

**SDS PODCAST
EPISODE 869:
AI SHOULD MAKE
HUMANS WISER
(BUT IT ISN'T),
WITH VARUN
GODBOLE**



- Jon Krohn: 00:00:00 This is episode number 869 with deep learning researcher Varun Godbole. Today's episode is brought to you by the Dell AI Factory with NVIDIA and by ODSC, the Open Data Science Conference.
- 00:00:18 Welcome to the SuperDataScience Podcast, the most listened to podcast in the data science industry. Each week we bring you fun and inspiring people and ideas, exploring the cutting edge of machine learning, AI, and related technologies that are transforming our world for the better. I'm your host, Jon Krohn. Thanks for joining me today. And now let's make the complex simple.
- 00:00:51 Welcome back to the SuperDataScience Podcast. Today I've got a brain stimulating episode for you with a hardcore AI researcher who's recently turned his attention to the future implications of the crazy fast moving exponential moment we find ourselves in. Varun spent the past decade doing deep learning research at Google across pure and applied research projects. For example, he was the co-first author of a Nature paper where a neural network beat expert radiologists at detecting tumors. He also co-authored the Deep Learning Tuning Playbook that has nearly 30,000 stars on GitHub, that's crazy. And he more recently authored the LLM Prompt Tuning Playbook. He's worked on engineering LLM so that they generate code, and most recently spent a few years as a core member of the Gemini team at Google. He holds a degree in computer science as well as a degree in electrical and electronic engineering from the University of Western Australia.
- 00:01:45 Varun mostly keeps today's episode high level, so it should appeal to anyone who, like me, is trying to wrap their head around how vastly different society could be in a few years or decades as a result of abundant intelligence. In today's episode, Varun details how human relationship therapy has helped him master AI prompt

engineering. Why focusing on AI agents so much today might be the wrong approach in what we should focus on instead, how the commoditization of knowledge could make wisdom the key differentiator in tomorrow's economy and why the future may belong to full stack employees rather than traditional specialized ones.

00:02:19 All right, you ready for this mind-altering episode? Let's go.

00:02:29 Varun, welcome to the Super Data Science Podcast. How are you doing today?

Varun Godbole: 00:02:32 Thanks. It's awesome to be here. Thanks for having me, man.

Jon Krohn: 00:02:36 I imagine you're in New York?

Varun Godbole: 00:02:38 I am, I am. It's actually not that bad right now, it's been a bit chilly, but yeah, looking forward to getting a bit warmer.

Jon Krohn: 00:02:47 We know each other from New York. We know each other from the gym actually.

Varun Godbole: 00:02:50 That's right.

Jon Krohn: 00:02:52 And so you also know that I haven't been at the gym now in over a month in New York, it's because I have been in Canada for the past month and people... So I'm near Toronto or in Toronto for the past month, and when I speak to people in New York, they're like, "Oh, it's so cold." I look at the New York weather every day and I'm like, "I would love to have the weather that you have today."

Varun Godbole: 00:03:15 You joke about that. I grew up in Australia, so I didn't really get this kind of cold weather and actually really like it. I didn't grow up with snow. I don't know. I find it

refreshing and kind of [inaudible] jolts you a little bit when you go outside.

- Jon Krohn: 00:03:28 I like that. You never went skiing in Tasmania?
- Varun Godbole: 00:03:31 No, no. Tahoe was the first time when I came to America and I nearly died doing a bunny slope, so.
- Jon Krohn: 00:03:41 How old were you when you first saw snow?
- Varun Godbole: 00:03:47 26.
- Jon Krohn: 00:03:48 Wow. That's cool. I'm sure we have lots of listeners that have never seen snow. We have listeners from all over the world, but that's pretty crazy for me because I was like one-day-old.
- Varun Godbole: 00:04:00 Oh wow.
- Jon Krohn: 00:04:00 When I left the hospital it was March, so I don't know, maybe it was snowing. Don't know. I'm not going to look it up right now, but I guess I could find out in an almanac.
- 00:04:09 Anyway, so as the listeners now know, we know each other from the gym, from a CrossFit gym that we both work out at, but the reason why you're on the show actually has nothing to do with that. And you don't even know this because I didn't tell this to you, but the reason why I put you on my list to get you onto the podcast is because I read about you in my friend Natalie Monbiot's blog.
- Varun Godbole: 00:04:32 Oh, interesting. I actually didn't know this.
- Jon Krohn: 00:04:35 Yeah, I know you didn't know that because I haven't told you. That is private information to me. Well, now, lots of people around the world. But yeah, at the end of January, Natalie Monbiot, who was my guest on this show on

episode number 823, amazing episode. She's an incredible speaker, and she talked about the virtual human economy, which was a really cool episode is about how virtual versions of you could generate an income and play a meaningful role in society and how that's already possible today, but it might be more and more common in the future. In fact, I don't know why I said might be, it will be more and more common in the future.

00:05:15 And so at the end of January, Natalie wrote a post. She writes a weekly email, and she has an Oxford degree in language and literature, so she writes pretty good blog posts and sends them out in her email newsletter and end of January, it had to do with you. In fact, I would venture to say that you're the inspiration for the week's whole post because the post is about instead of obsessing over AI agents, it seems everyone is in 2025, it says build systems that make us wise.

Varun Godbole: 00:05:46 Yeah, yeah. I mean, I think, yeah, that's right. I've chatted a bit with her about this, and I think a big influence for a lot of this was Professor John Vervaeke actually at the University of Toronto, and he's been doing a lot of work on the cognitive basis of wisdom, relevance, salience, and he has a bunch of amazing lecture series online that I think do an amazing job at providing synoptic integration of all the different aspects of cognition. And that's heavily informed my thinking, certainly around AI and LLMs over the last few years.

00:06:26 And yeah, I chatted to Natalie about this because we know each other through another mutual friend, and we talked a bit about this and a lot of people are really focusing on building agents, but I think what interests me personally is how we can use these systems to increase our personal agency in the world and our wisdom in the world. And you could argue that building an agent does give you more agency, but I think the framing matters in

the details of how product surfaces are constructed and the details of how you frame the problem.

00:07:07 And so what really excites me is how can I as a person be much more agentic in doing pro-social things and being aligned to what is true, good, beautiful, and doing good in the world, being a better person basically, being aligned to my aspirations, whatever they might be, whether it's I want to get fitter in the gym, or I want to have less anxiety, or I want to, whatever my instrumental goals are, how I can use these systems to do that and cultivate my own agency.

Jon Krohn: 00:07:38 It's beautiful. And part of why I wanted to have you on the podcast was I think it is so spot on, it's reframed my whole thinking about everything that I'm doing with AI. And even what I want to be talking about in keynotes or on the podcast, I think that it's such a great perspective and I'm going to evangelize it.

