



SDS PODCAST

EPISODE 807:

SUPERINTELLIGEN-

CE AND THE SIX

SINGULARITIES,

WITH

DR. DANIEL HULME



- Jon Krohn: 00:00:00 This is episode number 807 with Dr. Daniel Hulme, CEO of Satalia, and Chief AI Officer at WPP. Today's episode is brought to you by AWS Cloud Computing Services, and by Gurobi, the decision intelligence leader.
- 00:00:18 Welcome to the Super Data Science Podcast, the most listened to podcast in the data science industry. Each week we bring you inspiring people and ideas to help you build a successful career in data science. I'm your host, Jon Krohn. Thanks for joining me today. And now let's make the complex, simple.
- 00:00:49 Welcome back to the Super Data Science Podcast. Today you're in for a treat with the brilliant Dr. Daniel Hulme. Daniel is Chief AI Officer at the marketing giant WPP. He's also CEO of the AI consulting Services company, Satalia. He's entrepreneur in residence at one of the world's top AI research universities. That's University College London. He's co-founder of Faculty, and a speaker at the Singularity University. He holds an engineering doctorate in computational complexity from UCL.
- 00:01:18 Today's episode should be of interest to anyone. In it, Daniel details how and when Artificial Superintelligence may arise. The six types of singularity that ASI is expected to unleash, neuromorphic computing, how to align AI interests with human interests, and ways human work could be dramatically automated, not just in the future, but this very day. All right, you ready for this spectacular episode? Let's go.
- 00:01:49 Daniel, welcome to the Super Data Science Podcast. I met you at the time of recording, I think about two months ago. We met at New York University, they ran an AI masterclass, which I think we could also recommend to people if they're... I think it happens about every six months. Jepson Taylor, who has been on this podcast, I think probably more than anyone else, he's been on the

podcast like 10 times. And I think always under his previous name of Ben Taylor. So if you're looking for those episodes, they'd be under Ben Taylor. But now Jepson Taylor brought us together. We did a panel on, I guess it was just on the future of AI, and how it's going to impact business. And I loved everything that you said. In fact, my very first question is going to be following on from one of the things that you said on that panel. So welcome Daniel, how's it going man? And where in the world are you calling in from?

Daniel Hulme: 00:02:48 Oh, it's great to be here. Thank you very much. I'm calling from my little office in London, my home office, and also massive fan of Jepson as well. I'm going to have to try and remember what I said in New York now.

Jon Krohn: 00:03:00 No, don't worry. I've got it. I've got it for you. But really quickly before we get onto that topic... For people who are watching the YouTube version of this, you get to really enjoy this super kitted out closet that you work out of. It looks like it's the kind of thing, that if I was a kid in school, in elementary school, I would be sketching out what you have as your office all around you. And there's a really fun, there's even a ladder which you can see on camera, and I was hoping would lead to a tree house, it leads to some high bookshelves, which is still pretty interesting. But maybe we can work on the tree house for the next time you're on the show.

00:03:42 So the topic that I wanted to start off with, it's a light topic, it's a singularity. and so specifically one of the things that you talked about on that panel at NYU, was this idea of there being six singularities. And so that was something that completely blew my mind. In fact, for me it seemed up until that point, that inherently, in the word singularity, it almost sounds like single. And so I kind of thought there'd be just this single point, beyond which we cannot predict the future. But you talked about six

singularities then, so I'd love to open the floor to you to bring those up to our audience now.

- Daniel Hulme: 00:04:25 Yeah, yeah, of course. I guess you know the history perhaps of the word singularity? It comes from physics, it's to do with black holes and a point in time and space, that we can't see beyond. And it was adopted by the AI community [inaudible 00:04:37] Ray Kurtzweil, to refer to the technological singularity, which is the point in time where we build a Superintelligence. A good friend of mine, Callum Chace, coined the term the economic singularity, highly recommends his books. He's written a lot over the past three decades on what happens when we automate the majority of human labor.
- 00:04:58 And again, it's potentially a point in time that we can't see beyond. And having discussed the impact of AI with people over the past 10, 15, 20 years, you realize we're facing into interesting challenges or events in humanity, such as curing death, or even our environmental singularity. The point in time where we might either lose control of our ecosystem or completely gain control of our ecosystem. So I realized when I heard these words, technology, economic, environmental, that those were three of the words that are part of the PESTLE framework.
- 00:05:31 So if anybody's done a business degree, they'll probably come across PEST analysis or PESTLE analysis, which is essentially four or six macro words that refer to political, environmental, social, technological, legal and environmental or economic, I can't remember. And then I asked myself, "Well, is there a PESTLE of singularities? Are there points in time that we can't see beyond that can relate these six macro words?" And I am very happy to go into the detail if that would be helpful?

Jon Krohn: 00:06:06 I would. And also I'm going to try to make sure that I got these down right. So PESTLE , spelled P-E-S-T-E-L or L-E, I guess?

Daniel Hulme: 00:06:14 Yeah. Yeah.

Jon Krohn: 00:06:17 And it doesn't really matter in terms of getting all the words in there. So P for political, one of the E's for environmental, social, technological, another E for economic?

Daniel Hulme: 00:06:28 Yeah.

Jon Krohn: 00:06:28 And L for legal?

Daniel Hulme: 00:06:29 That's right.

Jon Krohn: 00:06:31 Okay, sweet. Yeah, let's dig into them.

Daniel Hulme: 00:06:33 Yeah, so I guess the political singularity, and when I was thinking about this, this was four or five years ago, around the time of Cambridge Analytica. And actually things like Brexit, and I was just seeing the impact that AI was having on our political foundations, challenging in fact our political foundations. And that raised the question of what happens if we start to live in a world where we don't know what is true?

00:06:58 So either we create a world where we don't have any faith in what is being presented to us, deep fakes, misinformation, bots, et cetera. Or there's a future world where we're able to know what is true, and authenticate what is true. And so the political singularity is the point in time where we either can or can't authenticate what is true. And I think one of the biggest fears that we have currently at the governmental level is, is us losing trust in our social systems, and what happens if we cannot determine what we're engaging with, that it is true.

00:07:31 And I guess that, that sort of extends a little bit more to not just our political foundations, but also the fabric of our reality. I know people who are currently being attacked by deepfakes of their children or their work colleagues. And in fact, two months ago, the CEO of WPP was deepfaked, and a clone of him was created and somebody tried to use it for malicious reasons. So the question is, what happens in a world where we don't know what is true, and how do we mitigate that risk? So that's political singularity.

00:08:02 The second singularity is environmental singularity. And we all know that the AI is increasing consumption, and as far as I'm concerned, consumption gives people access to goods and services that typically enrich people's lives. And we know it's putting pressure on our planetary boundaries, we know there are people that overconsume. But I guess there's two potential outcomes in the future. One is that we lose control over our ecosystem, or we gain control over it. And actually I'm very hopeful that if we apply algorithms, if we apply AI in the right way, to solving problems across supply chains, making it much more efficient, much more effective... We could significantly reduce the amount of energy that we need to run this planet.

