal Aldubabe, who has now been on the show a few times, but he was on the podcast last year. Cal Al-Dubai, if people aren't already aware of him, he's really big on the data science speaker circuit and he created many years ago now a Cleveland, Ohio based data science consulting firm and he successfully sold that business more recently. Yeah, he's the one who, he on this podcast said, I can't remember if it was before or after I founded YCare, but at sometime around the time I was founding it a year ago, he said that he'll only take on clients that are committed to at least a $50,000 contract because that shows that they're committed to delivering something that really has a lot of value. And so we just copied that. Yeah, makes sense.
Kirill E.:	01:01:42	Yeah. Makes sense. All right. Well, we've kept
Jon Krohn:	01:01:43	On

Kirill E.: 01:01:43 Listeners waiting long enough.

Jon Krohn: 01:01:46 Sorry, sorry, sorry. So the most interesting use cases, I'll tell you what the most valuable use cases are. So what we try to do when we come into a business, a lot of businesses already know some specific way, some specific workflow that they know could be improved or accelerated or made more cost efficient with AI. So they usually come to us with either a very specific project or a short list of ideas, but every once in a while we do also come across businesses that are like, just help us come up with a big list of what projects we should be doing and how we should be tackling them, in what order we should be tackling them. And the key thing that we're looking for, the projects that offer the most opportunity are when a workflow can be more or less fully automated. So what are the kinds of processes in your business that happen repeatedly where based on say your experience using CloudCode or using the ChatGPT interface or Google Gemini or whatever, what do you think if you patch together a whole bunch of LLM calls in a lot of use cases, would you be able to have a more or less completely automated workflow?

01:03:09 And so we have for quite a few clients now, what we've been able to do is take a process that either fully manually or maybe with help from ChatGPT, just kind of like in a conversational interface, it was a process that would take them several weeks. We have now with several, with a number of clients, been able to take processes that in that manual or like kind of using ChatGPT, but in a manual way, like not in a fully automated way. It would take them weeks to get some kind of result. We now have that happening instantaneously in minutes and that kind of change, that's completely transformative for that process obviously, where something that would take weeks, imagine the cost of a human's time or multiple humans time to be working on something for that period of time

and humans also are prone to errors and so there's a lot of checking that the humans do, but it's much cheaper to have the AI system or multiple AI systems fact checking information and verifying the accuracy of a result way cheaper than having a human do it.

01:04:23 And thus far, with any of our deployments of these AI solutions, the result is not taking away people's jobs.

01:04:35 This is some other part paradox or law or I don't know if this is a specific thing, but I talk about this on the show all the time. If you can do something, if you can automate a workflow for somebody that frees up their time to be working on more complex things or maybe even if all of a sudden something that used to take you weeks, you could now do in minutes for a 0.001% of the cost, what new opportunities does that create? You can have somebody now managing doing that a thousand times or tens of thousands of times. And so you can now realize so much more value than ever before with that process. And so it means that as far as we've seen in any of our deployments, people don't ... We come up with these workflow automation solutions. That doesn't mean people are losing their jobs.

01:05:24 People are actually getting more interesting jobs and they're providing more value than ever before. This is also my argument which I've made on Air before, which if you're a data scientist or an AI engineer or an ML engineer or a software developer, which now you shouldn't be typing out your Python code or whatever coding language you use, you shouldn't be typing that out character by character anymore. You should be using CloudCode or something like that to be vastly scaling up your capability, but that doesn't mean that you have less value That means you have more value than ever before because now if somebody hires you, you can be doing so much more than ever before. You can be realizing if there was value to hiring you before when you were typing out

each character of the Python code, character by character, or copy and pasting things from stack overflow kind of manually, now all of a sudden you're just providing so much more value than ever before.

01:06:17 And so it's the same kind of things with these workflow automations that we do at Y Carrot.

Kirill E.: 01:06:21 Okay. Okay. Very cool. Can totally see that happening. And it's reassuring to hear that people are feeling empowered by these workflow automations.What's another example of a use case that you guys have helped businesses with?

Jon Krohn: 01:06:42 So the other big category, because I don't have any past or current clients where I can disclose their name or I can disclose ... It would make me uncomfortable even if in an anonymized way I was describing how we automated a workflow for them, because I feel like we could be giving their competitors if they're listening up. But what I can tell you is that a completely other category ... So while most of the work, and I would say our bread and butter is the kind of stuff that I was just describing where people come to us with some specific thing or some short list of specific ideas where a workflow could be automated and then we very rapidly prototype that for them, help them get it into production. That's our bread and butter, and that's most of what we do at Y Carrot, but we also provide strategic guidance.