Varun Godbole: 00:08:00 Awesome, awesome. I genuinely believe that a lot of the world's problems, or at least certainly... Maybe I won't project this to everyone in the world, I'll talk about myself, a lot of my problems certainly come from not enough metacognition as opposed to too much of it. I think a lot of my own problems in my own life come from my life being insufficiently examined as opposed to just a certain sort of being too conscious.

00:08:32 And I think that the other reason I'm really excited is, or what I've seen in general in my own life, is that as you increase your own agency at doing things, I've found that I rapidly reach the limits of what that behavior entails. And I think it is in engaging with those limits that I feel like I've grown the most as a person in re-evaluating like what my priorities are, what I actually care about, my conceptualization of what my aspirational self should be like.

00:09:05 And so I think there's something really powerful there about the possibility of engaging in this reciprocal loop where we become more agentic as people, there are systems that afford that cultivation. And then as we do more stuff in the world and we use these systems, they get better through our interaction with it, which then mutually reinforces that. And I think there's something really powerful as we become more agentic, we reach the limits, we find our limiting beliefs, we find the limits of our imagination. And in recognizing and compassionately engaging with those limiting beliefs and those limits and those limiting behaviors, I think that we grow as people. And I think, I don't know, the best things in my life have come from that sort of human growth, and that's what I'm just really excited about these days.

Jon Krohn: 00:10:00 I think it's a great mission. I think it's feasible, and I think it is what all of us could be working towards in small ways or large ways as an individual as well, as you say, in the products that we build, in the AI systems that we build.

00:10:15 As listeners to this podcast, we're going to come back to that topic later in the episode, so this kind of wise AI stuff, "Hey, you're talking wise," that's what we mean, we mean AI that can really do a good '20s gangster impression. This idea of a wise AI system, we're going to talk about that more later on in the episode and some of your writing related to that.

00:10:42 I want to start with writing that you did five years ago. So five years ago you had a Nature paper, you were one of the first authors, you were co-first author on this Nature paper, and probably a lot of listeners know, but maybe not a hundred percent of them, that Nature is one of the most prestigious peer-reviewed academic journals that you could be published in. As you mentioned to me before we started recording, it has a pretty innocuous title, this

article. The title is International Evaluation of an AI System for Breast Cancer Screening. But that kind of belies why this is really interesting, doesn't it, Varun?

- Varun Godbole: 00:11:21 Yeah, there's a couple of things that's really interesting about this. So the team at the time was really interested in how we can use machine learning for various types of medical imagery. And this specific piece of work is for mammography imaging, which it's relevant to a lot of people's lives every year.
- 00:11:45 And what we demonstrated in this paper is that you can use a deep learning system to predict biopsy-confirmed breast cancers two or three years depending on the country within which you're the prediction in because different countries have different screening guidelines and so forth, you can predict those biopsy-confirmed cancers a few years in advance, and you can do this better than expert human radiologists. And so we benchmarked this on a very large retrospective screening data set sourced from the US and the UK, totaling like a few hundred thousand patients or something, like a hundred and something thousand patients, details in the paper.
- 00:12:31 And we also did a separate reader study with a expert radiologists who had nothing to do with the data collected, and the model will beat all of them as well. So the team was really excited for this to get published. And I don't know, I think AI has a lot of potential for medicine and making it more accessible and making it much more reproducible. So yeah, that was a really exciting thing to get the opportunity to work on.
- Jon Krohn: 00:13:02 Yeah. There's some other reasonably well-known names, as authors on this paper, Demis Hassabis, Mustafa Suleyman, co-founders of DeepMind. Pretty cool, man. And yeah, this is a good example of the kind of... It's an

interesting situation actually. We can tie this into the wise AI idea a little bit because here in this 2020 paper you're describing a system, an AI system that can replace humans on a task. But this isn't just about replacement, this is about augmentation and complementation, right?

- Varun Godbole: 00:13:40 Yeah, that's right. Because you'll find that, I mean, this wasn't really in the paper, but you can generally find that humans and AIs, they have different strengths and weaknesses. I think it's an area of active research.
- 00:13:54 Frankly, I'm a bit out of the loop on the literature on breast cancer screening because I've been working on different things in the last few years. But there is a world where the synthesis can be incredibly powerful. And I think that's what excites... That's really what excites me, and just using that. And there are countries where one way to look at it is the ratio of patients per mammographer in that country. Again, it's been years, so I don't have these numbers on me, but you could imagine that there are some countries where that ratio is pretty small. And then there are other countries where that ratio is actually... There are many, many patients per mammographer or radiologist.
- 00:14:40 And so I think there's just something really cool about the possibility of using technology to substantially improve access to what could potentially be life-saving screening for lots of people.
- Jon Krohn: 00:14:50 This episode of SuperDataScience is brought to you by the Dell AI Factory with NVIDIA, delivering a comprehensive portfolio of AI technologies, validated and turnkey solutions, with expert services to help you achieve AI outcomes faster. Extend your enterprise with AI and GenAI at scale, powered by the broad Dell portfolio of AI infrastructure and services with NVIDIA industry-leading accelerated computing, a full stack that



includes GPUs and networking; as well as NVIDIA AI Enterprise software, NVIDIA Inference Microservices, models and agent blueprints. Visit www.Dell.com/superdatascience to learn more. That's Dell.com/superdatascience.

00:15:36 For sure. So let's talk about some of the stuff that you have been working on since then. One of those things is the Tuning Playbook. So you describe yourself as passionate about having more systemic neural network development. So neural networks, so the kind of AI technology that would be used to facilitate breast cancer screening, that also facilitates all the kind of generative AI capabilities that we have today across text generation image video, all that stuff happens with artificial neural networks. And you have this Tuning Playbook that you released as part of a team at Google. And so tell us about that.

Varun Godbole: 00:16:17 Yeah, so a lot of the motivation behind this playbook came after the work on this mammography paper actually. At the time, and it's still kind of true today, training neural networks can be a very ad hoc, some might uncharitably call it alchemical, and it's kind of true. But it involves a lot of experimentation, a lot of empiricism and a lot of research to train and deploy a model.

00:16:47 And so something I was really interested and excited about is, well, at that point in time, I just trained a lot of models. I knew a lot of people that trained a lot of models, and it was like, how can we systematize this process? The broad research agenda that we were interested in is kind of like, you could imagine the transition from alchemy to chemistry or something like that. You could imagine... Systematization can be very, very helpful for engineering. And so frankly, on that paper, even though I'm the first author, the other authors know way more than me, and

everything I learned from that comes from them. And really, we kind of just got a bunch of our heads together and tried to write down what's worked, what hasn't worked, and we collectively have decades of experience training these models. And we wanted to provide a systematic approach for thinking about hyperparameter tuning architecture, just various aspects of model selection.