Jon Krohn: 00:08:45 And things like helping us figure out nuclear fusion? So being able to contain the plasma inside a nuclear fusion reactor, to keep the reaction going? Maybe help us devise new ways. There's interesting things by you in London there, there's a number of nuclear fusion groups in Oxford, including First Light Fusion, which it's propelling a pellet, at very high speed, that somehow it creates this collision. It's inspired by the way that rainbow shrimp can create those really crazy reactions to kill their prey underwater?

00:09:23 They generate this super high speed, attack with a claw that produces bubbles, and when the bubbles collapse, it causes their prey's shell to break apart. It's inspired by that. So these alternative ways of trying to come up with nuclear fusion, and that research came out of simulations. And so there's all these different ways that you could imagine AI helping us, with not only the kinds of things... The energy expenditure like you mentioned there, like supply chain issues, but also potentially with creating an abundant source of fuel on the planet.

Daniel Hulme: 00:10:05 Exactly. I think the point here is, that when we hear the word, the singularities, we often default to the negative. What happens if we create a post-truth world? But the point is, is that actually we could use AI, we could use technology to prevent the post-truth world. We could use an AI to help us get control over our ecosystem by creating new energy sources. And again, it's something that obviously DeepMind has also been investing in as well, is how to solve nuclear fusion.

00:10:30 So the third singularity, it is actually not my expertise. It's sometimes often referred to as the methuselahity, and it's the point in time where we cure death now. From what I understand, there are scientists that believe there are people alive today that don't have to die. And AI is obviously rapidly in advancing medicine, it's able to monitor ourselves and eventually we'll have nano bots potentially cleaning ourselves out.

00:10:54 And a bit like a car, if you stay on top of damage, that car will never ever break down. And I don't know what the world will look like if we realize there are people amongst us that won't have to die. It will change, I guess, how we educate ourselves or relationships that we have. It changes the complete dynamic of what a lifespan looks like.

- Jon Krohn: 00:11:10 And the way that you'd live a life. Because imagine if you can live forever, you wouldn't want to go outside, take any risks at all. You'd spend a good chunk of your thinking time just trying to make sure that you can keep it coming along, that you're not going to get hit by a bus or contract a disease. So it's an interesting world, a risk-free world.
- Daniel Hulme: 00:11:32 Well, actually there's obviously precursors to this. There's actually a company that I've just been informed of called, I think Tomorrow.Bio, and for I think it's like \$25 a month, you subscribe to this. Essentially this cryogenic service that if you die, they will freeze you. So for \$25 a month, if you die, they'll go and freeze you. Now you have to pay a lump sum of, I think it's like \$70,000 to freeze your head, or \$200,000 to freeze your body. Nobody knows if you can be unfrozen, but there are people actually now that are being frozen, with the hope that they become unfrozen at some point in the future.
- Jon Krohn: 00:12:08 That's wild. It's the kind of thing that I'm aware vaguely is going on, but it's really trippy to think about if a loved one was coming close to the end and you're like, "All right, we're just not going to end it. Put her in the freezer."
- Daniel Hulme: 00:12:23 Indeed. So the fourth singularity, actually one of my favorite singularities is the technological singularity. And of course, I think this is probably the most popular one. Which is the point in time where we build a machine that's a million times smarter than us. It's a point in time where we become the second most intelligent species on this planet. There's been a lot of books written about this, and I guess I flippantly advise people that when Superintelligence comes, look busy, be nice to each other and hopefully it'll bugger off to a different dimension.
- 00:12:56 And I guess my community and you would be able to tell me better Jon, based on all of the people that you've spoken to. But my community, I guess felt this wasn't

going to happen for another 30 or 40 years. We now think it might happen in the next 10 or 20 years. In fact, I found myself, forgive me for name-dropping, on a yacht with Elon Musk three weeks ago in Cannes, and we were talking about Superintelligence. Because I was pitching a company idea to him, and he thinks that we're a year away from Superintelligence. And I guess Sam Altman might suggest that we are four or five years away, but we could be close to building a machine that is a million times smarter than us.

00:13:34 And one of the things that I've decided to do, is actually lean into this question. I think I've been very lucky over the past 25 years seeing how technology has played out, and I've made some good guesses. But I think that actually machine consciousness is going to become very important over the next decade. And so I've started a research company on the side. Happy to talk about this later on, if you want?

Jon Krohn: 00:13:57 No, we can get into it a little bit now, this is Conscium, I guess?

Daniel Hulme: 00:14:01 Conscium, that's right, yeah. And so one of the big questions or a couple of big questions that we have in the field of AI is... The first question is the control problem. Could we build a Superintelligence that we can control? And my instinct is that the answer is no, but it's like trying to control a god. But there is another concept called the alignment problem. Could we align an AI, a Superintelligence with human values, or some value system?

00:14:28 It doesn't necessarily have to be human values, like altruism and cooperation and sacrifice, and all this kind of stuff. And I guess how people think currently about AI, is that you program something or you train something, and where it goes wrong, you try and put plasters in place



to try to mitigate those risks. It's like playing whack-a-mole, right? You have to knock it down every time it goes wrong.

- 00:14:54 And I actually have some ideas about how we might be able to evolve AIs. We've evolved, our species and other species have evolved over the past 2 billion years, where we've had to cooperate for our genes to be passed into the next generation. We've had to be altruistic, we've had to be sacrificial. And those are not programmed. They are embedded into the being of our bones, and our neurons. And I think that we could evolve AI in environments, and have those AIs have to survive by cooperating, by sacrifice, and all these kind of things.
- 00:15:26 And instead of programming AIs to have a value system, we might be able to embed value systems into their neural structure. Now, I'm not arguing that, that's going to be completely robust. It might be that they decided to not align with those values, but at least you would start out from a point where you had a value system embedded. And I actually think... There's a new emerging technology I've been tracking for the past few decades called neuromorphic computing. So Large Language Models are not really how our brains work, they're called neural networks, but they're really not how our brains work.
- 00:15:59 Our brains operate on the power of a light bulb. We learn very quickly, we're very adaptable. And these neuromorphic technologies that are appearing, also called spiking neural networks, actually could lend themselves to AIs that are aligned, but also might start to exhibit conscious behaviors. I think there's two big interesting questions that we need to face over the next decade. One is, humans are going to start attributing consciousness to things that are not conscious. That maybe problematic, it may not be. We've all seen the movie Her.

00:16:35 And the other question that we need to ask ourselves, is that if we inadvertently build machines that are conscious and not realize it, we could be putting those things in torturous or traumatic situations. And in some respects, we have a duty of care, not just to humans or to animals. We potentially have a duty of care to AIs, to make sure that they don't suffer. Nick Bostrom coined the term, mindcrime. So I think this is really an important set of questions we need to be facing into.

