SDS PODCAST EPISODE 969: THE LAWS OF THOUGHT: THE MATH OF MINDS AND MACHINES, WITH PROF. TOM GRIFFITHS

Jon Krohn: 00:00:00 Psychologists will tell you humans are terrible thinkers, irrational, biased, constantly making mistakes. Computer scientists meanwhile will tell you humans are incredible, inspiring enough to model the entire field of AI around. Today's guest lives in both worlds and he's figured out why they're both right. Welcome to episode number 969 of the SuperDataScience Podcast. I'm your host, Jon Krohn. Our extraordinary guest today is Tom Griffiths, professor of psychology and computer science at Princeton University and director of the Princeton Laboratory for Artificial Intelligence. He's the author of the megabest selling book, Algorithms to Live By, which we discuss in this episode. But the main focus is on his sensational brand new book, The Laws of Thought, digging into mathematical models of both biological and artificial intelligence, as well as the implications today and tomorrow. Enjoy.

00:00:53 This episode of Super Data Science is made possible by Dell, Intel, Cisco, and Acceldata.

00:01:02 Professor Tom Griffiths, welcome to the SuperDataScience Podcast. It's seriously an honor to have you on the show. Thank you for taking the time. How are you doing? Where are you calling in from?

Tom Griffiths: 00:01:10 Thanks. It's great to be here. I'm calling in from Princeton, New Jersey.

Jon Krohn: 00:01:13 Yes, yes. Where you are faculty, we're going to talk about that a fair bit in this episode. You're the director of the Computational Cognitive Science Lab at Princeton, which is a research group focused on understanding the mathematical foundations of human cognition. Fascinating and the central focus of today's episode. But you're also director of the Princeton Laboratory for

Artificial Intelligence, which sounds like it's a newer lab that's focused on innovative research efforts in AI and related fields. Do you want to tell us a little bit about your research?

Tom Griffiths: 00:01:48 Yeah. So the AI lab is our internal research accelerator for AI. So the university as a whole is trying to figure out how do we get the most out of AI for the kinds of things that we do on campus. And so our job in the AI lab is to find the problems that have unusual potential for impact from using AI-based methods and then give them the support that they need, people compute the two things that are the most common requests to be able to make it possible to do that kind of research. And then my own research is right at the intersection of psychology and computer science. So it's trying to understand how human minds work, using the kinds of ideas that come from computer science, but also trying to think about how to make computers smarter using the insights that we get from studying human minds.

Jon Krohn: 00:02:38 And so you've published books at the intersection of both of those things. At the intersection of AI in the human mind, you already have a bestselling book called Algorithms to Live By. We're going to get to that later in the episode. But first, I want to talk about your brand new book, which also is at the intersection of both of those fields. It's called The Laws of Thought: The Quest for a Mathematical Theory of the Mind. Fascinating. And in it, you trace the history of logic and statistics from George Boul to modern AI to define the fundamental laws that govern both biological and artificial intelligence. In this book, you open by noting that we're fluent in understanding the physical world, but far less so the mental one. Do you want to tell us a bit about that?

Tom Griffiths: 00:03:23 One of the surprises that people might have is that we're used to thinking about using math to describe how the physical world works. So if you go to school, you might take classes in chemistry or physics, you learn the laws of nature. And the premise of the book is that there's a complimentary set of ideas that really originated around the same time that people were thinking about these laws of nature being sort of mathematical principles describing the external world. And those are what we could call the laws of thought, the mathematical principles for describing our internal world. And some of the same people who are interested in using math to make sense of physics and these other kinds of things were just as interested in using math to understand how minds work. It just took a little bit longer to figure out how to solve that problem, really up through the 19th century until we had our sort of first glimmers of what solutions might look like, and the 20th century to turn that into an empirical science, what we call cognitive science.

Jon Krohn: 00:04:17 There's this word that's coming into my head as you've been speaking. And so I did an undergrad in neuroscience and a PhD in neuroscience, but it was a long time ago and I was focused on machine learning applications to the field of neuroscience. So this word coming into my head, I don't know how relevant it is here, but psychophysics.

Tom Griffiths: 00:04:37 Yeah. Psychophysics means the sort of intersection between psychology and physics that happens around perception, right? So if you think about what perception is, it's the marriage of these two things, like your visual perception, a photon comes into your eye and hits your retina, right? That's a physical process, but then you perceive that as a flash of light, that's a psychological process. And that was actually one of the places where the science of psychology started out was trying to make sense of these kinds of interactions between the physical world and our subjective experience. So the earliest psychology labs ran experiments where they would ask people to kind of report back what they experienced when they were experiencing different kinds of physical stimuli. And the rest of the science of psychology really evolved from that point.

Jon Krohn: 00:05:29 Right. So the psychophysics is really about perception, these kinds of first layers of neural processing in your books and in your research, it sounds like in general, you're more focused on deeper interplay, our subconscious or conscious thoughts, the continuous flow between our frontal lobe and our other sensory lobes and being able to have some kind of organized understanding of the world.

Tom Griffiths: 00:06:00 Yeah. So those first psychologists really, as I said, we're trying to characterize what was happening with perception, but we're also interested in what's happening in thought, but it turned out that thought was a much harder thing to study. And so at the start of the 20th century, another group of psychologists said, in fact, relying on people's subjective experiences here, we want to understand things like thoughts and feelings, that's not going to get us very far if we want to have a rigorous science. And so they said, let's get rid of those things. This is a movement called behaviorism. Let's just focus on the things that we can measure so we can measure people's behavior, we can measure the environments that they're in, and that's the route that we're going to have to making a rigorous science. And so they kind of just got rid of all of the things that I'm interested in, these kinds of internal processes of thought and so on.

00:06:45 And so part of the story that's told in the book is the story of how we got past that and got to the point where we could have a science of the mind, and it was really using mathematics that made that possible. So ideas like logic and the development of computers gave psychologists a new language for being able to talk rigorously about things like thoughts and a way of expressing precise hypotheses about how thoughts might work and how language might work and things like that. And so that then resulted in a new way of doing a rigorous science of the mind. And that's the foundation for the kind of work that I do where we ask these questions about the connections between human minds and computers.