- 00:17:57 And it's true, this playbook was kind of released, I believe, before ChatGPT came out. But I think that a lot of the things described in that playbook are still very true today because the intent of that playbook was to be a sort of fundamental look at how you should think about running hyperparameter sweeps, what sort of plots you should make, how you can be more systemically empirical with questions. I have this compute budget, these are the constraints of my problem, therefore, how can I systematically go through a bunch of steps and reliably reach a good outcome, and then what process should I have to do this over and over again.
- 00:18:38 And that's kind of what the whole playbook is about. And so it got popular at the time on the internet and I was pretty excited about it. And we released it as a Markdown file. So at the time, the standard way of releasing papers or ML artifacts like this was a PDF and archive, but we really wanted to release this as a Markdown file with I think Creative Commons license or whatever the permissive license is, because we really wanted the community to be able to easily fork it, modify it, come up with our own best practices and give us pull requests back, whatever, for it to be a sort of collaborative thing.
- 00:19:18 I think we weren't exactly clear, I don't want to overstate it, but it is cool that what ended up happening is that a bunch of folks decided to fork it, and I believe crowdsource translations in a bunch of different

languages, which are not endorsed by us because I can only speak English, but that was pretty cool and I think it's still pretty relevant today for people training models.

- Jon Krohn: 00:19:43 For sure. I think it's an invaluable resource, and I'm not the only one. It has 28,000 stars at the time of recording, which is insane. That's amongst the most stars I've ever seen on a project. So yeah, hugely impactful. Some amazing contributors on there. And so yeah, thanks to you and the Google Brain team as well as someone from Harvard University. Christopher Shallue.
- Varun Godbole: 00:20:08 Yeah, he actually used to be a Brain before he went to Harvard. Yeah, he's cool. Like I said, even though I'm the first author, the other authors, I really shout out to George Dahl, Justin Gilmer, Zach Nado, Chris Shallue, they're the real brains behind the outfit, and I was kind of just learning from them and kind of getting everything going. And yeah, I've learned... It was a lot of fun working with them and I'm grateful we were able to get that out there.
- Jon Krohn: 00:20:41 I teach an intro to deep learning course. I've been doing it for coming on 10 years now, and five years ago roughly, six years ago, it was published as a book, the curriculum that I developed as an introductory deep learning class. And something that every class always asks, once I explained that we can add some more layers, we can double the number of neurons in a layer or all of the layers. It was like, "Okay, but why? Why are you making those decisions?" And up until now, I basically always just said, "Well, you can either just experiment and find out empirically by experimenting with a bunch of parameters, or you can do some kind of search." The simplest thing is doing a grid search, so just setting up some parameters to search over, but there's also clever Bayesian approaches to homing in on what the ideal parameters could be.

- Varun Godbole: 00:21:36 Yeah. So this playbook is about that question pretty much. It tries to take a much more general approach. So it's kind of architecture-agnostic in the sense that it won't tell you this is when you should add a new layer versus this is when you should change the width of the layer. But it is about helping practitioners grapple with the question, here are the experiments I have now, what is the experiment I should run next? Because the assumption is that if you can set up the base case and a good recurrence relation, you can iterate your way to success.
- 00:22:15 And so there's a lot of thinking in the playbook about how should you think about setting up the right initial state for your experimentation and how should you think about, given the data that I have collected, what is the next experiment you should do? And I should emphasize this is meant to be a living document. That's also why it's the Markdown file on GitHub, we reserve the right to change our opinions and feedback is very welcomed and encouraged. And it's not the final answer, or... I won't speak for the authors, I won't pretend to be the arbiter of how everyone should tune their models, but it's just like we've been training models for a while, these are our 2 cents of how one could think about doing it. And so it's not really... That's the kind of vibe, right? And hopefully it helps people, if it doesn't, please click create issue or something and give us feedback. Yeah.
- Jon Krohn: 00:23:16 Nice. Well, thank you for this resource, Varun and everyone that you took information from to create this invaluable resource for all of us, Varun, it's brilliant. So after this interest in more systematic neural network development, or actually, I don't know if it's after, it could have been in parallel, you were also working on code generation research and I think this was also pre-ChatGPT, right?

- Varun Godbole: 00:23:44 Yeah, yeah, that's right. Yeah, it was actually... Was it after? It was sort of parallel slash after. It was like this playbook was a transition between me working on medical imaging and working on LLMs. Yeah, this was before ChatGPT.
- Jon Krohn: 00:24:01 It was before, because the Google research blog post about the impact of this code generation tool.
- Varun Godbole: 00:24:09 Yeah, we've been working for a while on that project, so yeah, that's right. That's right. That's right. It was before ChatGPT. That's right. It's coming back to me now, yeah.
- Jon Krohn: 00:24:19 I'll have this link for people in the show notes. There's a blog post from July 2022, which is three months, four months before ChatGPT's release, and the title of this Google Research blog post is ML Enhanced Code Completion Improves Developer Productivity, so definitely ahead of it.
- 00:24:39 And I mean, it's unsurprising now to hear this because probably all of our listeners, any of our listeners that are writing code, you've got to be using tools to help you. GitHub Copilot, Claude, Google Gemini, ChatGPT, there so many great tools out there for getting help with code completion, they're invaluable tools. But just two and a half years ago in July 2022, that wouldn't have been obvious. It might've seemed like a distraction or something that would have so many errors that would actually take you more time to wade through the mistakes that the code completion was making.
- Varun Godbole: 00:25:18 Yeah, I mean, at that point in time, I already knew... I mean, sequence to sequence modeling had been pretty big for a while. I guess it was first used in the early use cases with translation. But I don't know, this seemed like a really cool thing. The early results were really promising. A bunch of folks that I respected on the Brain

team were really into this, and it just seemed like a genuinely cool use case.

00:25:42 And I know, I don't know, it just made sense to me that I think something I'd seen before when we were working on the medical imaging stuff, or even at this point in time, I'd been doing deep learning research for a while and it was like, yeah, it was early, but it was also quite fascinating how quickly the models were getting better at just a lot of things. And every year the hardware got better, the size of the models got better, and once there were signs of life that, oh, wow, you can use this for engineering productivity, that just, I don't know, it just made a lot of sense to me. And yeah, it was a cool project and I'm really grateful that I got to work on this. And the team's really cool and I learned a lot from a lot of them about a lot of this stuff.

Jon Krohn: 00:26:40 So code completion, obviously we know that's hugely valuable today as a part of these general purpose LLMs that are out there, I already mentioned some of them Claude, ChatGPT, DeepSeek algorithms are making a lot of a splash right now at the time of recording this episode. And another valuable, and I've been using it a lot, code completion tool is Google Gemini. So Google Gemini is a great LLM. And you were on the Google Gemini team. You were working on co-generation as a part of that team I think as well, right?

Varun Godbole: 00:27:17 That's right, that's right. That's right. I was part of Gemini from the start. I was a member, a co-member of that team, and I was on that team right until recently when I left Google. And it was certainly an exciting time and it was cool. It was a lot of fun.