Jon Krohn: 00:17:07 Are you stuck between optimizing latency and lowering your inference costs as you build your generative AI applications? Find out why more ML developers are moving toward AWS Trainium and Inferentia to build and serve their Large Language Models. You can save up to 50% on training costs with AWS Trainium chips and up to 40% on inference costs with AWS Inferentia chips. Trainium and Inferentia will help you achieve higher performance, lower costs, and be more sustainable. Check out the links in the show notes to learn more. All right, now back to our show.

00:17:45 You did manage to get a lot of content into that. I am trying to race to keep up for it in my notes, and all the kind of follow up questions that I have. So I want to start with Consciousness, or I should also say that we've gone off on an avenue, we will come back, so don't worry listeners. We were going through that PESTLE framework corresponding to the singularity. So we'd gotten through political, environmental, and social, what you called there, the methuselahity, which is super fun.

00:18:12 And then we started on the fourth one, which is the technological singularity, which is the one that probably most people associate directly with the singularity. Their mind currently being expanded in this episode to these other kinds of singularities. But that is the creation of an Artificial Superintelligence, which exceeds our

capabilities. And so you have recently created Conscium, because of your keen interest in studying the consciousness of machines. I like that idea of avoiding, you said the Nick Bostrom term, what was it? Mind-

Daniel Hulme: 00:18:48 Mindcrime.

Jon Krohn: 00:18:50 Mindcrime. And so that's cool. I like that a lot. It is interesting to me that this was a recent conversation with Elon Musk and hopefully you are, and probably his thinking has evolved on this as well. No doubt people like you end up influencing him on these kinds of topics. But it was maybe about a year ago where he opined publicly that the way to avoid terrible outcomes at the singularity, is to create a super intelligent machine that is maximally curious. Which seemed to me to be maybe a little bit simplistic? As you described there, it is probably more like a complex whack-a-mole. And something that I think of in my mind, you can correct me on this, because you've spent a lot more time thinking about this than I have.

00:19:39 But when I try to reassure people about the creation of a machine that is more intelligent than us, and how that can probably be done in a way that is safe, there are other kinds of systems that humans have created over the past millennia, that are out of the control of any individual... Governments, militaries, the stock market. And so there are these kinds of things and bad things happen, and it is like a whack-a-mole of, "Oh, that was a mistake." So stock market collapse and lots of people lose their livelihood and people are jumping out of skyscrapers and we figure out, "Okay, we need independent central banks, in order for this to work." We need these kinds of financial regulations. We need to make sure that banks have a certain amount of security to avoid stock market collapse.

00:20:29 And I mean stock market collapse sounds bad, but governments and militaries have led to much bigger calamities. The 20th century had a huge amount of death and disaster from war. And so political systems get out of control, you can't get somebody out of power, and then all of a sudden they're rolling their tanks into your country, and killing certain people within your country. And so obviously things can go awry, but it does seem like that whack-a-mole that we're doing, iterates in the right direction.

00:21:00 With artificial intelligence, it seems like we're at a point where there are a lot of people trying to figure out where those biggest whack-a-mole problems are, in advance. And we are fortunate to, at this time in human history, have more brains on the planet than ever before. Those brains have more free time than ever before, to be thinking about things, because we're not just tilling the fields as an illiterate peasant, like we would've been most of our history, statistically speaking, you're just trying to subsist.

00:21:39 And so all these humans, all this thinking time where many of us, billions of us are connected in real time over the internet, hearing new ideas over podcasts, archive papers, GitHub code repos. And so there's this opportunity to get ahead of the biggest whack-a-mole issues around AI, and it leaves me feeling optimistic. So I just talked for a really long time on your episode. I don't know if you want to dive in, and [inaudible 00:22:03].

Daniel Hulme: 00:22:02 No. That was a lot. I see it sets me up really nicely for the final singularity which maybe we'll come on to in a minute. But just to go back to this idea of curiosity, which is what Elon mentioned, which is one of the reasons perhaps we are curious, is because it is evolutionary beneficial for us. There's an idea in consciousness called the free-energy principle that was

coined by a chap called Karl Friston. And what we're ultimately trying to do is, we're trying to minimize our perception of the world, versus actually what happened.

00:22:35 So we expect something to happen. It doesn't happen in the way that we want it to happen, and we need to learn or be conscious about essentially that delta. And so curiosity I think is a really important mechanism for us to be able to be an adaptive, and more effective species. And then one of the questions is, how do we evolve agents to be inherently curious and not just programmed to say, "What is that?" Ask the question, "What is that," if it sees something that it doesn't know.

00:23:06 I want it to have the inherent desire, innate desire to want to learn things in a way that I guess humans and other animals do. So I think absolute curiosity is a very interesting question to answer, over the coming years. And it'll be part of obviously Conscium's thesis as well. I should say that I did my PhD 20 years ago in modeling bumblebee brains. Bumblebees have a million neurons, and even looking at AI agents, biological agents that have very, very small brains, they can do phenomenal things.

00:23:40 And bees are curious, they have to of course explore their landscape, to be able to find and forage the new things that are beneficial for the hive. They also have sacrifice and altruism, and all this kind of stuff. So even with very, very, very small brains, we can do some lots of interesting experiments around these questions.

Jon Krohn: 00:24:00 I have to interrupt you, because I have to point out how fun the alliteration of bumblebee brains is, and it makes me really want to... I hope that your PhD was titled something like Building Better Bumblebee Brains.

Daniel Hulme: 00:24:14 Actually 20 years ago, building brains that had a million neurons was actually just impossible.

- Jon Krohn: 00:24:19 No, I know. I just couldn't think of anything to-
- Daniel Hulme: 00:24:20 Yeah, so I did-
- Jon Krohn: 00:24:20 ... literate it better.
- Daniel Hulme: 00:24:22 ... I changed my PhD to a completely different field in computer science. I don't know if you're familiar with the P versus NP problem? There are seven mathematical questions that have been asked by Clay's Mathematics Institute, they're called the millennium problems. One of them has actually been solved by a Russian recluse. Well, one of those questions is called the P versus NP problem, which is... Are there essentially algorithms that can solve what are called exponential problems?
- 00:24:53 An exponential problem is something that scales very, very badly, very, very quickly, and nobody knows if there are algorithms that allow us to do that. And the reason why I got interested in algorithms, is because it was impossible 20 years ago to build brains that even have a million neurons. And I thought that the bottleneck was, "How do we use better algorithms to make these direct brains learn more quickly?"
- 00:25:14 Turns out that Geoff Hinton at the time, was solving that problem, but on GP use. So using compute instead of necessarily algorithmic advances, and hence the reason why obviously Nvidia and GPUs are very, very, very popular now. Because they allow us to train big, big brains. I do think, by the way, there's an algorithmic element to this that's going to be beneficial in the future, but that's probably a different conversation.
- 00:25:38 So my PhD ended up being in both neural networks and also complexity theory, which allowed me to go and build a company that's been applying these technologies to solving problems in the real world. Anyway, going back to

the PESTLE of singularities, the legal singularity is the point in time where surveillance become ubiquitous. And I guess what I mean by this is, that we've probably all seen the Minority Report where you have these, what are called precogs that were able to see into the future a few minutes, and people would then use that vision to prevent crime before it happens.