Jon Krohn:	00:07:27	Fascinating. Could you give us one or two concrete examples from your research perhaps of exactly how you do this research? I mean, do you have people describing experiences they're having and then trying to figure out mathematical equations from that?
Tom Griffiths:	00:07:42	So most of what we're doing when we're sort of working in my lab is going back and forth between thinking about the kinds of computational problems that human minds have to solve and what the ideal solutions to those problems look like, and then running experiments where we test out how well does that line up with what people actually do. And I say experiments, what those really look like is maybe something that's a little more like a survey where we'd sort of show you something and we say, "What's your answer to this? " In sort of data scientist language, we're sort of giving people data, we're asking them to form hypotheses and we want to know what the conclusions are that people are drawing from those data. And so you can think about that as a problem where we're then trying to work out, if you think about the mind as an intuitive data scientist, right, what are the tricks and methods that the mind is using to make sense of the data that are presented to it?
Jon Krohn:	00:08:36	So undergrad students are typically the kind of the workhorse of these kinds of psychological experiments. They seem to do the line share of it. So in this kind of example, you would have undergrad students come into

your lab and they would be presented with questions and then you ask them to provide a response.

Tom Griffiths: 00:08:52 That's the way we used to do it. Nowadays, we run most of our experiments online and part of the reason for doing that is that we can just get much larger samples more quickly. So one of the ways that psychology has changed over the last, just even the last decade is by having access to increasingly large groups of participants, it allows us to ask these questions in ways that go beyond what we could have previously asked in the lab. And that's one of the things that we've really focused on in my lab is thinking about, okay, if you have access to much larger populations of people, how does that change the way that we do our science? We've done things like run experiments where we look at the decisions that people make, where we ask people to say, choose between different gambles. So you might get $5 as a sure thing, or you could have a 50% chance of $10 and a 50% chance of getting nothing at all, right?

00:09:50 And what choice would you make between those things? And so historically, those kinds of questions about how people make those decisions had been based on experiments that people had run with undergraduates and maybe there were a couple hundred decision problems that had been presented that were the basis for developing a theory. In my lab, we collected a dataset which involved more than 10,000 of these by asking people to do this online. And then we were able to use methods like artificial neural networks, some kinds of tools that come from modern machine learning to be able to analyze those data and come up with new theories.

Jon Krohn: 00:10:20 Right. So using modern AI systems to be able to analyze all of these data. The specific example that you gave there, I don't know if that's like a ... That sounded kind of like an economics example. It reminds me of Daniel

Cannonman type of experiments figuring out, I think since the '60s, Canamin and Taversky, we're focused on that kind of research where they're trying to figure out how econs, how perfect economic thinkers might make their decisions in the world. Does economics tend to be a big part of the kinds of questions you ask or is there a broader diverse set of questions that you carry out in your lab?

Tom Griffiths: 00:11:00 Yeah. So we're interested in questions that are really about how people make inferences from data and how they make decisions. And again, I think that maps pretty well onto the problems a data scientist faces. You have two kinds of problems. There's what kind of inductive inference is warranted by the data that you have, and then what action you take on the basis of that discovery. And it shouldn't be a surprise that those two things are pretty fundamental to the kinds of problems that human minds have to solve. We're going around, we're trying to figure out how our world works, and then we're taking actions on that basis. And so we study things like how people learn causal relationships, how people form categories, how people make inferences about what other people are thinking. These are all things that we can cast as problems of a kind of statistical inference.

00:11:50 And then we also study things like, yeah, under what circumstances do people make particular kinds of choices? One contrast with the common and Taversky tradition is that we're actually interested in asking whether we can make sense of those decisions in a way that actually seems maybe a little more rational, whereas I think a lot of the kinds of strategies that people use have been sort of held up as things that people are doing that are irrational. And if you think about it, so I live in these two worlds. One is the world of psychologists and one is the world of computer scientists. And the way we look at the mind is quite different in those two worlds. So if you go and talk to a psychologist, you're going to be told people are terrible decision makers, we're irrational, we do all these things that are kind of dumb.

	00:12:33	If you go and talk to a computer scientist, they'll tell you humans are impressive. We want to be able to build systems that can do the kinds of things that humans can do, right? Humans are inspiring in all of these ways. And so I try and bridge the gap between those two perspectives, understanding why it is that people make mistakes and how it is that people are nonetheless so impressive by thinking about human cognition as doing a good job of solving the problems that are posed by the world with the limited computational resources that are available to us.
Jon Krohn:	00:13:03	Really interesting. What a fascinating area to be involved in. You must really enjoy it a lot. And I mean, it shows in the quality of your writing and obviously your research output and those kinds of things. When people are doing these online surveys, how long does that take? Would it be possible for me or for our listeners to be able to do these? Can we just go to a website somewhere and do them?
Tom Griffiths:	00:13:27	Yeah. We mostly use crowdsourcing services to recruit participants for our experiments. And so anybody can sign up for those services. One of the main ones we use is called Prolific, and it's a service where people can sign up. You will see a listing of experiments you can participate in as well as other kinds of tasks. You get paid a little bit of money for doing each of these experiments. And so if it's something people are interested in, there's definitely opportunities to participate.
Jon Krohn:	00:13:54	Now, so that's what I was going to ask next is what is the incentive? And so if people are getting paid, then there's also a chance that people are trying to figure out how to

engineer LLMs to be doing this research. So how do you guard against that kind of fraud?

Tom Griffiths: 00:14:10 That's a legitimate challenge. It's something where every time we come up with a solution, people find a way of getting around that solution. So you can think about it as a kind of arms race between the experimenters and the people who are running the bots in our experiments. I have some former students who started a company that is actually focused on solving that problem. And so they have lots of nice tools for doing this, but it's when we just try and design our experiments in ways where we're asking certain kinds of questions that are going to catch out an LLM. I'm not going to tell you what those questions are because as soon as I tell you, we give away our secrets.

Jon Krohn: 00:14:49 Right. Well, yeah. So you have answered my question there. And it is a question of great importance and increasing importance. I mean, huge amounts of our economy now depend on, say, for example, data labeling services. There's millions and millions of people all over the world who are getting paid to label data for the frontier labs. So for Meta, for OpenAI, for Anthropic, for Google, billions of dollars are being spent to get people to solve ... If you think back a decade, it would have been very simple questions mostly like, "Is this a dog or a cat for a machine vision algorithm?" But these days, you could end up paying a PhD student in chemical engineering for a day's work to come up with a multi-step chemical engineering problem and answer for a multi-step reasoning algorithm. Algorithms like a one or 03 from OpenAI are kind of the popular kind of reasoning algorithms there.