Jon Krohn: 00:27:41 It must be amazing to, I mean, you don't need to go into actually any detail yourself, but I can imagine that it

would be amazing to be working on the team that would be right at the cutting edge of what AI systems can do.

00:27:55 So at the time of recording, for example, I'm looking at the LM Arena, which allows people to rank outputs and tied for first, statistically tied for first at the time of recording as the best general purpose LLM overall, aggregating across a ton of different metrics is Google Gemini 2.0 Flash, Google Gemini 2.0 Pro, ChatGPT-4o, DeepSeek-R1. And if you're willing to go beyond just those 95% confidence intervals, the statistical evaluations that give us that four-way tied first place in first place overall would be Gemini 2.0 Flash. So it's pretty cool to see that. I mean, you don't need to go into any more detail, but I'd imagine I'd be very proud if I was working on something like that.

Varun Godbole: 00:28:53 I mean, I was just one person in a large team. It was a lot of fun, it was really cool experience and the people there are really cool.

Jon Krohn: 00:29:04 And so I also know we can't go into very much detail on specifics, but I was curious to know because I was talking to you prior to recording about particular PyTorch libraries like PyTorch Lightning, which is something that I regularly use, and it seemed like you hadn't even heard of PyTorch Lightning. And I was like, "What? How can Varun not know PyTorch Lightning?" And then I was like, oh yeah, of course, because PyTorch is a Meta product and so probably Google doesn't use it. And I was like, "Oh, are you guys all using TensorFlow?" And you were able to inform me that isn't actually mostly what...

Varun Godbole: 00:29:39 Yeah. A lot of people use JAX. And there's actually Jacob Austin recently just posted a really great guide or tutorial on scaling models on TPUs and JAX and TPUs are used heavily inside of Google for modeling. Yeah. So yeah, I've been at Google for a while, and so I'm ignorant of what's

happened outside, because you're in the day-to-day kind of sprint mode of doing stuff.

00:30:11 And I think especially in the last few years, it's been, at least for me personally, kind of overwhelming working on machine learning because it's like every few weeks, some insane new announcement happens somewhere in the world about how someone has done some amazing new thing. And honestly, five or 10 years ago, it was a much simpler time in machine learning. Now it's just seemingly every week something is happening. And so, yeah.

Jon Krohn: 00:30:46 Excited to announce, my friends, that the 10th annual ODSC East (Open Data Science Conference East), the one conference you don't want to miss in 2025, is returning to Boston from May 13-15! And I'll be there leading a hands-on workshop on Agentic AI! ODSC East is three days packed with hands-on sessions, and deep dives into cutting-edge AI topics, all taught by world-class AI experts. Plus, there will be many great networking opportunities. No matter your skill level, ODSC East will help you gain the AI expertise to take your career to the next level. Don't miss out — the Early bird discount ends soon! Learn more at odsc.com/boston.

00:31:31 It is pretty wild. It is a spin. When I first started hosting the show four and a half years ago, I didn't always every week have some breaking story that I felt like I need to talk about. So sometimes regular listeners will know that I do two episodes a week, there's a Tuesday episode, a Friday episode. The Tuesday episodes always have guests, they're longer, they tend to be about an hour long, and on Fridays there's a lot more flexibility.

00:31:58 When I took over as host from Kirill Eremenko who had been hosting the show for the first four years, he called Fridays Five-Minute Fridays, and they were these short solo episodes and he would talk about life philosophy

advice and a lot of those episodes. And so I kind of carried on that tradition, but part why I was carrying on that tradition and kind of talking about habits you might like to develop or how to stay on top of your habits was because there wasn't always something new for me to talk about in data science and machine learning.

00:32:27 And now that's never the case. There's always so much that I could be talking about in Friday episodes, and I think it's been a couple of years now since... Sometimes I'll have guests that are tangentially related on topics. We had an economist on a few months ago, Nat Ware, to talk about why people aren't happy all the time despite it being such a great time to be alive compared to history. And so I try to have interesting episodes like that sometimes with guests.

00:32:54 But when it's me solo, I'm pretty much always doing research on a particular data science topic. That means usually today a machine learning or AI topic, I'm not doing that many episodes on a new data visualization technique.

Varun Godbole: 00:33:07 Right, right. It's been insane. I think it points to... I think there is genuinely this Cambrian explosion happening underneath us. And I do think that LLMs, which machine... Deep learning in general, but LLMs are just one of the most profound things that I've seen happen in computing, certainly in my short career. But when I look at history or I'm just interested in the history of computing, it's just insane what we've seen happening.

00:33:41 And it is just really fascinating how the way LLMs have even transformed machine learning itself, where one way you can view LLMs, or one way that I view LLMs is that there are sort of arbitrage of ML talent where it's sort of in the past if you wanted to be an ML engineer, you needed to learn PyTorch or JAX or whatever, you needed to learn

how to collect all this data. If the goal you were trying to solve was to put intelligence, quote, unquote, "into your product." And now seemingly what's happened is that you kind of need to be able to articulate the behavior you want in clear English in a prompt and have good evals to measure and characterize the performance of that inference call.

- 00:34:28 And it's sort of fascinating how it's turned so much of what used to be bread and butter machine learning into just prompting. And it's just fascinating how different the skill sets required seem to be for prompting effectively from the skill sets required for training models effectively. And so that's something I've been interested in for a very long time is what does it mean to be in a world where you can take inference for granted kind of thing. You can build systems on top of plentiful inference calls and what does it mean to be able to systematically tune your prompts and systematically manage your prompts as the models get better?
- 00:35:21 And that was actually one of the last things I published before I left Google, which was a playbook for prompting. And the playbook actually has two parts to it. The first part is sort of my high-level thinking, my high-level mental model of how you can think about pre-training versus post-training and how you can think about prompting. And the second half is clear prescriptions on this is what a good prompt looks like and this is what a bad prompt looks like.
- 00:35:52 But I think as I mentioned in the playbook, I think the first half is actually much more interesting and future-proof than the second half because the second half is written from where the models are today. Like prompting o1, or Flash Thinking, feels at least to me, qualitatively very different than prompting 4o or just Gemini 1.5 or something like that. But at the same time,

but that difference, I feel, is in the surface level form of the prompt itself as opposed to the mental models I have of what is actually happening under the hood and how that prompt is getting translated into computation. And obviously that's a little inscrutable because it's a model under the hood, but that's kind of what the first half of that tuning playbook tries to grapple with is how you can kind of think about that.

00:36:51 And the way I kind of approach that is by explaining to people that, and this is where it kind of gets a little philosophical, because it ties into some of the things I've been doing recently, which is LLMs are sort of anything machines in that they're general sequence to sequence machines. And nevertheless, when you ask an LLM to give you an answer, if you ask it how old is Jon or how old is some celebrity, it'll give you an answer.