- 00:26:16 And what would the world look like if we are able to understand people deeply, predict their behaviors? And that's an incredibly powerful position to be in. WPP is in the position of understanding perception and influencing perception. We want people to buy goods through ads, but that ability to understand people, understand perception and influence perception, is an incredibly powerful position. And we want to mitigate the risk of that capability being used by states, or bad actors from using that power to accumulate more wealth and more power.
- 00:26:54 So there's a good reason to be able to understand people, but you could also be potentially using these technologies for malicious reasons. The final singularity is actually my favorite singularity. It was coined by my friend, Callum Chace, which is the economic singularity. And I guess for the past 17 years I've been building AIs that have been freeing people from doing repetitive structured tasks. Those people have actually not lost their jobs. They've gone on to do more impactful, more purposeful, valuable things within their organizations. And my prediction is, is that over the next 10 years, we're going to see a [inaudible 00:27:32] explosion of new innovations, new opportunities for people to contribute.
- 00:27:36 I think jobs will be displaced and disrupted, but AI is like a new energy source, and that will allow humanity to grow. Going back to your point, there's more people now being more free to use their brains, use their minds, to come up with new innovations that actually makes

humanity grow to another level. I think, the point is, that beyond 10 years, I don't think really anybody knows what they're talking about. And I think you probably know the two extremes of the argument. The one extreme of the argument is that if we can free-up whole jobs because of using AI, we probably will.

- 00:28:08 The pressure to reduce costs, increase profits within our organizations, that pressure means that we probably will free-up whole jobs. And if that happens very quickly, if lots of jobs are displaced, our economies won't be able to rebalance and it could lead to social unrest. There are questions like UBI and four-day working week, and other mechanisms that might take the edge of some of those challenges. I mean the opposite and other extreme, or school of thought, is that we should be accelerating as fast as possible to the economic singularity. The idea is if you do free people from jobs, in theory, you are able to reduce the cost of those goods.
- 00:28:50 If you can reduce the friction from the creation and dissemination of food, healthcare, education, energy, you can bring the cost of those goods down to zero or almost zero. So imagine, if we applied our smarts in the right way over the next 10 years and we created a world of abundance, a world where everybody has access to the goods that they need, to survive and thrive. Now I know lots of people who don't have jobs. Those people are not staying at home, bored and depressed. They are usually using their time to spend more time with their family, to indulge in their hobbies, to travel, to enrich their lives.
- 00:29:28 But what happens if you push people hard enough and say, "Look, what would you do if you didn't have to do paid work, and everything was free?" They'll go and do all of those things I've mentioned. But if you keep pushing people hard enough, they'll say the same thing, which is, "I want to do something that contributes to humanity." I

think we all have an innate desire to want to make the world a bit better. And unfortunately, all of us are pretty much born into economic constraints, preventing us from doing that.

00:29:52 We're forced to live only for ourselves, and I think that if we use AI in the right way, we could create a world where people are able to live beyond themselves. And actually that's the purpose of Satalia, my company, is to create a world where everybody is economically free, to use their creativity to contribute to humanity however they want.

Jon Krohn: 00:30:12 Many decisions businesses face are massive, complex and heavily constrained. In these scenarios, mathematical optimization is often the best tool for the job. And Gurobi, which is trusted by many of the world's leading enterprises, is the go-to provider for fast at scale optimization. While filming episode number 723 last year, I had my mind blown about the wide range of scenarios, where optimization is the right solution for the job. Check out episode number 723 as well as the introductory resources for data scientists at gurobi.com to get up to speed on when optimization is ideal, and add this uniquely powerful tool to your data science toolkit. Then, coming up in August, we'll have a second episode on optimization with even more tips and tricks from Gurobi guru Jerry Yurchisin. Hope to catch you then.

00:31:01 Fantastically stated, I love that. That is a really well articulated vision of something that I'm absolutely onboard with as well. Something that I think about a lot, and that I hope to with literally things like this podcast, I hope to switch people's minds onto, "Oh, I don't just need to be building an algorithm to be able to predict a stock price movement, two milliseconds faster than my competitor. I could actually be helping with unlocking nuclear fusion, or making a big social impact."

- 00:31:32 And so that's a big part of what drives me, and you articulated it so well there, that an economic singularity would free basically everyone up to be able to think about, "How can I make the world better? Because you're not just worried about feeding your family, educating your family, making sure that there's a roof over your head. So very, very cool.
- 00:31:51 I have a question from a listener here. So I posted that you would be coming up on the show and Solomon Khan who is the founder and CEO of a tech company called Delivery Layer. He wrote a question that I thought was pretty funny. I said that we were going to be talking about the singularity, the development of Artificial Superintelligence, and how this would overhaul human society. And Solomon says, "How much does being an AI expert like you are, give you credibility to predict what society will look like after we achieve an Artificial Superintelligence?"
- 00:32:26 And so Solomon goes on to say, "To him, all the singularity people seem like they're just making wild guesses on things like tech expectations, societal implications." We could probably go through PESTLE framework and say where you're making wild guesses across all six. So I guess, if I was to rephrase that question a little bit, it would be... Why does spending time thinking about the singularity, potentially make you more able to predict what is going to happen after the singularity... If, one of the key things about the singularity itself, is that you're saying, "I can't see anything beyond it?"
- Daniel Hulme: 00:33:03 Well, I think I'd like to think that I try not to predict what happens beyond it. I try to give people the boundaries of what could happen. In each one of the six singularities, I don't know whether it's going to be a positive outcome or a negative outcome. But I try to give people the

parameters... It's a world of, for example, in the economic singularity where we have economic or social unrest, because of mass technological employment? Or a world where people have access to all of the goods that they need, so they don't have to have social unrest at.

00:33:35 I don't know what world we're going to create ultimately, it's down to us all individually to make good decisions. Hold ourselves account or hold governments account, leadership account, for making sure that we're making these good decisions. I hear this point a lot though, which is, why are technologists thinking about social questions, or philosophical questions?

00:33:54 And I think obviously technologists at the moment have a loud voice, because AI has happened. But I'm a massive advocate of bringing together the most diverse group of experts, to get them to surround a question and trying to answer that question. So I concur. I think we should be careful about listening to the people that are developing the technologies, even though that's the thing that's popular right now. And we should be engaging with historians and psychologists and philosophers, and all of the diverse aspects of humanity, facing to some of these questions.