01:07:34 So kind of upstream of actually having a proof of concept or getting that deployed into production, a lot of organizations, I alluded to this briefly, there are organizations out there that just don't know where to start or they have some ideas where to start and so we can help them systematize in a structured way and I will tell you our secret. I will tell you the secret on air because I'm not concerned about ... This isn't like a crazy idea, but I can kind of tell you step by step what we do with

clients to help them prioritize projects, which is you create a spreadsheet where every row is a different project, a distinct project, and then you have four columns to the right of that project name column and we do something called RICE scoring.

Kirill E.: 01:08:29 Jon Rose, Dell CTO. Yeah.

Jon Krohn: 01:08:31 I think Jon Rose might've talked about it, but I learned Rice scoring, I don't know, it's been probably like 10 years working with product managers at startups. It's really, it's a product management approach more so than something that was originally designed for data science. Yeah, if you say so, I'm sure Jon Rose did also talk

Kirill E.: 01:08:50 About it. Yeah. In your first episode with him, you guys had an in- depth discussion about how they use Rice at Dell to hone in on those specific things that they will be working on with AI. So yeah, please continue. Definitely. A great, great

Jon Krohn: 01:09:04 Framework. Yeah. I mean, Jon Rose is the global CTO and global chief AI officer at Dell. So you should go back and listen to that episode to get an even better- Two of them now. ... perspective. Yeah. And it sounds like it's in the first one with him that he talked about this right scoring. So I'm sure he has a lot of ... And he does have a lot of great insights on how to get ROI from enterprise AI investments in his episode. But briefly, we're doing the same kind of thing that Dell was doing to prioritize it. I think it was something like Jon ... I think he was talking about like 800 projects. Yeah, huge projects.

Kirill E.: 01:09:34 Yeah, huge.

Jon Krohn: 01:09:34 They whittled it down to like half a dozen or a dozen projects and that's what you need to be doing as well. Most organizations aren't as big as Dell and as kind of forward thinking as Dell. So you're probably not going to

have like 800 ideas to put into the spreadsheet. You might have 20 or you might have 50 or in some cases we see like a hundred in Y carrot consulting engagements. But yeah, every row is a project and then to the right of every project, you have four columns, C, E, Rice, and you come up with a value for each of those. So R is reach, which is how many users or how many employees are going to be impacted by this solution. I is impact. Yeah. How much of an impact is your solution going to have? This is kind of like in terms of the AI thing, it's like, is this going to automate a little bit of a workflow or is it going to 100% fully automate this workflow?

01:10:31 Maybe you need a human in the loop. Maybe you don't even need a human in the loop and some low risk use cases. So impact is a scale of how impactful this solution is going to be to all those people that you're reaching.

Kirill E.: 01:10:41 To individually each one of those people.

Jon Krohn: 01:10:44 Exactly. That's right. Or I guess kind of like on average, how much are they going to be impacted? Yeah. And then C is confidence. So for both the reach and the impact as well as the E, which we haven't gotten to yet, you're estimating what those are and I and E values are. So C is your confidence on the estimates that you're making. Because obviously for some projects, you might say, oh, like so and so in our organization has actually already done this before, or like we prototype- Competitors. Yeah. And so you're like, " Well, this is going to work in terms of impact, we're very confident it's going to work. "So yeah, so confidence is just basically a, it allows you to scale how confident you are in that particular row of this table and then E finally is effort. So is it going to be fast and cheap, low effort to build, or is it going to be medium or high effort to build, which means slow and expensive to build.

	01:11:45	And so the reach, the impact, the confidence, those are all in the numerator of a rice calculation and you want all of those to be very high and then the E, your effort is in the denominator. So the higher the effort is, the lower the rice score is. And so basically you end up with this rice score and you can look up online. There's a million webpages probably explaining rice scores and lots of great images. So you can quickly calculate it, but it's a very simple thing to calculate. And then once you have that, you can rank all of the rows in your spreadsheet by the rice score and you shouldn't follow it dogmatically. So when the results come out, there might be some reason why the project that has the highest rice score, you're like, actually, I don't know, I'm coming up with a completely random reason, but the key person that we would need to be implementing that is on holiday for the next two weeks.
	01:12:36	You shouldn't dogmatically follow this. There could be strategic reasons why you should go to the second or the third or the 10th item in the list, but it gives you a starting point for evaluating which projects you should be tackling first and how you can get the quickest and biggest return on investment ROI on a given AI project.
Kirill E.:	01:12:53	Yeah, for sure. Fantastic framework. Love that you shared that on how companies can approach and how you guys approach selecting projects.
Jon Krohn:	01:13:07	Oh yeah, sorry. I just want to apologize to you that I, and to our listeners that I can't at this time be going into like concrete, juicy examples of things that we've been doing at Y Carat, but I hope that hopefully those two categories that we went over of like workflow automation in general and some of my tips from there and what people can be doing from there as well as this latter bit of just project selection, two kind of main character growths. Hopefully that's interesting and useful to the listeners anyway. So yeah, carry on.