00:15:51 They require very complex questions and answers and you need to be confident about the accuracy. And so you're paying highly intelligent people a lot of money to create these data for these algorithms. And yeah, of course, everyone in that kind of system is incentivized to make their life easier, either getting assistance from LLMs as they do it or outright using them. You could imagine if you could figure out how to hack this, you could have it at a very large scale, LLMs harvesting money from these kinds of labs. So yeah, big, big fraud problem. And so I'm sure ... Do you know offhand the name of the company that your students have?

Tom Griffiths: 00:16:33 It's called Roundtable, although I should disclose that I'm a very small investor in the company.

Jon Krohn: 00:16:37 Roundtable. They've got a big market and a growing market to be providing solutions to. Really cool. All right. So we've talked a bit now about how your research works on understanding the human mind, and we've even understood a little bit on how you're using AI tools, things like neural networks to analyze the large amounts of data that you collect online, but you're also interested in going in the other direction. So using your understanding of how the human mind works to understand how AI systems work or to develop better AI systems. And so on that note, an interesting part of your book talked about symbolic AI researchers who were kind of the leading approach to doing AI for decades, but there was a fundamental mismatch between logic and intelligence that they faced, and it took decades to pivot toward probabilistic and neural approaches which now power our modern AI systems.

00:17:37 And so is there something that maybe today we're seriously overlooking like that? And how can we realize that we're making these big mistakes in the way that we're designing our AI systems and how can we overcome them?

Tom Griffiths:	00:17:49	That's a great question. I mean, I think it's not so much that we're necessarily making mistakes, but I think if you look at the history of people pursuing these different approaches to understanding intelligence, what you see is that there's a kind of trajectory where people try one thing and they're like, "Oh, this is the thing." And they go all in on that one thing and then kind of like push it as far as they can and then they hit some limit, right? And then they find another thing, there's another breakthrough that allows them to solve that problem and then they can sort of go all in on that and then keep pushing that until they hit some limit. And so I think that's just a normal part of these kinds of cycles where we come up with something that really seems like it's working well and it does work well, but it also maybe has some fundamental limitations that don't become clear until we've gone as far as we can using that kind of method.
	00:18:38	And neural networks in particular, because they're very hard to understand in terms of like if you write down a model and you write your code for implementing that model, it's really unclear what the limits are and what it might be able to learn or do. It's hard for us to know exactly what those limitations might look like. One of the places where we can get some ideas about what they are is by looking back in history, right? So neural networks, some of the challenges that they have are, they're very sensitive to their training data. Really a neural network is in our data scientist language, it's a non-parametric method, right? It's something where it's able to capture any function that you might potentially want to learn in principle is something you could learn with an appropriately complex neural network. And so it's able to learn in some sense anything.
	00:19:26	But as a consequence, what it actually learns is going to be very influenced by the data that it's trained on. And so there's a lot of things that fall out of that about the ways

in which our large language models sort of act in ways that might be counterintuitive to people, partly because of their sensitivity to their training data. So we have a paper where we looked at this with, when GPT-4 came out, people were very excited about it. They said, shows sparks of AGI. We wrote a paper that we called Embers of Autoregression that was about how the way in which that model is trained is actually influencing its behavior in ways that people might not expect. And so order regression is the way the model's trained that's trained to predict the next token, a word or a part of word based on the previous tokens or of the previous words that it's seen.

00:20:18 And so that means that the model is going to be very sensitive to how probable sequences of tokens are when it's producing answers. And you can show that that has some counterintuitive consequences. One of those is that if you just ask it to count how many letters appear in a sequence, it's more likely to give you the right answer if the right answer is 30 than if the right answer is 29. And the reason is that the number 30 just appears more often on the internet than the number 29, right? So just because of the raw statistics of those numbers, that's something that the model is very sensitive to in producing its output. And that's something that's counterintuitive for humans.

Jon Krohn: 00:20:55 That is counterintuitive. And in all the years that I've been in this space and obviously closely monitoring what's going on in LLMs as a professional and as a podcast host, I had no idea that that happened, but it makes so much sense that more probable terms, something like 30 relative to 29 is more likely to be predicted by the model. And it is interesting how now we don't really know what's going on behind the scenes when you're using something like ChatGPT or Gemini or Claude, you don't know how many different algorithms are being involved. And it seems like probably today, unlike when GPT-4 first came out, there is some kind of mathematical inference happening where you're getting an exact answer as opposed to a probabilistic one.

Tom Griffiths: 00:21:45 Yeah. So in fact, what's happened is that people who work with these models have noticed that they have this kind of deficiency. And so if you ask a modern model to do something like this, so in our paper, we focused on what we call deterministic problems, problems where there's exactly one answer based on the question that you ask. And we show that for those deterministic problems, the models still produce a probabilistic answer, right? They're influenced by the expectations that they have about the distribution of language. So if you're asking them for something which is a low probability answer, they're less likely to produce it. And so that characteristic still holds for modern models. We can show that they're still sensitive to the probability of the output. It's been a little bit attenuated, but if you actually interact with a model through the chat interface or something like that, and it sort of does all its thinking and does all these other things, one of the things that it'll tend to do if you give it one of these deterministic problems is write a piece of code to check its answer.

00:22:40 So it has been trained that for certain kinds of problems, it shouldn't rely on its sort of probabilistic backbone, it should do something which is appropriately deterministic in order to then give you the answer back. But that is also something that I think parallels something that we see in the history of cognitive science. So when people were working with rule-based models in good old-fashioned AI, they realized that they needed to have some kind of structure around this in order to make these models be able to do the kinds of things they wanted them to do. So this was building what were called expert systems in the 1980s based on production systems with these sort of rule-based models. So these had sort of big lists of if then rules. If this happens, then do this. If this happens, then change your goal to this. If this happens, and if your goal is this, then do this.

00:23:28 And so it had all of these instructions that they would follow, and they realized they needed to have some kind of architecture around that that would then allow this system to make good decisions by using all of those rules together. And you can see that being echoed in what's going on with modern language agents, where a language agent is really building that same kind of cognitive architecture around the core technology, which is the base language model.

Jon Krohn: 00:23:53 Really fascinating. I love how you tied it there to the historical expert systems. Something else in your book that straddles both human and machine cognition that we found fascinating in our research was in your chapter, probability is the new logic. You showed how dramatically conclusions can change based on prior probabilities. And you used an orangutan example that was easy to follow. Do you want to tell us about that?