00:34:33 And Conscium, for example, the consciousness company, I think I've managed to attract an incredible diverse group of experts that actually don't agree on what consciousness is. But the idea is that you get these different perspectives, these diverse perspectives to make sure that we've got the right eyes looking at this problem. So I'm a massive fan of diverse perspectives. I agree that technologists are not necessarily the right people to be asking some of these questions, but then I try not to answer them. I try to give people the parameters.

- Jon Krohn: 00:35:07 Yeah, it makes sense. I guess my answer to this kind of question is that even though there's a lot of uncertainty, spending some time thinking about this, and some time researching it... Surely, gives at least some higher probability of being able to make predictions. Or play the whack-a-mole game better, maybe even just in the run up to the singularity happening, to set us up for a good singularity than if we don't do it at all, if you deliberately stay naive, if you say, "You know what? This thing's going to be a poop show, and there's no point in even trying. I'm just going to actually spend all my time at the pub, watching footy, and that's going to be it. That's how I'm going to spend out my days till the singularity. That's my favorite thing to do." Which sounds pretty good, but in that scenario, you're surely not going to have any insights. You're surely not going to make any impact. And so at least if we try, it could bring some benefits.
- Daniel Hulme: 00:36:16 Indeed. I guess one thing to mention is that I don't know when and how these singularities are going to be created. What might argue we're already potentially living in some of these singularities now. But in terms of the intelligence level of these AIs, one thing that I try to do, is think about how the intelligence of these AIs are going to play out over the next four, five, six years so that I can make sure that we're making good decisions from a business perspective, to capitalize on those intelligences.
- 00:36:45 And in terms of my predictions, I feel like at the moment, Large Language Models, are a little bit like intoxicated graduates, and I think that over the next six months those graduates are going to graduate to a master's level. Reasoning is going to become really important, and the next sort of battleground for AIs. I think that maybe another 18 months after that we might get to a PhD level, and I think that OpenAI has already announced that GPT-5 is going to have PhD level intelligence, or capability. Where you're able to give it a complex task,

and it can be able to go out there and create its own experiments, set its own hypothesis, and solve that problem.

00:37:28 Maybe 18 months later, it will have a postdoc level and another 18 months we'll have a professor level that has the creative capacity, the depth of thought of all of the professors that we have across the globe. So that's sort of my prediction in terms of the intelligence development. I can't claim at all how that's going to impact humanity. So what I'm not doing is, I'm not saying, and that will cause these things to humanity. Which I think was what the question was. I think I've got the right to sort of predict how I think these technologies will evolve, but I don't necessarily have the right to show, or determine how that will impact business and society.

Jon Krohn: 00:38:13 Yeah, great points. A related question from Jepson Taylor himself. So first of all, he says that this episode is going to be epic, and I agree, it probably already is. And then Jepson wants your data estimate as to when we'll have a Superintelligence?

Daniel Hulme: 00:38:31 Oh, it's a really tricky one, isn't it? So going back to this idea that we'll have potentially a professor in our pocket the end of this decade, which one might argue is AGI. I think there are different definitions of AGI. One is as smart as all of us, or have the capability of all us now. A professor Large Language Model, can't cut hair or it can't necessarily drive cars. So anyway, but let's assume that we'll have a professor in our pocket by the end of this decade. Now the natural question you would ask a professor is, "Go and build an AI that's smarter than you."

00:39:06 And there's a concept called the fast takeoff, which within minutes, days, weeks, we might end up in a situation where you've got AIs creating ever smarter AIs. So I think

that if my prediction is right, and we are going to track these sort of intelligence levels over the next five years, it could be that in 2030, we'll see Superintelligence.

- Jon Krohn: 00:39:34 It is possible. And it is possible, we'll see. We'll see if scaling continues to work as well as it has up till now, it might. We might need other kinds of techniques. It's interesting. I hadn't heard that... Opening I was making claims on what GPT-5 would be able to do, like have a PhD level intelligence. It stands to reason, based on everything that we've seen so far with scaling. So taking the transformer architecture, scaling it up, more training data, more model parameters, we see at each step of the way, GPT-2, GPT-3, GPT-4, big advances in capabilities. It stands to reason that, that will continue with GPT-5. It's interesting to be making kind of expectations around what that level will be, because you don't know till it's done. And so GPT-4 did lots of mind-blowing things, that people didn't anticipate it would be able to, and same with GPT-3. So it'll be interesting, it'll be interesting.
- Daniel Hulme: 00:40:25 And I think reasoning is also going to play a big part in this, at the moment. It's a terrible example, but if you have in your dataset the ability to confidently predict that Socrates is a man, and confidently predict that all men are mortal. Unless you have it in your dataset to confidently show that Socrates is mortal, it's not very good at making that inference. Which is why it hallucinates. What reasoning will allow us to do, symbolic reasoning will allow us to make those inferences. And I think that it will also mitigate a lot of these issues around hallucinations, and things like that. But I think we're going to see a step change in intelligence, over the next 18 months.
- Jon Krohn: 00:41:04 Is Methuselah a man?

- Daniel Hulme: 00:41:07 Methuselah? I think it was coined by actually Aubrey de Gray. Aubrey de Gray, I think it comes from... Yeah, I'm not sure where it's come from. It was Aubrey that coined it, who's obviously one of the godfathers of longevity.
- Jon Krohn: 00:41:19 Yeah, yeah. I saw him speak in person at the St Gallen Symposium in Switzerland in 2013 or '14, and it was mind-blowing. He was the first person... It was a didactical session, where the person interviewing him was a journalist, and was skeptical about these ideas of immortality. And Aubrey de Gray, to every question, it was clear that he'd encountered this kind of resistance, these kinds of questions a hundred times before. And so he had such well thought through answers with clear analogies. I'll try to find a way to... I think that, that is on YouTube, so I'll try to dig that up-
- Daniel Hulme: 00:42:00 Great.
- Jon Krohn: 00:42:00 ... and put it in the show notes for people. Yeah, Aubrey de Gray, fascinating thinker, and he brought his mom to the Symposium. So I'm pretty sure, now that's a decade ago, but I'm pretty sure that's who it was, and had drinks afterward. And I think, if I remember correctly, was grabbing a scotch and he was like, "This is my mom."
- Daniel Hulme: 00:42:24 That sounds like Aubrey. I think he's a good fan of scotch and maybe that's one of the secrets is, maintaining those loving connections with your family. I think there is an indication that-
- Jon Krohn: 00:42:33 No doubt.
- Daniel Hulme: 00:42:33 ... social connections is super important to longevity.
- Jon Krohn: 00:42:37 No doubt. All right, so we've come a long way here. We used the PESTLE framework, and then kind of came off that to go off on tangents. Which were largely tangents

that I wanted to cover anyway, in this interview. So the PESTLE framework for the singularity, we've now gone over all six of those letters in the acronym. So political, environmental, social, legal, and economic singularity. Thank you for taking us through all of those. And in doing that, we also got into talking about alignment. And so there's a lot of discussion around alignment in AI, aligning AI to common human values and goals.