00:37:29 But the thing is, this thing is not an embodied thing in the world. It's getting its factuality from the data sets that it's been trained on and so forth. And so that's kind of what that playbook tries to explain is that one way you can think about it is that there is no such thing for the purposes of an LLM, often objective fact. So for example, for the Lord of the Rings fans out there say, I come up with a proposition, Aragorn is the king of Gondor, right? Well, actually, Jon, are you a fan of Lord of the Rings? I don't know if...

Jon Krohn: 00:38:09 I have watched all of the Lord of the Rings films many, many years ago.

Varun Godbole: 00:38:15 Okay, okay.

Jon Krohn: 00:38:16 And I read The Hobbit, I didn't like reading The Hobbit.

Varun Godbole: 00:38:20 That's fine.

Jon Krohn:	00:38:21	I found it dull and linear.
Varun Godbole:	00:38:23	Okay, so let me ask you this question then. Suppose I come up with the proposition Bilbo Baggins is a hobbit, right? That's my statement. Is that true or false?
Jon Krohn:	00:38:34	I would say it's true.
Varun Godbole:	00:38:36	Okay, show me, where is Bilbo Baggins? Show me where Bilbo Baggins is.
Jon Krohn:	00:38:43	He is in the Shire.
Varun Godbole:	00:38:45	Shire isn't real, man. Where is Bilbo... How could it be true if the Shire isn't real?
Jon Krohn:	00:38:51	J.R. Tolkien is the God of determining what is real in the Shire.
Varun Godbole:	00:38:59	Right, right. So in the world of the Shire, it's true, but on earth it's not true unless we're saying fiction is somehow real. And the point there is that every proposition you make has a certain backdrop of assumptions behind it. And so that's why in the playbook, I called it cinematic universe, because people are used to that idea now, right? Because different cinematic universes, there's the DC, there's the Marvel, whatever.
	00:39:32	But the point is that Bilbo... The truth value of Bilbo Baggins is a hobbit is contingent on which cinematic universe you're in. If you're in the DC cinematic universe, that proposition maybe it doesn't even make sense. It doesn't even have a resolvable truth value. In the cinematic universe of the Hobbit, it does. And in the cinematic universe of the earth it doesn't. So now what happens when you take a sequence to sequence model, and one way you can think about it, or one way I think about it is that the internet or all of written text corpora

or something like that are an approximation of the set union of all cinematic universes in existence.

- 00:40:19 And the issue is that if you train on all of these things and then you do next token prediction on a statistical model based on that, there'll be many modes in that distribution. But the mode you care about, that is the cinematic universe that you care about, might not be the mode that the model has fitted to. And so we've just described pre-training, and so that's the problem with pre-training. And so what you need to do now during post-training is you now need to shape that distribution to the one you actually want with a specific cinematic universe where there are always two participants, one is the AI assistant, one is the human. The model does not have access to the internet. There are some assumptions about the model's interaction with the world. There are some assumptions on who the user is. There are some implicit and passive assumptions about what is or isn't a fact.
- 00:41:21 And so that's why I think that if you give a pre-trained model or any model like an under-specified prefix, like, "Hello, how are you," unless you've post-trained it properly, you just get a nonsense generation because it doesn't know which cinematic universe you're in. Or if you ask it how old is insert celebrity name, it doesn't know if you're talking about the real celebrity, are you talking about some fictitious universe? Are you talking about when that proposition was true X years ago? Are you talking about now. There's actually a lot of cognitive scaffolding necessary for all of this humanity to work.
- 00:42:02 And so the playbook, the first half of the playbook is... It's an attempt, I don't know if I succeeded to kind of explain this a little bit more. And then that informs the way you think about prompting, because then the implication of that is that one way to think about prompting is, well, for

post-training, the data for post-training comes from human raters, right, that are creating demonstrations for... So one way to think about what these post-trained LLMs are is that they are statistical models, role-playing as human raters that are role-playing as AI assistants. Does that make sense?

Jon Krohn: 00:42:52 Yeah, that does make sense.

Varun Godbole: 00:42:55 And so when you take that a step further, it's like, okay, how do you prompt a model effectively? And this is where people I've seen on the internet also talk about, it's all about putting the right context in the model, and I think that's true in some reductive fashion.

00:43:13 But really what we're talking about here, or one way to think about it is say you write down a prompt and imagine if you just picked off a person off the street and now this person has access to the sum of all of human knowledge, but it has no idea who you are, it has nothing... It's like imagine if when you put a prompt, there is a person picked off of the street who has access to the sum of all human knowledge, and every time a prompt comes in, they're going to read the prompt and they're going to read the data you put in the prompt and they're going to decide what the response is.

00:43:51 So when you frame it that way, obviously that's not happening under the hood, but when you frame it that way, what would you actually put in the prompt? How would you communicate to such a person effectively? And I found that when people... There's this funny, you know the rubber ducky effect where it's like before you come to me for advice, try to explaining your problem into a rubber duck. And then what I found is when people complain that, "Oh, my prompts aren't working," and they're like rattle off all the ways it's not working and all the ways that the model doesn't understand what they

want. And when I tell them, "Have you tried just taking these complaints and putting that in the prompt?" And usually the prompt does way better.

00:44:43 Because another way to think about it is, I linked to it in the playbook, is there's a Wikipedia page about this. I believe sociologists or something call the difference between high context and low context cultures, like forms of communication that are much more explicit versus forms of communication that are much more implicit. I've found that the most effective way of prompting these models is to be extraordinarily explicit with them.

00:45:11 And another book I really recommend in that playbook is Nonviolent Communication, which I learned of in therapy in couples therapy. I didn't do couples therapy to get better at prompting models, but needless to say, for all the viewers out there, but it turns out that effective interpersonal communication and being explicit with what you want in terms of observable behavior from another person is actually not that dissimilar from being more effective at prompting these models.

00:45:45 And so anyway, all that's in the playbook and that sort of shape a lot of... I've been thinking about just cognition for a while, and then that's sort of shaped the things I've been thinking about for the last few months mostly.

Jon Krohn: 00:46:01 That's amazing. That was really funny. I love that the Google software engineer goes to couples therapy and what comes out of it is he becomes better at prompt engineering. It really works.

Varun Godbole: 00:46:17 Well, yeah, and I think this is also why I think I've seen... I've found anecdotally that the best people at prompting know the least computer science or something like that. And I've actually found this inverse correlation between the most knowledgeable machine learning engineers and

their ability to prompt. This is all anecdotal, this isn't science, this is just vibes. But I think there's a truth to that because it turns out, it seems to me that being able to prompt effectively, the skills required are very similar to the skills required for effective interpersonal communication in a context where the two participants don't have a lot of assumptions about what the other knows. And that is challenging.