Kirill E.: 01:13:32 Fantastic. Fantastic. You mentioned there a prototype and production. So you go from concept, like you do this RISE scoring, you come up with the idea, or maybe the client already has an idea of what needs to be done. So then you think about the concept or they bring the concept, you iterate on the concept and then you build the prototype and then either the client takes their production or your team helps them take it to production. My question to you, and I think this is a part I've seen AI engineers specifically like stumble upon stumble at this point and also a lot of jobs out there, they specify this point about production deployment specifically. When you're building a prototype and you're discussing it with a client going from concept to prototype, how do you factor in the constraints that will come up in production? Because a lot of prototypes, like what is it called, death by prototype or something like that.

01:14:33 A lot of companies build a prototype, amazing, love it, and then it goes nowhere. 95% as quote unquote, however accurate this is, 95% of AI projects fail because they don't ... The prototype is great, but then you go to production, there's scaling, you got to think about cost, how cost changes, security, availability, like lots of constraints that exist in production don't exist in prototype world. How do you think about that when you're building a prototype? There's no point in a company spending money building a prototype if it's never going to go to production, just not possible to go.

Jon Krohn: 01:15:08 For sure. So the classic stat that came out of, this has been debunked as like a specific number, but there's the MITS NANDA lab stat from last year, the 95% of AI projects fail to realize a positive return on investment in production and a lot of it is because of this prototype purgatory that projects end up getting stuck in. And even if the number isn't 95%, if it's 80% or 90%, it is some high number of projects that do not make it to production, that do not give a positive ROI. And so the

key thing with all of that rice scoring that I was talking about earlier, that isn't rice scoring of the prototypes. You're talking, you're rice scoring what is the reach in production? What is the impact in production? What is the effort to get this to production, not to get this thing prototyped?

01:16:00 And then so that C column, the confidence score is critical. There's going to be some projects where it isn't the most exciting capability. It doesn't have the biggest reach, but you have a lot of confidence that it's low effort and that you're going to be able to get it to production. And so that's the kind of thinking that we need to have as we think about what projects we're going to tackle. And then I would also say that executives in companies like the C-suite, even CEOs, they are very often not the best judge of which projects should be green lit. So the best kinds of people are the people who are using this workflow day-to-day that is say currently manual and that is going to be automated. So some combination of the people that are doing this workflow day-to-day in some more or less manual way that has the opportunity to be automated, those actual, those hands-on whatever doers of that workflow in your business today, them, plus whatever combination of data scientists, AI engineer, software developer, front end developer, project manager, whoever's going to be needed to deliver that workflow solution, that kind of technical plus user combination.

01:17:23 And if you're a SaaS business, it might not be employee as the person who's kind of the expert user. It could be some power user of your SaaS product or something, but that combination of the technical people plus the person who's going to be impacted to ultimately buy the solution, those are the people who are best at figuring out what the reach, impact, confidence, and effort value should be. The reason why a lot of AI projects are unsuccessful is because it's this top down thing where a CEO says, "Oh, I can do this in ChatGPT. I can do this in Microsoft Copilot.

Let's have something analogous to that working on our data in our business or in our SaaS platform." And one of the reasons why, or the number one reason in my experience that I've seen these kinds of projects fail is because what's missing is the underlying data in that organization to make that workflow hum along.

01:18:24 So when that CEO is in ChatGPT and kind of prototyping some idea, they're able to prove to themselves that LLMs have some particular capability. But if your enterprise doesn't have at scale the actual underlying data to be able to make, to allow that LLM capability to be realized, it doesn't matter how good LLMs are because the underlying data aren't there. I don't know if that makes sense, Carol, but that's the number one reason why I see projects fail.

Kirill E.: 01:18:58 Or the data is there, but like you had one of the guests in, I don't know, maybe start of this year say that you can't have a Ferrari if you don't have the highway, right? If you have the data, but the data pipelines are not there, the data is not clean enough, like the data is not reliable enough. So you have to first build the highway to material. You have a ton of bricks and concrete and support beams, great, but that's not a highway. You have to first put that together into a highway and then your Ferrari can go fast.

Jon Krohn: 01:19:28 That is a really, really good point. So that is another reason why. Yeah, absolutely. So either the data don't exist at all or you've just been throwing the data as unstructured documents into a lakehouse or data lake and now all of a sudden you want some production system to be running on that in real time for your users or for your employees. And so absolutely you're 100% right. All of the bricks and cement ingredients, all of that great Australian sand is sitting there in Australia that we needed in Dubai.