00:43:19 Like you said, there's different ways of thinking about this. Is it just maximizing curiosity like Elon Musk suggested some time ago? What are the kinds of things that we need to do, to get an AI aligned with our goals? But in the way that our society is set up today, we tend to reward short-term profits, over things like long-term sustainability. And you have mentioned elsewhere, the potential for an ultra capitalistic system, that avoids the pitfalls of the traditional capitalism that you're in now. Can you explain to our audience how decentralization and tech, can create a more efficient and equitable allocation of resources?

Daniel Hulme: 00:43:58 Yeah. So there's two things here I want to talk about. Well, maybe we can come back to the alignment problem in a bit, because I want to share an idea about a way that we might be able to capture human morality. So I'll come back to it in a bit. But I was, for a long time... So I'm very interested in decentralization, because I've worked in large organizations, I've worked in small companies, I've had my own startup. And I think a lot of companies start in the same way. They want to be decentralized, they want to get processes and bureaucracy, and hierarchies out of people's way, so they can go and do the things that they need to do. But of course, as you get bigger as an organization, you end up creating hierarchies, putting in these choices and structures, that slow organizations then down. They prevent them from being innovative.

- 00:44:49 And so for the past 15 years, I've been trying to figure out how could we use AI to create decentralized organizations? What I mean by that is, organizations that are able to identify the best diverse group of experts, to be able to make the decision. Whether it's feedback or hiring, or firing or pay, rather than those decisions being made in the hierarchy, can you identify those people? It is often called a liquid organization, or liquid democracies. And I think that now AI can, by the way, I think that we are able to now ingest all of the digital footprint that exists across an organization. And an AI could make sense to it, and they could say, "Look, you've worked very closely with that person. You are very knowledgeable about their domain, you understand about the company strategy, so therefore you should have more rights to a say in their salary, than somebody else."
- 00:45:40 So I'm very interested in how to do that. And the reason why I'm interested in how to do that, is because... Let's assume that over the next decade AI will free people up from tasks and maybe even whole jobs. One of the things that I'm interested in, by the way, that's a good thing, because it means that we can reduce the cost of goods. But if it happens very quickly, as I've mentioned before, it could create some sort of a social mess. So how can we use AI to identify what granular pieces of work people could be doing, and recommend and enable those people to do that work?
- 00:46:14 So rather than it being a job role, and you understand people's skills, their plethora of different ways that they can contribute, and then decentralize that across different tasks that need to be done across an organization. And one of my dreams, was to try and figure out how could you scale up to a planet? How do you create a planet where we are free to be able to go and contribute in different ways, and facilitated by AI? And

that is something I still think deeply about, and we're still working on it in Satalia.

- 00:46:45 But what's interesting, what's happened over the past several years is the birth of open-source. And I think this is sort of a much more interesting and semi-related concept. Which is, there's a reason why Meta open-source Large Language Models, it's not necessarily for the good of their heart, it's because it doesn't hurt their business model. If Google and Microsoft open-sourced their models, then it potentially has a revenue impact.
- 00:47:14 Meta can do it, because their business model is not dependent upon it. And what happens with Meta is, is that they're able to now capture data, more data, which is, as we know, makes AI smart, data is valuable. But also they can access talent, that are contributing to those open source technologies. And what I'm seeing is this impulse, more and more, to actually open-source innovations, rather than just sitting on them, until somebody else comes along and out innovates you. You open-source it, you make it available to your competitors, et cetera.
- 00:47:48 Because what's valuable is the data and the talent. And ironically, it's that open-sourcing, which is a very sage business decision. It's that open-sourcing, that also creates a world of abundance, like Large Language Models, which are perhaps the most intelligent technology we've ever built, are pretty much free now to everybody. Or they will be free to everybody soon. That power is in the hands of everybody. And there's a possibility where... I don't know how many phones you go through every few years, but I go through a couple of phones every few years.
- 00:48:23 Now, if we had those phones, if we gave those phones to people that don't have computers, if those phones were



connected via the internet... Then not only is education then free, because pretty much all education is free on the internet now, anyway. But those people then have access to these technologies to get them to go and create new innovations. And so this idea of open-sourcing and decentralization, I think is a very powerful idea. But what we're seeing is actually the capital model, the capitalistic model, forcing the open-source movement.

- Jon Krohn: 00:48:57 Since April, I've been offering my Machine Learning Foundations curriculum live online via a series of 14 training sessions within the O'Reilly platform. My curriculum provides all the foundational knowledge you need to understand modern ML applications, including deep learning, LLMs and A.I. in general. The Linear Algebra, Calculus and Probability classes are now in the rear view mirror, but Statistics and Computer Science classes are all still to come. Registration for both of the Stats classes is open now, we've got the links in the show notes. Intro to Stats will be on August 21st, Regression and Bayesian Stats will be on September 11th. If you don't already have access to O'Reilly, you can get a free 30-day trial via our special code, which is also in the show notes.
- 00:49:42 Beautifully said, a wonderful vision. So these ideas of getting our used phones into the hands of people all over the world, an internet connection, and then access to free education. And then those LLMs can be taking whatever source language that educational resource is in, Mandarin, English, what have you. It could be translated into whatever dialect that person has in their maybe historically marginalized region, and be taking advantage of, "Okay, if I see this symptom, this is how to treat it." And, "Oh, there's a local plant that actually has that treatment." Or, "This is what I'm seeing on my farm, this is how I can grow better crops."

00:50:28 It is really exciting, and I think it is gradually happening more and more. Something that I often think about, and I guess this is related to the economic singularity being realized, is... And again, maybe this is related to the mission of this podcast, I guess, which I have never thought about stating explicitly. But this idea of, if we can be switching our attention more actively to trying to make a big positive impact, as opposed to trying to maximize just economic income.

00:51:04 It's interesting. It seems to me there's so many people who are interested in making a social impact, there's so much vast potential. And yet when you think about your typical week, you, me, anyone listening... How much of that time is actively spent ideating on, "How can I make the biggest impact? And what can I do with my skills that will make the biggest impact this week, this month, next year?" I don't know. I don't spend as much time on it.

00:51:36 It seems like it's almost something that even through education systems, you spend a lot of time on things like partial derivative calculus. Which okay, cool, important, but automatable, there's, there's so many kinds of these social things of like, "How can we structure our week so that we're eating well, getting enough sleep." And maybe thinking about how we can solve, getting people into the habit of getting a journal out, and thinking about how they can make a big positive impact. Instead of thinking about mortgage rates and buying a bigger home, and getting the next car?

Daniel Hulme: 00:52:17 I think that's a really good point because, and then people say to me, "Look, Daniel, there's always going to be people wanting more stuff." And I think there are people that want more stuff. Most of the people I know that want to have a gold-plated Lamborghini, they need therapy. They need to work through those issues. Most of the people I know that have become independently wealthy, they

know that once you reach a certain amount of wealth, having more stuff makes you unhappy.