- Jon Krohn: 00:47:12 Yeah, yeah, yeah. Great insights there. We've now talked a lot about what you were doing in your years at Google on the Google Brain research team, working on Gemini. Very cool stuff. Most recently this LLM prompt tuning playbook.
- 00:47:30 But a couple of weeks ago or a month or so ago, at the time of recording, you have effectively retired, for lack of a better word, you are now... You are keeping yourself busy. And one of those things that you're keeping yourself busy with in retirement is writing. And so going back to the top of this episode and how Natalie Monbiot's post about you and your work brought me to asking you to be on this episode, the specific post that she referred to was called From Knowledge to Wisdom: Value Creation in the Age of LLMs. And this is all about wise AI, the thing that we started talking about right at the top of this episode.
- 00:48:11 So tell us generally about what wise AI is your vision here? I mean, you already talked about it at the beginning of the episode with respect to this idea of increasing human agency as opposed to just being focused on agentic AI systems. So using LLMs, using AI to enhance human agency and to make us the wisest best versions of ourselves, helping us with our metacognition. So actually maybe in the outset of today's episode, you already kind of talked about that enough. There may be other things that have occurred to you that you'd like to add now.

00:48:46 But one of the things that I would definitely like to get to is that there are some complexities that arise, there are some challenges that arise from LLMs, and there's a lot of labs that are explicitly chasing artificial general intelligence, which is this idea of a single algorithm that can do any of the learning that a human could do.

00:49:13 But in many ways it would be different than human intelligence, in some ways more powerful than human intelligence because of the huge breadth. So you were telling me, I think it was prior to us starting recording about T-shirt sizes and where you have the width of the T-shirt and the height of the T-shirt. Maybe you can kind of go into that analogy and explain that. But basically, these LLMs, these AI systems as we approach AGI as more and more tasks can be done by LLMs, the marginal cost of creating so many things, code, art, it trends towards zero. And that presents really interesting implications for humans for what our purpose is, how we value ourselves, what we should be doing.

Varun Godbole: 00:50:05 Yeah, there's a lot to unpack. So let me zoom out a bit. So first of all, I think I am not aware of a single consensus definition of what AGI is. It feels like it's one of those things where the goalposts keep shifting on what it truly means.

00:50:31 And so zooming out when I wrote that post, it was just a stream of consciousness. I'm frankly still trying to wrap my head around a lot of these ideas. And it's sort of like that post was sort of open source blogging or I'm blogging in real time about what I'm thinking. And the way I wanted to approach it is rather than predicting where things are going to go or what is the sci-fi future with AGI look like, just gaining more clarity on where are we now? What is happening right now, looking from a first principles basis on what the relevant curves are right now, and then extrapolating them just a little bit and

then seeing, okay, what are the implications of that if you do that? And then extrapolating them just a little bit more and seeing the implications of that and so forth.

- 00:51:30 And especially because there's just so much hype and snake oil about what's going on with AI. And so I wanted to start with the actual numbers and curves. And so the first thing that occurred to me is that people have said this a lot over the internet is that with every year, two things are simultaneously happening. The models are becoming a lot better on all the benchmarks, like Claude today versus last year, Gemini, ChatGPT, these models where they are at the time of recording versus where they were 12 months ago, way better, just generally way better.
- 00:52:12 But at the same time for that same functionality, the cost of inference has just dropped by orders of magnitude. Context windows have substantially increased. In some sense, latency has improved because we're seeing smaller models that can do the same thing the previous models could do. And all of this frankly mirrors what we've seen elsewhere in computing, is that first the big expensive thing comes out and the cheaper things come out.
- 00:52:41 And it's not just computing, it's like the Tesla Roadster comes out before the Model S and before the Model 3. So that's where I started. And then one of the things as you pointed out at the start of the episode, is one of the things that people have been using these models for is for productivity enhancements for software engineers like Cursor, Lovable, all these startups, Replit, Copilot in GitHub. And so if you plot that curve, it is pretty astonishing the way that curve is moving. And then if you just extrapolate it, you've already started... I've already started seeing this phenomena where people that weren't traditional software engineers or wouldn't have identified as software engineers, the act of producing software is just much more approachable to them now because these

tools exist not just in the act of generating the code itself, but understanding what a piece of code is doing, understanding some obscure thing. Just the practice of generating and delivering software is becoming much more accessible.

- 00:53:53 So that's the curve. So then I started thinking, okay, well, there seems to be a feedback loop driving these two things, and that's a separate thing that we could potentially talk about. But the important thing is that that curve there seems to... There's seems to be a lot of momentum and force behind that curve. And so if you start extrapolating that curve to its logical conclusion, what do you get? What does that look like? And two things started falling out.
- 00:54:28 The first one is that just as how you had the emergence of full stack engineers after the open source LAMP stack came out a few years ago, my suspicion is that what you'll start to have is the emergence of what I'm calling full stack employees or something like that where for example, imagine you were... Imagine, go back to a pre-LLM world where you had some UX designers, a PM and a bunch of engineers on a team, and let's say the PM wants to think about shipping a new feature. Well, in the old world, they would've had to write up a PRD, get some UXer to make a mock, get the engineering manager to agree and prioritize... It would've been this huge thing to align the team towards that new feature.
- 00:55:21 The world we're rapidly approaching is that the PM, assuming that team is using AI assisted tools, can just create a prototype that maybe you wouldn't actually ship, but they can just go to the eng manager, or they could go to the team with generated mocks in a generated prototype showing the interaction, showing what it could look like, and maybe already doing some user research with the prototype and substantially de-risking it before

an "engineer," quote, unquote, ever actually spends a single second thinking about whether it should actually be prioritized or not. That's just a very near term thing. And so what I'm describing there is a sort of fluidity to the way roles are likely going to work.

- 00:56:08 And so if you keep extrapolating from that curve, if you assume that the models are going to continue getting better and cheaper, then what you start seeing is it's not clear to me what the organization of the future looks like. Right now many organizations are organized kind of functionally in the sense of there are the PMs in one org, the UXers in one org, the engineers in one org. Maybe there are different permutations of how... But that role of a PM or UX and eng is reasonably well-defined in terms of consensus in the industry.
- 00:56:47 But if you now reach this world where these AI-enabled tools can just accelerate you in these unforeseen ways, that injects a lot of nebulousity or fluidity in how these teams are organized. And so the model I've been thinking of recently is that's what I mean by full stack employee, where they're thinking not just within their narrow parochial domain, they're thinking across the stack end to end in how that PM or that unit of people delivers value in the broader organization, right?
- 00:57:23 Because one way you can view these roles is they're a bundle of skills. And it's like, I don't know who said it, it's like there's always either bundling or unbundling is how you create value. Maybe what we're seeing is the beginnings of an unbundling of the aggregated set of skills that you see inside of a team so that they can be dynamically re-bundled in different ways depending on the idiosyncratic needs of that organization or team. And that's pretty crazy when you think about it.