Tom Griffiths: 00:24:17 Sure. I mean, let me maybe start one step back, right? So one of the things that you mentioned was that people started to realize that logic wasn't going to do all of the things that they wanted to be able to do. And so part of that is that when you think about the kinds of problems that human minds have to solve, you can divide them up into two kinds. One kind is what we could call deductive problems, right? Where there are problems where you have enough information given to you that you're able to work out exactly what the answer is, right? You have all the information you need in order to solve the problem, right? So if I give you a math problem like two plus three, you've got all the information that you need in order to work out what the solution is, right?

00:25:00 Or even if I give you a problem with an unknown in it, right? So X plus two equals five, pretty easy for you to work out what X is there. And that's an example of a deducted problem. And so those symbolic rule-based systems were really good at solving deducted problems. You can think about that as you have algebra, logic, chess, these were sort of things that demonstrated the power of these kinds of systems. Where they didn't do so well was another class of problems, which are what we could call inductive problems. And these are problems where you get some information, but it's not enough to know what the answer is with certainty. And it turns out a lot of what humans do is induction, where if it's perceiving the world, photon hits your retina, what happened in the external world? You've got some data that came in, but you don't have enough information to know exactly what it was that was the three-dimensional structure in the world that provided that information to you through your retina.

00:26:04 If you're hearing somebody say something, you're hearing some noisy signal, you're forming an interpretation of it, you're trying to think about what someone was thinking based on the actions that they took, you have some data in the form of their behavior, but you're making guesses about what their mental states are. If you're even something like learning a language, you might have some hypotheses about the structure of a language, your data are the things that you've heard in that language, and you're trying to use the data to constrain those hypotheses. And so all of those are inductive problems, and that's kind of where that logic, symbolic rules approach hit the wall, is trying to explain these different important aspects of human cognition. And so probability theory is the tool that we can use for solving these inductive problems, because it gives us a way of talking about how data should influence our beliefs, even when we don't have enough information to be certain.

	00:27:01	And it does that by essentially allowing us to assign degrees of belief to different possible worlds that we might be in. And then as we get data, those data help us to narrow down those possible worlds in a way that allows us to update our beliefs, reallocate those probabilities across those possible worlds, and be able to draw some kind of conclusions even if we're not able to reach certainty. And so the example I give in the book is it's from an old detective story. I don't know if this is the story, the murders in the rumor, which is one of the first detective stories. I don't want to give any spoilers, but I guess I have to. So if you haven't read the story and you want to read it, you should skip ahead a little bit. But the key thing that happens is that the detective, Monsieur de Pan, figures out that the culprit is an orangutan and does so based on some clues like somebody heard a voice, but they didn't recognize the language, there was some tufts of orange hair that were found, and then that the room had to be accessed via a window that was high up, that would be hard to climb to.
	00:28:24	And so, none of those facts gives you certainty that the culprit is an ranguage, right? No detective is necessarily going to go to a crime scene and see those things and say immediately, "Oh, arangue tangs again." But each of them gives you a little bit of evidence for that conclusion. And then probability theory is a tool that we can use for saying how we should integrate those pieces of evidence in order to update our beliefs.
Jon Krohn:	00:28:49	Right. And so basically some of the things there that would be very unlikely for a human to be able to do, like being able to get into the window, to be able to get into the room while they also have red hair. So all of these probabilities multiply together, the most likely conclusion you can draw is that it's an orangutan.

Tom Griffiths: 00:29:09 Yeah. And the way that we express this in probability theory is via what's called Bayes Rule, which is named after the 18th century minister, Thomas Bayes, who kind of first made the suggestion that the math of probability theory is the way of describing how it is that we should go about changing our beliefs. And Bay's rule says, you start out believing something to some extent, and we'll sort of call that the prior probability of that hypothesis. And then in order to work out how much you should believe it after seeing some data, what's called the posterior probability, you take that prior probability, multiply it by how likely those data would be if that hypothesis were true, and then do that for all of your hypotheses and then just sort of divide them by the sum of those quantities just so that you end up with something that ends up being a probability distribution.

00:29:48 But the key idea is take the degree of belief you had before and then multiply by the probability of the data if that hypothesis were true. So if it's something that a human would find very hard to do, it's very unlikely that a human could have reached that window, but it's very likely that a rangutan could, then that's going to increase your belief in it being an Irangutan, even if your prior probability of Arangutans was very, very small.

Jon Krohn: 00:30:09 Now, that kind of taking advantage of prior probabilities, it seems like the way that we train our large language models today, whether it's the pre-training supervised learning or the post-training reinforcement learning, it doesn't seem like we're taking advantage of priors in any way.

Tom Griffiths: 00:30:25 Yeah. And in many ways, that's a consequence of a lesson that people learned in machine learning that I think is potentially going to come back and bite us soon. So the first point here is that this is exactly right in terms of if we want to characterize why it is that human children can learn language in five years and our AI system requires more like 5,000 or 50,000 years of language input to reach the same state. That difference comes down to priors, to having different expectations about what the structure of language are like. So inside our human heads, there's a whole bunch of stuff that makes it easier for us to learn the kinds of languages that we end up speaking, and that stuff is not inside our neural networks. So there's a sense in which that's by design, which is that prior to about 2012, the dominant approach in machine learning was one where a human would engineer, say, features that went into your machine learning system, and then you would have a machine learning system that operated on top of that.

00:31:32 And you can think about those features that were being engineered as a source of priors, a source of inductive bias that was sort of setting up the kinds of solutions that the model could find. And that was really because we didn't have a lot of data, and so you needed to engineer in those biases in order to make the system work. And then in 2012, I sort of tied things to that point because that's when the ALExNet, this sort of convolutional neural network won the image net challenge, right? And that was a model which was a very large, complicated neural network, although small by modern standards, that was trained on a very large amount of data, although again, relatively small by modern standards, right? But here, the insight was that you didn't need to hand engineer the things that were going in, you didn't need to hand engineer those inducted biases.

00:32:24 If you had enough data, you could learn to solve this problem in an end to end way. The neural network could actually find a better solution than the human engineer could, if you had enough data in order to get there. And so since then, machine learning has been operating in this paradigm of very weak priors, very weak conducted biases. The only kinds of biases you put in the system are through the choice of architecture you make, right? Maybe you use a transformer or a convolutional neural network or something like this. And then you just throw enough data at it that it's able to end up solving the problem. And so that paradigm is very different from the human one, right? You can kind of think about the machine learning engineer's problem is, I'm going to start out here, I need to get here, so I need to get enough data to get me from here to here.