Kirill E.: 01:19:56 Yeah. It's interesting because the patterns that you're mentioning that you see in B2B in your engagements, I see the mirror of those patterns in B2C in our training programs because the two professions, the two backgrounds that are most common that take up the most recent, the current project that I'm most excited about, our AI engineering challenge, for example, is software engineers, software engineers, backend developers. So those people building the, they're going to be building the prototypes, they're going to be deploying the solutions, they're going to become the AI engineers, but also data engineers because they understand that companies are now more and more demanding of them to build the right highways for AI to travel upon, like quote unquote, and they need to understand how does AI actually work in order to be able to build the infrastructure for AI. So there's no wonder that the software engineers/developers and the data engineers are the two professionals, types of professionals most interested in this topic of AI.

01:21:09 Yeah.

Jon Krohn: 01:21:10 There's also for both of those kinds of professionals, there's a huge opportunity right now to be scaling yourself up by using LLMs to make your work faster. We had recently in April, we had Matt Glickman on the show from Genesis Computing. He's a CEO and co-founder of that business and their specific niche is automating data engineering. And so yeah, there's huge, huge, huge demand and huge opportunity with scaling up data engineering for sure.

Kirill E.: 01:21:37 Fantastic. I love that for a lot of these tech professions, there's like two parts about AI. One is using AI to make yourself more efficient, but the other thing is using AI, understanding how to use AI in applications, how to build things that leverage AI. So very interesting time. Have you heard of this, like what you were talking about, I had one

of these topic penciled in for our discussion, what you mentioned just now also resonated with that forward deployed engineer. Have you heard of that?

Jon Krohn: 01:22:12 Yeah, yeah, for sure. So I think Palantir was probably the first big business that was using forward deployed engineers and getting a lot of valuation of market value from using forward deployed engineers, but now we're seeing it more and more anthropic OpenAI. They're using the forward deployed engineer model and I think this goes to show how this is kind of an example of how earlier in this episode I was saying how these tools that allow me as a data scientist or as a software developer, as an AI engineer to be so much more productive in terms of writing code than ever before and evaluating my code than ever before. Just because you have a tool that can do that, you can't automatically at this time and not for presumably years to come, you're not going to be able to go to anthropic.com or chatGPT.com and say, "Automate my business." You need to have people getting into the nitty gritty, figuring out what data highways need to be built, prioritizing the highest value projects, working with stakeholders to figure out exactly how that should work so that it actually meets requirements and that it successfully automates a big chunk of a given workflow.

01:23:37 So yeah, these forward deployed engineers allow organizations, whether they're frontier labs like OpenAI or Anthropic or other software providers to be able to actually make an impact with their solution. Because that is a big thing today. I think that the days of kind of the pure SaaS software as a service model being the kinds of 80% margins that they enjoyed over the last decade or two that really got venture capitalists excited, that kind of pure SaaS model in this new era breaks down because something that you can do as pure SaaS is now pretty easy to just use CloudCode to do it automatically for you.

	01:24:22	So I think, and that's part of why I founded Why Carat is I believe that there is going to be opportunity for many years to come in figuring out how to bridge these incredible AI capabilities that we have with the real world problems that people face. And so I think, and I'm not the only one who thinks this, Hussein Kasai, who was co-founder and CEO of the biggest British startup exit in history, he's also been on our podcast. I can't look up the number right now while I'm speaking, but Hussein Kasai has been on the show. He co-founded Onfido like 15 years ago and that ended up being the largest startup exit in British history, even more than DeepMinds sail to Google. And yeah, he's also a huge proponent of this idea that the future opportunities lie not in Pure SaaS, but in this kind of forward deployed model or this kind of consultative model where it's a blend of your software solutions with really getting those data pipelines working, getting the solutions integrated into the real world workflows.
	01:25:29	That reminds me of one last guest that we had on recently and she was one of my favorite guests that we've had in the past year, Larissa Schneider. She's co-founder and COO, chief operating officer of a business called Onframe that also follows this model, a Bay Area based company that has been doing a sensational job of saying, "Okay, let's have a relatively limited number, a large but relatively limited number of software solutions that need to be tailored with Ford deployed engineers or with consultants or whatever to a particular enterprise in order to be effective." And yeah, they're seeing huge success. It's really exciting to watch on frame go.
Kirill E.:	01:26:08	In a nutshell, what is a forward deployed engineer?
Jon Krohn:	01:26:11	Oops. Thanks for asking
Kirill E.:	01:26:15	Me that. I mean, I read about it literally yesterday. I was like, "Oh, that's such a cool title." And I heard at the start

of the year, AI evals was the hottest thing in AI and I wonder if now it's still AI evals or it's now more forward deployed engineer because OpenAI just invested $4 billion into buying a company specifically to enable that capability for themselves of forward deployed engineers.