- 00:57:59 And then at the same time, because these models are just generally available, what I've also started wondering, I'm noticing is that not just is the cost of delivering software for fixed complexity going to go down rapidly proportionate to the quality of the models at the cost of inference or whatnot, but the cost of reproduction is also going to go down. So say for example, you create some website, you make some SaaS thing or whatever, and you're actually not super differentiated in the market, the cost of me reproducing that and assuming I can somehow... I don't want to differentiate with you either for whatever reason, and I want to race to the bottom with you, even if you assume the LLM and whatever will have certain unit economics, there's a really fascinating dynamic there where lots of products and services suddenly become much more competitive than they used to be in the past unless they're extraordinarily differentiated.
- 00:59:11 And so a natural conclusion that I went, and again, these were all assumptions, there's a lot of chains of ifs here, but I started wondering like, oh, wow, does the software industry start to resemble the music industry way more where there are one or two artists that can change the hotel prices in any city they visit, and almost everyone else is at the very opposite end of that spectrum, unless they're providing an extremely differentiated product or service that allows them to maintain competitive advantage and keep those profits going.
- 00:59:51 And then that brought me... This is kind of like... I'm trying to thread all the things you mentioned. This is what actually brings us to wisdom because where do good ideas come from? Because if what we're saying is that in the world of bits, not in the world of items, that maybe has a different set of dynamics to it, but we can talk about that. I know a lot less about the world of items and what it takes to manufacture and whatever, but if what

we're really saying is that the economic curve that seems to have a lot of momentum behind it. And again, this isn't a binary thing, right? It's going to be a gradient, but also discontinuous where maybe if you're a small enough team, what that means is you just don't hire a PM for a while if you can just all do it yourselves. You just don't hire too many engineers if you can do... It creates... As the curve, even though the curve itself may be smooth, the consequences might be very like non-smooth and discontinuous in the impact it has in the market and that's very difficult to predict.

- 01:01:02 But if you keep going along this curve, and if you see that, oh, okay, the value of differentiating properly in the market is only going to increase, and that means that it's going to be greater and greater at premiums on kind of insight, the systemic cultivation of insight, where do good ideas come from? Well, it comes from human beings.
- 01:01:29 And then you can also talk about sci-fi, about models do it or whatever. There's a bunch of philosophy there, we don't have to go into that unless you want to. But for this thought experiment, let's just say that there are still human stakeholders at the end of the day in the incorporation of the company and whatever. Where do humans get good ideas? Well, it turns out a lot of the literature, or at least I'm still very much a student of this, and like I said, I'm still learning this, but it seems that there's a lot of similarity between the various processes for the proactive cultivation of wisdom and the proactive cultivation of curiosity, open-mindedness, exploration, overall adaptability.
- 01:02:12 And so if you take all these ideas seriously, it's sort of like in a world where knowledge is becoming rapidly, rapidly commoditized, the key differentiator is to the extent that you are wise as a person, how pro-social are you, how valuable are you to the community? And how

self-authored are you? Right? And by self-authored, I don't mean I have the license to be [inaudible] but are you, in the Joseph Campbell sense, following your bliss, right? Are you voluntarily individuating and going along your life path and cultivating wisdom?

- 01:02:57 And so that kind of ended with talking with Natalie and talking to a few folks. That's why I think what excites me about this technology isn't the act of kind of merely automating things because that very rapidly... That framing takes you to a world of, I think... One way to look at that frame is it takes you to a world of perfect competition potentially, which is not what you want, if you want to... For lots of reasons, that's not the world you want to live in. But if you want to lead... Create differentiated value in the economy, what that takes you is you don't want... The framing for me isn't how can the models, how can these AI systems be more agentic, it's how can I be more agentic?
- 01:03:46 And it is by me being more agentic, I will hit my own limiting assumptions, behaviors, whatever, and then therefore grow in wisdom. And so how can I integrate this technology as part of my cognitive stack to be a more effective human being in meeting my aspirational goals, whether that is showing up in a certain way for my family, friends, community, et cetera, for my business. And so that's kind how I tie all that together. And then once I started, it's still... There are large sections of this map that are very unclear to me. But once I started seeing this thread, I've been interested in meditation and wisdom practices and stuff for many years now, therapy.
- 01:04:33 But that became really powerful for me because I started getting curious, how can you start connecting these things together and what does it mean to connect them together and what does it mean... And that ended up having deep insight also and how to effectively wield these

models because I forget who said this, it's like the medium is the message. We use natural language to talk to these things. Language was made by humans, for humans. These models obviously aren't human, and language makes lots of cognitive assumptions about humanity. And so I think it takes you to some really interesting places in how you can interface with these machines much more effectively to create value in the world. And so that's kind of my ramble of the through line through all this.

- | | | |
|----------------|----------|---|
| Jon Krohn: | 01:05:24 | I asked you a very long complex question, and so it makes sense that it took a very long answer, but also super fascinating. I think that quote, the medium is a message is actually a Canadian communication theorist, Marshall McLuhan. |
| Varun Godbole: | 01:05:38 | Right, yes. Right, right, right. |
| Jon Krohn: | 01:05:40 | And yeah, everything that you said there is fascinating. It's interesting to me that you got to this idea of people needing to pursue wisdom in order to be able to differentiate and basically save the economy. I love the idea. |
| Varun Godbole: | 01:06:01 | I don't know that I'd frame it in terms of saving the economy. I would say something... Here's how I frame it. It's not clear what the limits of automation are. I genuinely don't know. And it's that it's possible this is the best the models are going to get. I personally don't believe that, but I also don't know how much better they get, I just don't. |
| | 01:06:34 | And I think, this is getting a bit philosophical. Again, this is what I find so fascinating is the philosophy which I kind of never really cared about, is just feeling so much more tangible for me in my life right now when I think about how to solve real problems with these models. And |

it's not sort of like, how can we save the economy? To me it's more like the economy in my mind, the economy is there or should be there for human flourishing. That's what I personally care about, right?

Jon Krohn: 01:07:11 Sure, yeah.

Varun Godbole: 01:07:12 And it's like I can't change the economy writ large. Frankly, I don't even want that kind of power, I don't think I could shepherd it wisely. But I know that in my own life there are ways in which I'm likely self-deceiving, the ways in which I'm foolish, the ways in which I want to be better aspirationally. And it's just reflecting on the fact that I don't know what's going to happen in the future.

01:07:40 But what I do know is that the best things in my life have come from working on myself and kind of working on myself for showing up better every year in therapy. That has been the best thing that ever happened to me.

Jon Krohn: 01:08:00 To be the best prompt engineer that you can be

Varun Godbole: 01:08:02 For myself, right?

Jon Krohn: 01:08:04 Yeah. Right.

Varun Godbole: 01:08:06 Well, I think that's something deep here, right? It's like even prompting, there's something really deep. It's like how often before you solve a problem, you prompt yourself, okay, what are the things I need to think about? And there's something inside of us that starts generating the answer. I find that kind of fascinating, and I find it fascinating how this finally really grounded models of cognition for us to understand ourselves better. Because what I care about is I want to try to be a better person. It sounds a bit cheesy, but that's after reality smacking me around a bit over the years, I've learned that is actually the most useful thing to do with my time.