00:33:08 The cognitive scientist problem is I know people learn from this much data, so I have to figure out what all of the things are that are filling up the rest of that gap, right? And we can actually quantify that, right? So if it takes you 5,000 years of speech to train your large language model and the human child learns it from five years of speech, then there's some prior distribution, some information, some inducted bias that accounts for 4,995 years of speech, that's a lot of data that has to be encapsulated there. And so part of what we do in my lab is try and figure out what those biases are on the human side and try and develop technologies that can help us to make neural networks that express those. So one of the techniques we use is something called metalearning, which is a way of learning the initial weights that you put in your neural network in a way that isn't a hard bias, like that previous machine learning approach of constraining the kinds of features that the model would get.

00:34:02 It's a soft bias because it says, "I'm going to start you out in a particular part of the space of solutions. And if I can start you out close to the solution you need to find, you're going to find it easier to then discover that particular solution you want. " And so that meta learning technique tries to find good initial weights for a neural network by creating neural networks that are able to solve a bunch of learning problems where they're learning from much less data for each of those problems.

Jon Krohn: 00:34:30 Fascinating. Yeah. You probably have a perspective on where and when this dramatic shift might happen, but it's obvious to me that the current approach that you described there of say taking 5,000 or 50,000 years worth of training data for an LLM to be able to have the capabilities of a five-year-old human, there's this huge gap and it sounds like something like meta-learning or the other things that you're working on in your lab could be the solution. And I guess it could kind of sweep us out of nowhere just as the transformer architecture did and is now ubiquitous. And so it'll be interesting to see where that goes and the part that you play in it. Actually, my next question that I had already pre-planned follows really nicely from this organic discussion that we've been having now for the past five or 10 minutes. And so in your research, you frame human intelligence as rational adaptation under tight limits.

00:35:26 So explaining biases, heuristics, and few shot learning, so learning from a few examples, which is what LLMs don't do, as optimal use of limited time, limited computational resources, and limited bandwidth for communication. So it seems like that ... I didn't even really get to a question there, but the question that I was going to have, you kind of already started now getting into, but maybe with that context, you can develop on your answer even further. The question is, how should those constraints guide the design of the AI systems that we now build?

Tom Griffiths: 00:35:57 I think one thing is it's not necessarily the case that they have to. So I would think about those constraints as a good way of characterizing what the difference is between human minds and our AI systems, right? So human minds solve a set of computational problems that are characterized by these constraints. We live only a few decades, give or take. We have to do all the things we're going to do using just a couple pounds of neurons, right? And we can only communicate with one another by making weird honking noises or wiggling our fingers on a keyboard, right? Those are the kinds of constraints that humans operate under. And so as a consequence, we have to be able to learn from small amounts of data because we're just not going to live long enough to see lots of data, right? It's not like AlphaZero AlphaGo, which is trained on many human lifetimes of go games, right?

00:36:59 You can't do that. Human GoPlayers are going to be limited by what's possible to learn in a human lifetime. You have to be able to recognize the structure and problems so you can use previous solutions that you found or come up with good sort of efficient strategies for solving those problems using the limited computational resources you have. So unlike Deep Blue, a human chess player can't search 100,000 positions per second, they can search much less than that and they have to focus on finding the good ones to search. And you can't overcome those constraints by being able to transfer the data that we've experienced or pool our computational resources directly because we have to operate under those bandwidth limits. And so we have to come up with good conventions for allowing us to do things like work together and to pull data, things like writing language in the first place, but writing science, starting companies, right?

00:37:59 These are all mechanisms that we can have for being able to coordinate the resources that are available to us. And so that set of things really characterizes in some sense what it is to be a human as an intelligent system. And our AI systems are not subject to any of those constraints, right? They can learn from many human lifetimes. They can add more compute when more compute is needed. They can transfer state or share the data that they train from. All of those things are much more straightforward for your AI systems. So one consequence of that is that you might expect that the AI systems are just going to be a different kind of intelligence from human intelligence, right? Because they're not operating under the same constraints, they're not solving the same computational problems. They're going to be able to solve problems potentially really well.

00:38:42 And in some cases, they're going to solve those problems better than humans, but they're not necessarily going to solve them in ways that are recognizable to humans or ways that are the same as the solutions that humans find. And so that sets up sort of two possibilities. One possibilities we say, that's fine. We just want to engineer the best systems we can. We should just keep on doing what we're doing, keep going in that direction, and we're going to get this kind of divergence between what human intelligence is like and what AI is like, and that's okay. The other approach you could think about is one of saying, oh, well, actually we want some of those attributes of human intelligence to be inside our AI systems. In particular, it might help us make AI systems that just make more sense to us in the kinds of solutions that they find.

00:39:26 And so if we want to do that, then we really need to understand the things that are going on on the human side and think about how it is that we transfer those over. And so we talked about metalearning is a trick for building systems that have inductive biases that are maybe more like humans. That's certainly an avenue that you can go down if you want to do that, but engineering inductive bias is hard. Scaling is much easier than engineering inductive bias. And that's the reason why we're in the current sort of moment that we're in an AI is that that's a technique that sort of works and it's an engineering problem. We know how to solve it, but engineering inductive bias really requires having an understanding of the problems that we're solving and what reasonable inductive biases for those problems are. And it's something that people are still trying to figure out how to do well.

Jon Krohn: 00:40:11 Yeah. Let's talk about this inductive bias thing more because I think this is a really important point and it's something that will probably be obvious to a lot of our listeners as we explain it a bit more. When we think about the LLMs that we have today, they can't, as far as I know, maybe you can correct me on this, but I think they're quite limited in the extent to which they can induce some kind of conclusion beyond something that is in their training data. So that is quite unlike human cognition where we don't just depend on the information given. We can use inductive biases, the priors we were talking about earlier, assumptions to make some kind of guess as to what's happening in situations that are quite different from any that we have ever encountered in our life experience before. And as far as I know, any of the AI systems that we have today really struggle in that kind of circumstance.

Tom Griffiths: 00:41:09 I think in some ways, I would say this is a little more of an open question to what extent this is possible. The things I can say about it are, there are certain kinds of generalization that we know large language models can do. So some of my colleagues here at Princeton have a nice demonstration of this where you can sort of get them to compose ideas together, right? So you can give them a set of properties and then they can produce something new, which is the consequence of combining those things together. And it's certainly not something that ever existed because those things hadn't been combined together before.