00:52:46 And I know it's a cliché, we see all of the time on the internet, people that have made lots of money. But having another car, having another bedroom to clean, actually, it is genuinely a burden. And I think that there's already been studies that have been done on this. Which is, getting to that level is not as high as what people think. And I think even in the U.S., they did some studies to say something... It's still a huge amount of money, but if you have \$700,000 in your bank, you sort of feel economically free.

00:53:17 Now the question is, could we use AI to create that abundance? Could we use AI to make energy free, and people have access to food for free, et cetera, et cetera, so we don't have to worry about paying for it? And the evidence that I have, based on at least the community that I engage with, is that people tend not to want more stuff. What they want to do is, they want to have authentic connections, they want to contribute to humanity. And I'm not saying that we should all contribute to humanity, you can do whatever you want with your time.

00:53:47 But let's try to get to the point where everybody is free to do what they want. I'm a Trekkie, I like Star Trek, and I think Jean-Luc Picard said, that in that time that he existed, he said, "The acquisition of wealth is no longer a driving force in our lives. We work to better ourselves and the rest of humanity." And that's really the world that I want to create.

Jon Krohn: 00:54:09 The Next Generation is the best series, right?

Daniel Hulme: 00:54:11 That's the... Exactly.

- Jon Krohn: 00:54:12 Yeah. I feel the same way too. It's so tough, because I think we're probably about the same age, and so that's kind of the one you grew up with. And so it's so easy, it's like, "Oh, that's also the best music." It's these obvious biases towards what you were a teenager for, as you start to become an adult and you have your own independence and consciousness. But I love Star Trek, The Next Generation.
- 00:54:35 I've been thinking a lot lately that I need to rewatch it, because there's so many great lessons there. And like you're saying there, Jean-Luc Picard is leading one of the best starship's that humans have, that Earth has. And after dinner, he goes and practices the flute, or the clarinet.
- Daniel Hulme: 00:54:55 Yeah, indeed.
- Jon Krohn: 00:54:56 And you think about furthering yourself and how noble that is, how valuable it is. As opposed to getting on another call to... Another Zoom call, to pursue another potential source of revenue?
- Daniel Hulme: 00:55:11 I don't know about you, but I've got two little girls, a 3-year-old and a 5-year-old, so I've got an excuse now to re-watch Star Trek. They've got all of that to look forward to.
- Jon Krohn: 00:55:20 Nice. Very nice. Yeah. So we've come a long way here in the conversation. Most recently we've been talking about alignment, and this liquid structure that you champion at Satalia. And that is becoming maybe easier and easier to administer, thanks to AI. And you mentioned there that you could, today already. And I agree with you, that we have the capabilities in LLMs alone, to be able to... If you were to ingest all of the important parts of your company's digital footprint, and I think it's that. It's the gluing the pipes together, it's a data engineering problem

of getting all the slack conversations, all the phone calls, all the Zoom calls.

00:56:04 But if you get all of that information, and you're feeding it into a Large Language Model, the attention spans are great enough. And you could frame the problem in such a way that you could digest each of those streams into digestible parts. And then those summaries can then be digested by a follow-on LLM, and it could be making recommendations. Like you said, on who should be getting promoted, or any of these kinds of decisions that we have inside an organization. Or certainly advising on them.

Daniel Hulme: 00:56:34 You reminded me about the other thing I wanted to talk about, which it was the alignment and morality. Do you want me to just-

Jon Krohn: 00:56:37 Of course.

Daniel Hulme: 00:56:37 So one of the initiatives that Conscium is kicking off, is an initiative that we call Moral Me. You might have come across the MIT Moral Machine? MIT launched this website that essentially was the trolley problem, where you have a choice... Do you kill a robber that's just robbed a bank, or two children crossing the road? And it gave you lots and lots and lots of scenarios, and it would crowdsource the answers.

00:57:08 So what we're doing is, we're actually going to launch an app soon called Moral Me, which allows people to submit moral dilemmas. And that could be any moral dilemma that you can imagine. Let's imagine you're sitting on an airplane economy class, and you've been upgraded to business class. Do you decide to give that to your neighbor, who's never been in business class or something like that? That's a moral dilemma. So we're going to allow people to upload, submit moral dilemmas,



and Reddit, by the way, is a fantastic source of moral dilemmas, if you're a Reddit fan.

00:57:42 And then we're going to obviously then crowdsource people's answers for those moral dilemmas. And of course your demographics, your religious beliefs, your age, your personal circumstances will determine or dictate, how you answer those moral dilemmas. But the idea is that there'll be some things that we agree on as a species, there'll be some things that we don't agree on as a species. And then we might be able to identify patterns, or groups of areas where people disagree. And maybe even start a dialogue to try to get people to come to some agreement.

00:58:17 But what's really nice about this is, that you create a whole lot of test cases, and the human benchmark against those test cases, those moral dilemmas. That we can then run a Large Language Models against. So we can then submit your Large Language Model, we can run it across all of these moral dilemmas, and see how the Large Language Model diverges from human beings. And we want to use that as one mechanism, to be able to test the morality of AI.

Jon Krohn: 00:58:43 Very cool. So if we can get these things right, if we can get the alignment right and we can get these data engineering problems solved, and we can be having LLMs potentially not just helping us with business decisions, or actually literally making independently business decisions in our businesses, but potentially across all walks of life. And so you've, in your many talks, a word that appears a lot, or a concept that appears a lot is, decision and optimization in the context of business and AI. You've mentioned that humans are terrible at decision making, and that companies have decision problems, not data science problems.

- 00:59:25 You've emphasized that the real value of AI lies in decision making, therefore, and not just pattern recognition. And so I'm going to read a quote to you that our researcher, Serg Masis pulled out for me. So a recently deceased noble laureate Daniel Kahneman, in his book Noise, he wrote that, "Humans are unreliable decision makers. Their judgments are strongly influenced by irrelevant factors such as their current mood, the time since their last meal and the weather. And so all this chance variability of judgements we call noise."
- 01:00:02 So if Daniel, if we want to be able to convince human experts, judges, doctors... That the studies show, if they've just had a coffee, if they've just had a lunch break, judges for example, are much more likely to let somebody off, or give them a more lenient sentence. So clearly irrelevant factors are influencing the decision-making. How do we convince these experts who have spent their life believing that they've been developing expertise, and that they are the expert, they should be trusted? How do we allow them to recognize their susceptibility to noise, and to lean on AI instead?
- Daniel Hulme: 01:00:43 Oh, that's a great question, isn't it? I think first of all, you need to be curious and you need to go and read, Thinking Fast and Slow, and Noise. And Dan Ariely's, Predictably Irrational, there's plenty and plenty of books out there that show our dysfunctions in our ability to make certain decisions, as well as our irrational or logical fallacies that we fall victim to. So there's a lot of evidence out there. What I'd like to happen, and maybe what will happen over the coming years, is that we're hopefully going to more and more have digital assistance like Pi, Inflection's Pi. On our phone, or helping us in answering questions, and doing things for us.
- 01:01:26 Those digital systems actually are not digital systems. They are digital twins, they're learning of you, your data,

and ultimately they're going to try to represent and understand you. Those digital twins though, are going to be potentially be a lot more rational than you. And so perhaps what we could have, is a situation where we have a digital representation of you, that is guiding you and helping you make sure that we don't make poor decisions. Me with my marketing hat on now at WPP, we have to not only figure out how to learn and activate human beings, you have to learn to activate and market to AI. Which is a completely new paradigm shift for marketing.