01:08:43 And so I think if it is less of saving the economy and it's more of like how can we as individuals have the tools to engage in human flourishing at least within whatever we control. And then to me, the economy is just the emergence of whatever that is, right?

Jon Krohn: 01:09:03 Yeah. Yeah. And that's where I was hoping... That's where I was going to get anyway, which is to say that that I think is the noble goal here, and that's, as I said at the outset of the episode that is now the message that I'm evangelizing as to where we should be going with AI is to be allowing for human flourishing, to be supporting us as much as possible. So, brilliant.

01:09:24 Yeah. So I'm going to have a link to your newsletter as well as this particular post, From Knowledge to Wisdom: Value Creation in the Age of LLMs. It seems likely that by the time this episode is out, you will have other blog posts related to this as you have more thoughts and flesh them out more fully. I love that, since you are now on sabbatical or in retirement, you describe yourself-

Varun Godbole: 01:09:49 sabbatical, sabbatical. You said retirement, and I just clenched when you did that. Sabbatical. I'm on a break. I just want to chill out for a bit, but then I want to work.

Jon Krohn: 01:10:00 So as part of that vibe driven development, I realize you're mostly talking about probably software development, but you're also developing a lot as a human it seems. And so on that note, I would love to hear, I think based on what you told me before we started recording, you might have a book recommendation for us that would help us all with our personal development.

Varun Godbole: 01:10:20 Yeah, I would highly recommend Siddhartha by Herman Hesse. It's a book that I have found myself reading multiple times. I think that book is really deep. Yeah, I would highly recommend that book.

- 01:10:39 And if actually as a bonus, if you want, there are some people who like videos more than they like books. If what you want are videos, I would highly recommend the lecture series, Awakening From the Meaning Crisis by Professor John Vervaeke, from Jon's Canada. He's a professor at the University of Toronto, and I think he's got some really great stuff. And even though it says Awakening from the Meaning Crisis, it's about meaning and salience and so forth, I will say that watching that lecture series has substantially leveled up my thinking around LLMs, prompting and cognition.
- 01:11:23 And it's a sort of similar connection in spirit where it might seem kind of left field, but it's similar in spirit, how Nonviolent Communication was helpful for me for prompting more effectively or thinking about interpersonal communication patterns. In the same vein, what that lecture series will afford anyone who goes and watches it is a much deeper kind of exploration of your cognition and how you make meaning and what you find relevant, and therefore, I think it will help you have more empathy, if that even makes sense for an LLM, or to think more systematically about how to prompt them more effectively, how to wield them more effectively, how to more effectively design product surfaces to allow your users to interface with them more effectively, to give you more empathy for how your users may want to interact with these systems.
- 01:12:23 It seems kind of philosophically, but I would actually say for the people that actually want to build good experiences and products and think about these things rigorously, I think the Awakening from the Meaning Crisis lecture series by Professor John Vervaeke is an excellent resource.
- Jon Krohn: 01:12:40 Nice. Thank you for that suggestion as well. So Varun, you had a lot of fascinating thoughts to share with us in

this episode. Of course, I'm going to have your newsletter for people to subscribe to. It's a Substack, so it'll be free to subscribe to, at least for now, maybe someday it's going to be a huge paywall.

- Varun Godbole: 01:12:59 Probably not. The Substack, just for folks, just to manage the expectations of anyone who clicks subscribe, this is not my job, I don't want it to be, I just want to get better at writing and I enjoy it. And so this is just a space where I kind of write down what I'm thinking about. Different posts will have wildly different levels of editing, and it's mostly just a way for me to share what I'm excited about with my friends and to create a space with people to comment back about what they found exciting or not, and I don't foresee it becoming... Yeah, I want to keep it fun and, yeah.
- Jon Krohn: 01:13:42 Very nice. Well, in addition to your newsletter, how else can our listeners follow you after today's episode?
- Varun Godbole: 01:13:48 I think the newsletter is probably the best way. I have a LinkedIn account and a Twitter account, but I pretty much never tweet. Sorry, X, I never X. I never tweet. I don't know. I never do that. I lurk on there. And even on LinkedIn, I never really post unless I feel like I've written something that friends have told me that they liked, and then I'll post it on there for other people in case they like it. So LinkedIn and Twitter, but I don't really use them as much. Substack is probably the easiest way kind of get in touch because you can just reply to the subscribe and I'll get an email from it, so it's like... That's probably the easiest way.
- Jon Krohn: 01:14:36 Very nice. Awesome, Varun, thank you for taking the time today. I really enjoyed today's episode. It was seriously mind-expanding, and yeah, hopefully we'll catch up again in the future and see what you're up to.



- Varun Godbole: 01:14:46 Yeah. Thank you so much for having me. This was fun.
- Jon Krohn: 01:14:49 On an episode with Varun Godbole. In it, he covered how skills learned from therapy, like being explicit about needs and avoiding assumptions translate directly to effective AI prompting. How AI tools are breaking down traditional role boundaries, potentially leading to full stack employees who can work fluidly across different domains using AI assistance. The economic implications of AI making knowledge increasingly commoditized, potentially leading to winner take all dynamics similar to the music industry. Why focusing on enhancing human agency and wisdom through AI might be more valuable than pursuing autonomous AI agents. The importance of metacognition and self-examination in an AI augmented future where personal growth and wisdom become key differentiators, and how understanding human cognition and meaning-making can lead to better AI interactions and product design.
- 01:15:43 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 Varun's social media profiles, as well as my own at superdatascience.com/869, and next week I'll be speaking at the RVATech Data and AI Summit in Richmond, Virginia. I'll be doing the opening keynote, that's on March 19th. It's a day long conference, there's a ton of great speakers at it, so it could be a great opportunity to meet in person and enjoy a great conference, especially if you live anywhere near Richmond, Virginia.
- 01:16:19 All right, thanks of course to everyone on the Super Data Science Podcast team, our podcast manager, Sonja Brajovic. Media editor, Mario Pombo. Partnerships manager, Natalie Ziajski. Researcher, Serg Masís. Writer, Dr. Zara Karschay. And our founder, Kirill Eremenko.



Thanks to all of them for producing today's mind-altering episode.

- 01:16:36 For enabling that super team to create this free podcast for you, we are completely dependent on our sponsors, so you can support this show by checking out our sponsor's links, which are in the show notes. Otherwise, share this episode with someone who'd like it. Review it on your favorite podcasting platform. Subscribe if you're not already a subscriber. Edit our shorts... Edit our videos, sorry, into shorts if you'd like to. But most importantly, just keep on tuning in.
- 01:17:01 I'm so grateful to have you listening and 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.