Jon Krohn: 00:41:41 For sure. That is something that kind of that blending of content and style is very impressive. So something like, give me, explain this Shakespeare sonnet in the style of Snoop Dogg or something like that, which for sure you end up with an output that is unlike anything we've ever seen before.

Tom Griffiths: 00:41:58 Well, I mean, they're sort of challenging this idea that they can't do something that's not in their training data, right? So I think that's not strictly true. They're able to do this kind of ... We call it compositional generalization, right? They're able to compose things that they've seen and make new things and they can do that up to some depth. I think the interesting question is, are they able to produce things that are genuinely novel in the sense of maybe not even being recognizably composed out of things that you've seen before, right? Like come up with a completely new artistic style rather than something which is a blend of styles in a way that's novel or something like that. And I think that's a question where it's still much more of an open question. We in my lab have looked at some questions related to this.

00:42:48 So I have a postdoc, LLU, who's been looking at some related things, looking at creativity in these models. And I can tell you hot off the presses some of our results there, but one of the things that she's found is that if we ask LLMs to come up with sort of novel product ideas, they actually do slightly better than people in terms of producing ideas that people rate as novel. And interestingly, a strategy that works for people, which is asking them to sort of do this by forming an analogy, right? So you say, come up with a new idea for a toaster where it's using principles that come from how octopus suckers work or something. That's a good way of making humans extra creative, but it turns out not to work for large language models. And one reason why that might be the case is that the language models are already operating with such a broad set of concepts that they don't need that extra external stimulation in order to come in.

00:43:49 And so there are ways that it seems like they can do some kinds of creativity. I think the thing that we haven't seen is something where they're sort of like genuinely coming up with something that we can recognize as a sort of breakthrough. One of my favorite thing to do whenever a new model comes out is like ask it to come up with ideas for scientific studies that we could do. And sometimes they get a little bit interesting, but I think it's also worth remembering that because the models are trying to produce things that have high probability, they're trying to produce high probability sequences of tokens, they're going to be pretty conservative in the things that they do. So they're not going to suggest something that's completely wacky because they've not been trained to do that. They've been trained to produce things that are like with high probability based on the things that they've seen before.

00:44:34 And so if you think about an important part of creativity is just like being completely out of the box, that completely out of the box behavior is something which is maybe inconsistent with the way in which they're trained.

Jon Krohn: 00:44:43 Right. We're back to that more higher probability of the answer of 30 as opposed to 29. And it reminds me of a friend of mine who's one of the best selling authors of our time. I won't name him because this is kind of getting into behind the scenes conversation with him, but he was interested in, he is interested, I suppose, in using LLMs to unearth rare information in his field. And so kind of the stuff that's common has already been in his books or other bestselling authors in his field. And so he wants to be able to dig up rare ... And it's very hard to get ... One way of getting there is by just asking for more and more and more and more and more examples. And then eventually maybe some of the examples start to become relatively rare, but it's very hard to get that kind of output.

Tom Griffiths:	00:45:38	That's right. That's what we would expect because it sort of pushes against the way in which they're trained. One trick that we use for actually getting this to work is they, because what they're doing is calculating a conditional probability, right? So they're taking the information that they've seen before, they're conditioning on that information, doing their belief updating about what the things that would come next would be. You can push them into lower probability parts of the space by doing something where you're creating prompts compositionally, right? So it's much less likely to sort of come up with something spontaneously that's an unusual thing. But if you write a little bit of code that sort of procedurally generates prompts where you're putting in an unlikely combination of things, then you say, "Okay, tell me about this. Find an example of this. " And then it will then be able to say, "Oh yes, a good example of this is this.
	00:46:27	" And so you can actually use it to search the space of improbable by things by constructing prompts that are pushing it out into the improbable corners of that space.
Jon Krohn:	00:46:37	That is interesting. Good tip there for a lot of us in ways that we're trying to engineer LLMs to do these kinds of rare outputs that we might want for whatever reason, including for creative reasons. I'm going to get into curiosity in a moment because that's a fascinating topic, but just before we get there, I want to draw the line under something you were talking about maybe 10 minutes ago and we've kind of gone off on other parts of conversation, but something that you talked about, which is something that's so obvious to me now that you've said it out loud,

and I guess in some ways, I've always kind of known this to be true, but I've never perceived it so clearly as when you were talking about it, is this idea that it's kind of silly for us to be so focused, many of us, me, I come from a biological neuroscience background and then that, over the course of my PhD, it became more and more AI focused.

00:47:31 And so myself, I think a lot of people look for biological inspiration, particularly in human mind for how we can be designing AI systems, but because the constraints are so different for an AI system versus a biological system in terms of ... The example you gave perfectly there with deep blue of being able to think about hundreds of thousands of possible next moves in under a second would certainly be possible with today's systems. That's obviously something that a human can't do. And so I guess it probably makes sense. We might as well look to human intelligence or biological intelligence for inspiration in some ways, but we've got to remember that we have completely different kinds of constraints in the system that we're applying it to. And so what ends up working really well, it could be something like scaling. It's mind blowing to me what scaling has been able to achieve, taking this transformer architecture, scaling up with more data, more compute.

00:48:31 When I finished my PhD in 2012, if somebody had asked me if we could have AI systems with the kinds of capabilities that we have today in 2026, I would have said, "That sounds like science fiction to me. I doubt it.

Tom Griffiths: 00:48:45 Yeah. No, I think a lot of people would have agreed with you. It's very easy. This is one of the challenges of the field of AI is that people pretty rapidly adapt to normalizing whatever the last big advance was. So some AI practitioners sort of say, "Oh, the definition of AI is whatever the thing is that we can't do yet." So you keep moving the. Exactly. So I think this sets up an interesting challenge. So you have these two paths, you can keep building systems that are different from us and leaning into the ways in which we can make use of all of the properties of computers, or you can try and build systems that are maybe more like us and sort of think about how it is that we navigate the constraints that we navigate. So I think you could think about this as being a little like as a sort of classic analogy that's made in AI, right?

00:49:46 If you think about like you want to build a plane, then going and looking at how birds fly is useful, but there's going to be some point where you're going to say, "Okay, now we've got the basic principles and now what matters more is building like faster jet engines." And you're not going to learn how to build faster jet engines by looking at how birds fly. And so that's that root of saying, "Okay, we've taken inspiration from biological intelligence. We kind of figured out what the basic principles are. Now let's build this thing." I think the way that I think about that is that one of the most fundamental similarities between what birds are doing and what planes are doing is that they're dealing with the same laws of nature, right? That there's a set of principles of aerodynamics which are shaping what's going on in the bird and are going to shape what's going on in your plane and understanding those principles of aerodynamics is something which is fundamental if you want to build better airplanes.