Jon Krohn: 01:26:42 Those are such different kinds of things. It's hard to imagine how, I mean, I guess you could literally look up the Google search frequency of forward deployed engineer versus AI evals, but they're such different beasts and they're also so complimentary. Your forward deployed engineers doing AI evals, maybe that's the hottest thing. It's the interaction term of those two things. So to define forward deployed engineer quickly, like Palantir was doing for years or OpenAI would do today, it's where you have an employee of Palantir, say, go into the Department of Defense or a police department. Palantir does surveillance stuff.

01:27:32 So this Palantir person goes into the Chicago Police Department or a team of forward deployed engineers from Palantir, goes into the Chicago Police Department and they're employees of Palantir and it's some big typically like seven, eight, nine figure contract that Palantir has sold to the Chicago Police Department that is paying for those forward deployed engineers from Palantir to be working alongside the IT department or engineers at the Chicago Police Department and they're there full-time, they're there in person typically. And so they kind of feel like part of the Chicago Police Department team, but they're actually employed by Palantir, even though the Chicago Police Department is paying for them. So it's a great model for Palantir because they're getting margin on there and they're also charging for the software. So you're charging for the engineer being there on site, but you're also charging for the software that they're using, which is yours.

Kirill E.: 01:28:38 Got it. So it's kind of like a tech term for what then consulting would normally be secondment, like they just second it over to the client.

Jon Krohn: 01:28:46 Yeah, exactly.

Kirill E.: 01:28:48 And I think it's picked up because one of implementation of AI, like what you're doing at YCART, for example, you deliver a solution, but if there's not enough capability within the client, you can't just use AI, you can't just integrate AI on its own. You need to teach the people that will be using AI, like you alluded to earlier, that how do they get the value of increased efficiency of them using it? How do they become more productive? How do they build a habit of using that AI? So they need somebody there to constantly remind them until it becomes second nature.

Jon Krohn: 01:29:26 For sure. And in addition to that, I think some organizations also are, there's a demand for them paying consulting firms so that they have the confidence that the decisions they're making relative to AI usage amongst their employees are secure and the best value future proofed at least for some years to come. So yeah, absolutely.

Kirill E.: 01:29:51 Awesome. Well, we have gone way over on time with this podcast. I still have like half a page of questions I could ask you. So maybe We'll have to do a second Jon as the guest appearance. I do want to ask you

Jon Krohn: 01:30:05 This question. We'll do it for episode 10,000 in one.

Kirill E.: 01:30:09 Yeah, that would be funny. We'll be a hundred and something years old each.

Jon Krohn: 01:30:14 That'll be nothing by

Kirill E.: 01:30:15 Then. Yeah, for sure. So there is one question before we get to the wrap up questions. We talked about more technology, more data centers, more intelligence, more

electricity, more capabilities, more productivity, a lot of more coming into this world. It doesn't seem that it's making us happier. It seems

Jon Krohn: 01:30:33 Like

Kirill E.: 01:30:34 Paradoxically the more easy things become, the more unhappy. And you mentioned one out of three people experience depression, anxiety in their lifetime. There's charts showing that once social media came out, anxiety and depression skyrocketed with ... There's a correlation between social media prevalence and teenage anxiety and things like that, depression. And recently came across this interesting insight that Christopher Nolan, the famous director of Inception and Interstellar, now what's it called?

Jon Krohn: 01:31:11 What's another movie that starts with in?

Kirill E.: 01:31:16 Odysseus just came out as well. He doesn't even use email and doesn't have a smartphone. He has a flip phone and he doesn't check email. And when I heard that it just made me ... I don't know how he feels about it. I think he's probably happy, but it made me feel envious in a good way. I wish I could not look at my phone, not have to check. So from everything you've seen and the guests you speak with, your background in neuroscience, what do you think we have in store for us as humans in terms of our emotional wellbeing with the world that's coming up with all this AI and prevalence and growth?

Jon Krohn: 01:31:50 Yeah, really good points. At the beginning of this episode, we talked about how I have these rosy memories of what it was like in academia and now I'm in some ways trying to go back to that kind of life, but maybe I actually can't get there because maybe what's changed isn't that I'm no longer in academia, but that we have had these social media, smartphones, AI models, email, this kind of always on culture, this feeling that we're always behind AI

moving so quickly. Maybe what I'm yearning for is because especially when I was doing my undergrad, we still had all physical textbooks. I recently did a tour of my undergrad university, Wilford Laurier here near Toronto and they ended, they kind of took me on this alumni tour and the alumni director who took me on the tour, she had it end in what is called the university bookstore, but there are no books.