- Jon Krohn: 01:02:10 Unlock a market to AIs, as the customer?
- Daniel Hulme: 01:02:15 Yeah. Because AIs are going to start making decisions, purchasing decisions on our behalf.
- Jon Krohn: 01:02:19 Right. Right, right. Of course. Oh, my goodness. Yeah. So let's talk about that quickly too. So you're the Chief AI officer of WPP, which is one of the largest media conglomerates in the world. And it's interesting, these kinds of media conglomerates like WPP... My first data science job was at Omnicom.
- 01:02:36 And it's interesting, when you're in the industry and you know B2B companies, you've been in the corporate world for a while... You start to recognize these huge media companies. Which have 100,000 employees each, and there's a handful of them around the world. But until you get into the corporate world and you see that, you're completely unaware of these huge media conglomerates. And WPP, I think, is the biggest of them all?
- Daniel Hulme: 01:03:03 Yeah. Yeah, I actually, don't tell anybody this, don't tell people across... I didn't know what really WPP did, until I joined them. We didn't really have any background or pedigree in media marketing communications. We knew about how do you identify frictions, and how do you apply AI to solving those frictions. And we were very fortunate

that their leadership recognized that we could apply that to media marketing communications, and it's just been a phenomenal success. So yeah, my job now is to coordinate AI across 100,000 people across the globe, which is great. And I should also say that WPP is a co-founder investor in Conscium. They are really interested also in facing into some of these big questions.

- Jon Krohn: 01:03:40 That's really cool. All right, so listeners, this is on you, as each individual listener, you can't let anyone know that Daniel didn't know about WPP, when his company Satalia was acquired.
- Daniel Hulme: 01:03:54 I still claim not to know anything about media market communications, it allows me to ask stupid questions.
- Jon Krohn: 01:04:01 I love it. So it's interesting to hear you say there, about how you need to start marketing to AI decision makers, which is really interesting. I don't know if you can go into that a little bit, but something else that media marketing companies need to do is, they're responsible for creativity, so down to literally the ad copy and how ads are going to work online, or on TV. All these kinds of things are devised by humans in companies like WPP.
- 01:04:32 And so how can AI enhance creative processes within these organizations? And it seems like potentially your experience as a lecturer at Singularity University, kind of come full circle here, around the singularity. Yeah. How can AI unlock creativity in the workspace? We can, I know we're getting short on your time here, so this can be our final big question?
- Daniel Hulme: 01:04:59 Well, I think one of the reasons why I actually joined WPP, is because I really love their mission and their purpose, which is to use the power of creativity to make a better future. And I think a better future is a world where everybody is free to be able to live beyond themselves.

Using their creativity to then contribute to humanity, and often people think that, that's called a utopia. And I think none of us would agree what utopia is. We probably agree what a dystopia is, but I really like this concept called a protopia. Which is the idea of a system that is engineered and that can just get better and better.

01:05:32 And I think the more people we can free, the more we can enhance and unlock the creative capacity of human beings, the more that they'll use that energy source to make the world better. So the WPP is all in on creativity, and I'm very, very fortunate that my day job is figuring it, how do I enhance, accelerate human creativity? Not only across obviously 100,000 people, but the question that I'm asking myself as well is, "How can we do that, for the whole of humanity?"

Jon Krohn: 01:05:59 Very cool. Quickly, book recommendation for us before we let you go?

Daniel Hulme: 01:06:02 Oh, I have here on my shelf, Artificial Intelligence: A Modern Approach. But this was my textbook when I was undergraduate back in 1999. But actually I think that Behave, by Robert Sapolsky. I think understanding who you are as a human being, what motivates you, what you're good at, bad at, to allow you to make better decisions, is a really great book. Behave.

Jon Krohn: 01:06:30 Awesome. Great recommendations. And how should people follow you after this episode? You're pretty big on LinkedIn, is that the place to go, or where else?

Daniel Hulme: 01:06:39 LinkedIn, email, happy to share my number, anything. Grab a coffee anytime.

Jon Krohn: 01:06:43 Wow. What a tremendous offer there. That's wild. All right, Daniel, thank you so much for taking the time. It has been mind blowing to have you on the show.

Hopefully we can catch up with you in the not too distant future. Hopefully we won't be in a post-singularity world, before you and I next catch up on the air. It'd be great to hear how things have evolved, and how your thinking has shifted over that time.

Daniel Hulme: 01:07:06 Thanks, Jon. It's been great.

Jon Krohn: 01:07:12 Wild, fascinating conversation with Daniel today, and he filled us in on following the PESTLE framework, the six types of singularity that would be unleashed by ASIs. Specifically political, environmental, the social methuselarity, in which we could live indefinitely. The technological singularity, the legal one, and the economic one. He also talked about how neuromorphic computing makes computer chips work more like human brain cells, potentially allowing AI systems to eventually learn from a single example on the power of a light bulb, like our human brain does.

01:07:45 And he also talked about how work could be transformed. Not just in the future, but this very day, by engineering all the relevant corporate data to flow into LLMs. So that AI systems could advise human decision makers in real-time on things like commercial, and HR decisions. 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 Daniel's social media profiles, as well as my own, at superdatascience.com/807.

01:08:12 Thanks of course to everyone on the Super Data Science Podcast team, our podcast manager, Ivana Zibert, media editor Mario Pombo, operations manager Natalie Ziajski, researcher Serg Masis, writers Dr. Zara Karschay and Silvia Ogweng, and founder Kirill Eremenko. Thanks to all of them for producing another spectacular episode for us today. For enabling that super team to create this free



podcast for you, I'm so grateful to our sponsors. Thank you so much. You could support this show by checking out our sponsors links, which are in the show notes. And if you'd ever like to sponsor an episode of the show, you can see how to do that at jonkrohn.com/podcast.

01:08:49 Otherwise, please share. Please review this episode on your favorite podcasting platform, or YouTube. Subscribe to this podcast of course, if you aren't already a subscriber, but most importantly, just keep on listening. I'm so grateful to have you listening and I hope I can continue to make episodes you love for years and years to come. Until next time, keep on rocking it out there, and I'm looking forward to enjoying another round of the Super Data Science Podcast with you, very soon.