00:50:40 And so that's the way that I think about what we do as cognitive scientists and how it's useful for thinking about AI is that if we're able to figure out what are the broader principles that are relevant to characterizing intelligence, what are the laws of thought, those laws of thought are the things that we're going to expect to be in parallel between natural intelligence and the artificial intelligences that we're building. And having expressed things in those mathematical terms gives us the opportunity to say, "Oh, here's where an insight can transfer in one direction or the other." But there are different levels at which those insights can transfer. They can be the most abstract thing of what the problems are we have to solve and what's the right way of solving those problems where things like logic and probability theory and base will show up. Or it can be at the level of, oh, how do we actually build a system that produces that intelligent behavior?

00:51:28 And that's where things like artificial neural networks and so on are useful. But I think it gives us the framework that we need in order to then work out, how do we have this conversation turn into a productive one that can go in both directions and how do we recognize when it's not useful to necessarily use those biological principles?

Jon Krohn: 00:51:48 That's a great framework for trying to understand when we should be using biological inspiration in our AI systems related to that, so curiosity. One of, if not the most famous industrialist of our time created his now near the frontier AI lab through this idea of maximizing curiosity. And in a paper from 2020, you model curiosity as rational information seeking. In your book, using Herbert Simon's parable of the Ant on the Beach, you suggest that much of intelligent behaviors complexity comes from the environmental complexity rather than the agent's internal sophistication. So I've talked about a bunch of different things there, but the kind of the common thread is that some people think the curiosity and modeling that in AI systems could be a key driver of getting to big breakthroughs in AI systems. And you've obviously spent a lot of time thinking about curiosity and what it means.

00:52:56 And so I'd love to just have your insights on that.

Tom Griffiths: 00:52:58 Yeah. So my work on curiosity, and this was with my former graduate student, Ratched Dubai, was looking at how we can think about what the computational problem is that curiosity solves, right? And the way that we set that up is essentially what you're doing when you're being curious about things is you're gathering information which is potentially going to increase the value of the actions that you're able to take in the future, right? So this turns out to be a reasonably good model of humor and curiosity is that like our curiosity gets captured by those things where if you think about it, it's not necessarily something that ... Like you're not necessarily just like getting distracted by things that you've never seen before because you constantly are seeing things that you've never seen before, right? And you're not getting distracted by things that you've seen a lot because you've seen those things a lot and they're no longer interesting.

00:53:54 What you're sort of curious about is often things that you've seen a few times and you think, "Oh, maybe I should be paying attention to that. " And you can explain why you have this sort of inverted U-shaped curve of curiosity from this perspective of you're trying to find the things which are going to sort of maximize the value of information for the actions you take in the future because something that you never see, you're also never going to see again. So if you've never seen it before, you probably will never see it again. So it's not really worth learning anything about it because this might just be a one-off. And something that you've seen a lot, you've already had lots of experience with, and so you've already learned how to deal with that thing. But if you see something a few times, then that's an indication that that thing is going to occur again in the future.

00:54:36 And because you've only seen it a few times, it's probably not something you know a lot about. And so for that reason, it's something that should capture your curiosity.

And we show that this actually gives us a pretty good model of how it is that people actually choose what they're going to engage with when they're learning about new things. In the AI setting, there are people who have explored curiosity as a way of defining an objective function for an AI system where traditionally you might think about building an AI system by saying, "Here are some rewards that you can get. " And you get a reward when you solve a problem and you get a reward when you navigate the maze or you get a reward when you get the football and the net or whatever it is that you're trying to do. And there's an alternative way of doing that where instead of making the rewards dependent on things that happen in the external world, you make the rewards dependent on things that happen internally.

00:55:21 So you get a reward for learning something about a thing that you may not have seen, that you've seen a few times. You build in a reward function which aligns with that relationship where it's like you're sort of decreasing the entropy of your beliefs, right? The sort of decreasing the uncertainty that you have about things in proportion to how often you see those things and then it produces that same kind of inverted U shape. And so that's a reasonable way of thinking about creating an AI system that's going to go out and explore on its own because it's going to be driven by trying to identify those pieces of information that are likely for it to be using in the future. And so I think that's the reason why you see that same concept manifest across these different places in psychology and computer science.

Jon Krohn: 00:56:10 Really interesting. I'd love to be in your lab doing the research that you're doing. This is such cool stuff. Getting near to the end of the technical questions that I have for you, I think this is the pen ultimate topic area that I have for you. And so this is around evaluating systems. So this is a bit of a long question, but as you lead the Princeton lab for AI, one of the concerns is to ensure these systems are aligned with human values and that we have robust collaborative systems that understand the nuances of human judgment. You've previously argued that LLMs and humans show the same core interpretability challenges. So opaque training data, inaccessible mechanisms, complex internal structure. And so LLMs and humans alike are internal life is best studied through behavior. I mean, that goes back to right at the beginning of this episode talking about the way that you evaluate intelligence through surveys and that kind of thing in humans.

00:57:12 However, you caution against anthropomorphizing AI systems because they're likely to be alien in various ways, as we've talked about in this episode. So things like different kinds of constraints mean that they're going to behave in different ways than us. And so how can we think about designing our benchmarks for evaluating AI systems so that we have some better sense that we're actually getting at an understanding of what's going on as opposed to pattern matching. So there was a really long question, but you're nodding your head, so it seems like you follow the thread through.

Tom Griffiths: 00:57:51 The first thing I'd say is I'm not a big fan of benchmarks, right? Benchmarks are an engineering device rather than a scientific device, right? So if you're making a benchmark for an AI system, normally what that means is you're saying, "Here's some capacity that I want my AI system to have, and I'm going to make a set of tasks that I can use to measure whether the system has that capacity," then I'll be able to say, "Okay, it got 80% of those tasks right. So we'll say it's scored 80% and that's how well it's doing in that particular capacity." That's a little different from the way that we approach studying LLMs in Mylab, where we're more interested in what are the targeted experiments that we can do that are going to give us information about what's going on inside them, what the kinds of mistakes that they make are, sort of deeper understanding of that system.