01:33:05 There's a few. They have anatomy textbooks. There's like on shelf of books and then there's a few books that you would buy as like the history of the university or whatever. It's something that's like a tourist or whatever, like a memento, a token. It's not a textbook. And the bookstore is now just full of merchandise sweaters and key chains and stuff. And there was always a little bit of that, but now it's everywhere. And at the beginning of my undergrad, I had a flip phone that just had snake on it and I could go to a study room with my calculus physical textbook and a notepad and I could do calculus problems for hours. And the only interruption that I would have is other students in that study room, maybe just like looking over what they're doing or maybe they could actually come talk to me or whatever, but it's like real human interaction.

01:34:06 And these are people in the science building that I'm seeing all the time and it really had this feeling of connection. So maybe if I went back and wasn't academic, maybe I couldn't even recreate that feeling that I miss. So yeah, I'm really hearing what you're saying here with your question. And I definitely, I suffer from it and I'm jealous of Christopher Nolan too. I try to have periods. It was recently Memorial Day weekend in the US, which is a three day weekend in May, Saturday, Sunday, Monday. And I tried very hard to not go on Slack, to not check my emails, to not go on social media, to not respond to texts that I don't need to respond to over the weekend. It's not making plans with somebody that day. And I really

enjoyed that. I really enjoyed those three days. But then on Tuesday, when I came back, there were things like Kyong Yung Cho, who was on this show recently, it's I think the most popular episode of 2026 so far, episode 997 with Professor Kyong Cho, this amazing, one of the most excited AI authors ever, AI academics ever.

01:35:18 He was in- Attention, right?

01:35:19 Yeah. Yeah, he co-authored- Co-invented attention.

01:35:22 Yeah, with Yashua Bengio and I forget the first author's name, which is a shame because that's actually the first author, but the three of them co-invented attention and he was giving a lecture, a five minute walk from where I'm staying in Toronto and I could have ... Yeah, so I'm actually technically, I'm in a suburb. I'm in Waterloo outside Toronto and he was a five minute walk to go see him deliver this big lecture in a lecture theater. And I didn't know that because I didn't check, like he posted it on LinkedIn over Memorial Day weekend and I missed it. And so it's like this weird, like you feel like you're always missing out if you don't check these things and obviously like I can't, I'm the CEO of Y Carrot. I can't not check my work email or not respond to work emails because the business is going to die.

01:36:10 I can't not check Super Data Science podcast Slacks because then we're not going to have any podcast episodes. I'm trapped in this and it's difficult to see a way out of this like needing to kind of be always on. At least with work stuff, there's mostly this expectation that in most businesses you don't need to be responding on holidays or evenings. But yeah, I don't know. I don't know how you get away from this feeling of being left behind with social media. I guess Christopher Nolan is in the fortunate position of being extremely successful and he can afford probably multiple assistants and people supporting him and someone must be like monitoring

		emails or monitoring social media feeds and when something that he needs to know about happens, they can phone him up on his clip phone.
Kirill E.:	01:36:59	He said they print out emails if he really needs to see an email, they print it
Jon Krohn:	01:37:03	Out for
Kirill E.:	01:37:04	Him.
Jon Krohn:	01:37:04	Exactly. So
Kirill E.:	01:37:06	That's kind of like what I'm thinking. Do you think AI will help us get more, everybody become more like that or it will make things even worse than there is now?
Jon Krohn:	01:37:14	I think it's basically up to you as an individual. I think probably in the same way that social media and digital ads and TV ads before that, basically the vast majority of people are being taken advantage of in a system because they're not creating the space to say, "Okay, is me scrolling through my social media feed here doing me net positive for my emotions and my finances or is it hindering me? " And it's hard and especially like the platforms, whether it's TV ads all the way through to social media experiences, they have engineered the product managers or the marketing people have come up with ways that they make these experiences really compelling. It's hard to step away. It's so much easier. Yo get that quick dopamine hit and you keep pulling on that dopamine lever long after it's provided any actual positive sensation because you're hopeful that that next video that you flip to on TikTok, that next super data science podcast episode is going to be more interesting than when Jon was interviewed by Kirill.
	01:38:31	So it's hard to step out, but I think, and probably people who are listening to this podcast, I don't mean this episode necessarily, but to this podcast in general, you

		are the kind of people who are trying to figure out a way to have AI systems, AI agents be like Christopher Nolan's assistants and you could literally have it phone you on a flip phone or print out emails for you and you talk to what seems like a human voice. There is no technical limitation to engineering that solution for you today and there are maybe some set of SaaS products that you could chain together or open source solutions that you could chain together to deliver that for you, but it requires time and it requires, you need to get past that fear of making mistakes. There's going to be things like, I'm going to miss Yangyon Cho lecturing.
	01:39:21	If I do this, if I step away from social media and have an AI system do it, there's going to be things like that that I miss unfortunately, but I've got to think, okay- I don't think
Kirill E.:	01:39:29	You would have missed that one.That one would have probably come up in ways.
Jon Krohn:	01:39:34	Just depends how good, just how well I engineer the system. That's right. I can mess it up. I only have it phoning me about things that happen in my direct messages and noth on my feed. It depends on how you engineer it. So I think that there's opportunity and maybe we'll even get to a future with abundant energy and abundant intelligence where in the same way that a lot of countries in the world have universal healthcare and affordable education, maybe part of that will be guaranteeing that you have access to agentic systems that make you like Christopher Nolan, that allow you to be freed from this frantic pace of technological change and allow you to feel kind of grounded and centered and maybe that is a future. I hope that that is a future. And I think that in the long run that's possible, but I think in the short term, in the coming years, it's up to you.
	01:40:38	It's up to you to do that for yourself.

Kirill E.: 01:40:40 Great way to end this episode. Great call to action. Thank you, Jon. It's been a pleasure interviewing you again so many years later, learned quite a lot of new things and as usual, what's your book that you'd recommend to our listeners?

Jon Krohn: 01:40:59 All right, Curol. So I've got a nonfiction recommendation and I've got a fiction recommendation. Will you let me get away with that?

Kirill E.: 01:41:09 Sounds good. Let's do it.

Jon Krohn: 01:41:10 All right. So my nonfiction recommendation is actually from a recent guest on the show. It was episode 975 with Zach Cass. So Zach Cass was the head of GoToMarket at OpenAI when ChatGPT came out and it looks like you're looking back at your bookshelf. So you probably have his book, the next Renaissance. Yeah, Ren AI Scents. Renew Bookshelves. Yeah. This one. See it from that kind of chartreuse cover. Yes, yes, yes, exactly. And that episode, I guess if you want to see whether you should pick up this book or not, listen to episode 975 and see if you like Zach and his way of thinking, but he aligns really nicely with this techno optimism that I have and his book is a very easy read. It's 200 pages with big font and he does a great job. He's in a way that I- I was going to

Kirill E.: 01:42:02 Say, you'll probably read it faster than listen to the episode.

Jon Krohn: 01:42:06 It's possible, but yeah, you'd have to be a really good reader. This is probably like a 10 hour book read, but it's a one hour podcast episode. It's something I'm planning on giving it to some family members, especially I have one family member that I have in mind who thinks that everything in the world is bad now and things are going to be worse and I'm like, I don't know. I think things are really good and getting better. And I think it's kind of your social media feed, the news cycle that is making you feel

like everything's bad and everything's getting worse and that translates into your inner reality. So in your inner reality, things really can be getting worse. It can be getting darker if you feel like that all the time. And Zach Cass in his podcast episode, 975 of this podcast, as well as in his book, the opening is about exactly this, about the social media and newsfeeds and how we feel like things are getting worse, but then he does a really great job of explaining why things are great and in a big part, thanks to AI, things are probably going to be even better in the future.

01:43:06 So that is a recommended read. It's not like, I think on the podcast before I've recommended Sapiens by you of Alnora Harari and those are very diferent kinds of books because Sapiens is quite dense. It's really like you're like, wow, there's so much to learn from Yuval Noah Harari on every page.

01:43:27 The next Renaissance is not like that. It's like short to the point, but you kind of get all the key ideas out about where we are today and where we're going in the future with AI in an exciting and positive way. So that's my non-fiction rec. And then my fiction recommendation, so regular listeners will probably already know that my favorite fiction author is Kurt Vonnegut. He is funny and something that I only recently pieced together because I recently read one of his ... I was working through all of his books in chronological order and I'd never read ... His earliest books are what he's most famous for. So the most famous book, which a lot of Americans have to read when they're in high school, I understand is Slaughterhouse five by Kurt Vonnegut. I actually, that doesn't even get into ... Definitely it's not in my top five of Kurt Vonnegut books.

01:44:24 I don't actually ever even recommend reading Slaughterhouse five. Kat's Cradle I've recommended on the show before, which is an amazing book by Kurt

Vonnegut. I've also recommended on the show before player piano, which is particularly interesting to read in this AI era. But so like I'd say in general, Kurvanage is super interesting and because I've been working through things chronologically in one of his last novels he talks about a letter, he talks about kind of like common themes through his books. And one of the things that's really interesting, and I didn't piece this together about like, why do I love Kurt Vonnegut so much? And I think one of the things is his books don't have villains. There's no bad guys, which is I think what a lot of real life is like. Most of us, there are some really bad people out there, but I think the vast majority of people are good people trying to do good things.

01:45:16 And what happens in Kurt Vonnegut books is that like random happenstance leads to like catastrophe or bizarre interesting situations where like everyone is trying to just do the right thing, but like society collapses and everyone dies anyway. And so that's kind of like a, I don't know, I really like Kurt Vonnegut for that. So Kurt Vonnegut in general, but specifically I'd recommend Kat's Kradle.That's my favorite book by him.