00:58:43 And so that's less about making a big set of tasks and seeing what percentage of them they get right. And it's more about designing a set of problems where the particular problems that they get right or wrong are then diagnostic of something which is going on inside that system. And so that's the difference between, like I said, approaching this as an engineer and approaching it as a scientist. Really, we use a lot of the methods of cognitive science and psychology to study these systems where we're conducting experiments that are designed to elucidate what the principles are that are working under the hood.

Jon Krohn: 00:59:16 Right. Well, great. From human cognition to understanding AI systems to benchmarking them, this has been a fascinating episode. Much of what we talked about in today's episode is from your brand new book, The Laws of Thought, which at the time of us recording isn't even out yet, but at the time of publication, this episode's publication, it should be live, it should be available for order in most markets. And based on how well your previous book, Algorithms to Live By Did, you can expect that this is going to be a mega bestseller just like that one was. I don't have insights into exactly how many copies it sold, but it has over 6,000 reviews on Amazon. So you can expect orders of magnitude more in sales to Because of that kind of figure, which is just wild. In our field, for these relatively technical books to be selling like that, I mean, that is really remarkable, Tom.

01:00:10 Congratulations on that. And so I just also wanted to talk a tiny little bit about your book, Algorithms to Live By Here at the end of the episode. From your early days as an avid fencer to co-authoring that book, you have long

		applied computer science to the messiness of real life. Your book Algorithms to Live By Reframes Everyday Human Procedures as computational strategies. So optimal stopping, explore, exploit trade-offs, computational kindness. And growing up, you described fencing as involving interesting computational problems, even trying to break complex moves into sequences. So how did those early problem solving experiences shape the way you now think about cognition and algorithms?
Tom Griffiths:	01:00:53	For me, it's very natural to think about the world in computational terms. So the other formative experience I had growing up was playing a lot of tabletop role playing games. And you can think about what you're doing in a tabletop role playing game is essentially defining a system which is supposed to somehow approximate real life. So it has all of these things that you're putting in there about what the chances are of succeeding at taking various actions and how you get better at doing things and so on. So you can actually think about your role playing game as encapsulating some kind of psychological theory of cognitive progress as you're learning new things and so on. And so I think the thing that we did in Algorithms to Live By was really saying that those computational principles have much broader applicability than people might expect. So I think we had two kinds of readers for that book.
	01:01:49	Part of what accounted for it selling relatively well was that we had people who we originally wrote the book for, which is people who are used to thinking in computational terms and us telling them, "Oh, here's how you can apply this to everyday life." And then we had another audience of readers who kind of understood how everyday life worked, but wanted to learn a little bit more about computer science. And so I think it was really the synthesis of those two kinds of things that really caught people in coming at it from these different perspectives.

The intersection of psychology and computer science that that encapsulates is very much the same space as the current book, but it's sort of coming at it in a different way where it's the same thing of trying to think about how do we formalize something, but you can think about in algorithms to live by a little bit more, we were trying to formalize something about the external world and the new book is about how do we formalize our internal world?

Jon Krohn: 01:02:48 Right. The big question that has been the focus of your research this whole time. Really fascinating. So again, algorithms to live by, of course, the mega bestseller that you can grab. But most of today's interview was getting into these internal processes of both humans and machines. And so if today's conversation was fascinating for you, you'll want to grab Tom Griffith's brand new book, The Laws of Thought, the Quest for a Mathematical Theory of the Mind. Thank you so much for taking all this time with us today. Before I let you go, I ask all of my guests the same last two questions and with a book author like you, this second last question seems almost silly, but do you have any book recommendations for us other than your own books?

Tom Griffiths: 01:03:35 Yeah. So what I would recommend is my co-author on Algorithms of the By Brian Christian wrote a wonderful book called The Alignment Problem, which touches on other things that we mentioned here about how it is that you create AI systems that are aligned with people in terms of the values that they have and the actions that they're going to take. And I think it's a really good introduction to both some of the principles behind modern AI and some of the challenges that we face in deploying it.

Jon Krohn: 01:03:58 Perfect. Thanks for that recommendation. And for folks who love this conversation today, in addition to buying the laws of thought, how can they be following you for your thoughts in real time in social media and that kind of

Tom Griffiths:	01:04:12	Thing? I'm not very good at social media.
Jon Krohn:	01:04:15	Well, then this podcast will be a rare opportunity for folks to get in behind your brain.
Tom Griffiths:	01:04:19	So I have a LinkedIn page and you can find my lab on Twitter and BlueSky. So it's the Computational Cognitive Science Lab. So it's Cocosi_Lab on Twitter and BlueSky.
Jon Krohn:	01:04:34	Nice. We'll have a link to those all in the show notes, of course. Professor Tom Griffith, thank you so much for taking the time out of your no doubt, very busy schedule to record this fascinating episode with us. And yeah, when you release your next book, we'd love to have you on again.
Tom Griffiths:	01:04:50	Thanks for having me. This was a great conversation.
Jon Krohn:	01:04:55	Exceptional episode today with Professor Tom Griffiths. In it, the Princeton professor covered how the laws of thought are mathematical principles for describing our internal mental world, how psychologists tend to view humans as irrational decision makers because they focus on our biases and mistakes, while computer scientists find us impressive because we accomplished remarkable feats with extremely limited computational resources. He talked about how LMs are trained to predict high probability token sequences, which means they're more likely to output 30 than 29, say, in accounting task. How a human child learns language from roughly five years of speech input while LLMs require the equivalent of thousands of years, a gap representing the inductive biases and priors built into human cognition and how just as airplane engineers took inspiration from birds, but

eventually focused on jet engines rather than feathers, AI researchers should understand when biological principles transfer usefully and when machine specific approaches like scaling offer a better path forward.

01:05:56 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 Tom's social media profiles, as well as my own at superdatascience.com/969.Thanks to everyone on the Super Data Science podcast team, our podcast manager, Sonja Brajovic, media editor, Mario Pombo, our partnerships team Natalie Ziajski, our researcher, Serg Masís writer, Dr. Zara Karschay, and our founder Kirill Eremenko.Thanks to all of them for producing another stellar episode for us today for enabling that stellar team to create this free podcast for you. We're deeply grateful to our sponsors. You can support the show by checking out our sponsor's links, which are in the show notes. And if you'd ever like to sponsor the show yourself, you can find out how at johnchrome.com/podcast. Otherwise, help us out by sharing this episode with someone who'd love to learn from it.

01:06:46 Review the show on your favorite podcasting app or on YouTube. Subscribe, obviously, if you're not already a subscriber, but most importantly, just keep on tuning in. I was 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 SuperDataScience Podcast with you very soon.