Kirill E.: 01:45:43 Fantastic. Thank you, Jon. I've never read Kurt Vannegan, but sounds really interesting. All right, last question. Where can our listeners follow you and find you if they want to know more?

Jon Krohn: 01:45:56 For sure. Thank you for asking that Quro, following the template that you created for me when I took over some hundreds of episodes ago. We continue to do that on the show always asking for a book rec and how to follow people. And yeah, so the primary social media for me as it is for most of our guests these days on the podcast is LinkedIn. Do feel free to connect with me on LinkedIn and if you mention that you basically, if you say anything that shows to me that you're not trying to sell me something or trying to hire me as long as you can just basically say

like, "I liked your podcast or I like your YouTube channel or whatever." Anything like that, I'll definitely accept your connection request, but you can also just follow me there if you want to. Anyone can follow on LinkedIn these days.

01:46:47 I also have a personal YouTube channel. So in addition to the Super Data Science podcast, YouTube channel, I have a personal one which has in recent years videos on Agentic AI and for the most part kind of hands-on videos for Agentic AI. Going back a bit further, my core content was around the math of machine learning, so the linear algebra the calculus. And I am planning on, so for people who have been following my personal YouTube channel and are disappointed that it's been many years since I've been creating that math content or producing videos on my personal channel with regularity like we produce the podcast episodes here now that Sonya is doing so much operations for me, have a solution and it actually involves the media editor Mario for this podcast where I'm committing to creating video content on a regular basis. So I'm not committing to a specific date for that happening, but in the near future, possibly even when this episode is out, I will kind of have a regular, ideally even weekly cadence of publishing videos there as well.

01:47:49 So yeah, LinkedIn, YouTube, I have an email newsletter that you can sign up for on JonCrown.com. That's another thing that I have been bad about in the past year about doing. But again, that's something with Sonya that we should be able to get going again and yeah, those are the best ways to follow me.

Kirill E.: 01:48:05 Love it. And don't forget to subscribe to this podcast. This is where you can find Jon every week, twice a week.

Jon Krohn: 01:48:12 Yeah, that's for sure. That is for sure and for all the way through to episode 10,000 and beyond.

Kirill E.:	01:48:20	To infinity and beyond as the Toy Story character says. Fantastic. Thanks so much, Jon. It's been a pleasure. I loved interviewing you, learning from you and keep rocking it. As you say, keep on rocking it out there with the Super Data Science podcast.
Jon Krohn:	01:48:33	Thank you, Kira. Well, yeah, such an honor to have you interviewing me here and I hope people enjoy the tables being turned a little bit at me talking. I feel like I talk way too much, but maybe they listen to the episode at 1.5 speed or something and the insights were useful that way. You're a really great interviewer, Kirill. I can't believe that you asked me to be the host of this amazing podcast that you founded and that you got to such an amazing state in terms of business operations and popularity. It completely changed my life and mostly for the better.
Kirill E.:	01:49:11	Fantastic. Thanks. You're too kind. Thanks, mate. It was a pleasure.
Jon Krohn:	01:49:15	I hope you enjoyed today's role reversal episode in which I was the guest and this podcast founder and former host, Kyril Eremenko, hosted it. Something special we were doing for episode 1001. In this episode, we discussed how studying neuroscience at Oxford was driven by a fascination. My studying, neuroscience at Oxford was driven by a fascination with how the mind emerges from chemicals and biology and why I now regret leaving academia for a hedge fund. I talked about how coding tools like Claude Code have eroded the technical mote I built over my career, but this paradoxically makes skilled scientists and engineers more valuable than ever because we can now deliver far more. I talked about how Javon's paradox means that making AI cheaper and more efficient will increase total demand for compute not shrink it, how Rice scoring helps organize pick winning AI projects and how a well-developed AI agent could act like a personal assistant that gives us Christopher Nolan like Focus.

01:50:12 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 Kirill's, social media profiles, as well as my own at superdatascience.com/1001. There is a fun episode number for everyone, but especially the computer scientist set there. 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 and our founder Kirill Eremenko. Special thanks to him for conducting such an energizing interview. Yeah, again, hope you enjoyed it. Thanks to everyone on the team for producing another super 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 sponsor's links, which are in the show notes.

01:51:06 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. Otherwise, share this episode with folks that would like to, I guess, hear an interview by me review the episode on your favorite podcasting app or YouTube. Subscribe if you're not already a subscriber, but most importantly, I hope I'll just have you listening again. I'm so grateful to have you listening and I hope I can continue to make thousands more episodes for you to love for years and years to come. Until 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.