



SuperDataScience

**SDS PODCAST  
EPISODE 995:  
END-TO-END  
FOUNDATION  
MODELS FOR THE  
ENERGY INDUSTRY,  
WITH JAZMIA HENRY**



- Jon Krohn: 00:00:00 Today's guest had a paper accepted at Nurips, the world's most prestigious AI conference in 2023 that made a claim that nearly everyone thought was ridiculous, that scaling bad data can make models worse, not better. Two years later, the entire industry is quietly admitting she was right. Welcome to episode number 995 of the SuperDataScience podcast. I'm your host, Jon Krohn. My deeply impressive guest today is Jazmia Henry , holds degrees from Tulane and Columbia and is partway through a PhD at Oxford. She's held data and AI roles at Morgan Stanley and Microsoft and now she's a member of the technical staff for AI and machine learning at Collide, a Texas-based startup that's building AI infrastructure, including all aspects of specialized foundation models for the energy industry. In this episode, you'll hear about her cutting edge academic contributions and her exciting work at the energy industry frontier.
- 00:00:54 Enjoy. This episode of Super Data Science is made possible by Anthropic, Acceldata and Cisco. Jazmia, welcome to the SuperDataScience Podcast. How you doing today?
- Jazmia Henry: 00:01:06 Doing great. Lovely.
- Jon Krohn: 00:01:07 I am so glad to hear that. I'm not sure if this will make it through the editing process. There might be some kinds of filters that get applied, but right now I can hear birds chirping in your background. Like many of them, it sounds wonderful. Where are you calling in from?
- Jazmia Henry: 00:01:22 Charlotte, North Carolina. Right in the suburbs. Perfect. It's
- Jon Krohn: 00:01:27 Beautiful.
- Jazmia Henry: 00:01:28 Yeah.
- Jon Krohn: 00:01:29 I've only been there once. I was there for a couple nights. We were shooting an ad for Nvidia and Dell that was for



Bloomberg TV and it was so hot. We had some outdoor shots and so the shooting started at like 6:00 AM to try to get it done before it was insanely hot because it was over a hundred degrees Fahrenheit that day.

- Jazmia Henry: 00:01:50 Oh yeah, that must have been the summertime. Yeah. Oh
- Jon Krohn: 00:01:53 My
- Jazmia Henry: 00:01:53 Goodness.
- Jon Krohn: 00:01:54 And I had to wear a suit for this shoot. I know.
- Jazmia Henry: 00:02:01 Did you pair it with a cowboy hat? You got to do the whole vibe.
- Jon Krohn: 00:02:07 I just paired it with sweating and the makeup artist had to keep coming and try to dab it away and make it look natural. No, but Charlotte was super nice. Great food, super friendly people. Anyway, let's dig into the technical stuff here. You are an exceptional researcher, practitioner of AI in the real world. I was so excited to get you on the show. To give a little bit of context, let's talk about where you work. So we actually, I don't think we've ever had an episode. We're coming up on a thousand episodes of this show. I've been hosting for the last 500, 600 of them.
- 00:02:45 And certainly in the 500, 600 I've been hosting. We've never had an episode focused on oil and gas, which is really important. And I think maybe for some listeners, even for myself, we think clean energy alternatives would be great in the future, but for the foreseeable future, we are going to need petroleum products as a big part of our energy mix, no question. And with geopolitical events right now, at least at the time of recording, there's a lot of interest in oil and gas and shipping in particular. So anyway, I thought it would be particularly interesting right now to get you on the show. So tell us about your company and what you do there.



- Jazmia Henry: 00:03:25 Yeah. I work at Collide. I had the benefit of getting reached out by some amazing oil and gas professionals and our leader, his name is Colin McClelland and he's been working in the fracking industry for like over 20 years. When you count up all of the members on our team that work in oil and gas, they have over a hundred years of experience of doing work either from being a geochemist all the way to being a petroleum engineer. And so it was really cool because the idea is you got a whole bunch of people who came out of the boomer generation who were working in oil and gas and they had a particular way of doing it. And then that's when you have a big drop off from the '90s to early 2000s and then things pick up again during the show revolution and the 2010s.
- 00:04:12 And in between that, you have just a lot of knowledge that's just getting lost as those older people are disappearing, younger people don't know where things are. And so the idea is like what happens if we take AI and we put it through all these documents, 90% of the documents in oil and gas is all unstructured. So it's not like, oh, okay, well, we can go through and just put everything in Excel sheets. A lot of the work is handwritten documentation. It lends itself into leasing and titling. It lends itself into safety and regulation. And so what happens if we use AI to go across all of that, make sense of it, synthesize it, but also make things safer for those people who are the reason why we have access to electricity and are able to enjoy data centers, which are the things that give us access to AI in the first place.
- 00:05:03 So it's a lot of cool work and definitely something that if you would've asked me 10 years ago of the industry that would've really made me energized, I would've been like, "Oh yeah, finance," because that's where I was before. But now it's been amazing to be able to work with people who are able to say, "Hey, the AI that you guys are



working on actually saves lives." And that at the end of the day helps you sleep at night. That feels amazing.

- Jon Krohn: 00:05:29 That is really cool. I had a brief two-year stint where I was working in finance and well, I was going to kind of say the opposite because for me- Yeah, it can be stressful. Well, it wasn't so much that. What I was going to say is for me, it was hard to stay motivated about making money for its own sake. We were trading financial derivatives and they were actually, it is interesting, interestingly related to what you're doing because I was mostly trading things in the crude energy complex, so oil and gas futures. Yeah. And some options, mostly futures. And we were never taking delivery of oil barrels or pulling oil out of the ground, but it was a very liquid market. And so for applying algorithms to it and trying to get some kind of return, there's a lot of opportunity and more liquid markets.
- 00:06:27 So that was what we picked to work on. But yeah, for me, it was tough to kind of stay motivated. Whereas actually when you're pulling oil out of the ground and selling that, you're providing something useful to the world economy in a way that you're not when you're doing ... And there's all kinds of things that you could be doing in finance. I mean, obviously there are lots of aspects of finance that are providing utility. I mean, there's financing for any kinds of projects, whether it's energy projects or otherwise happens by financial firms for the most part and that is something valuable. But the stuff that I was doing, it wasn't something that the world economy needed. So anyway, I can see how it would be motivating potentially to be in your industry. So tell us a bit more about your role and what you're doing at Glyde.
- Jazmia Henry: 00:07:17 Yeah. So I am the foundational model engineer there. Every day is something different because we are very early stage and so we're ramping up with much larger customers right now. But the original idea is if you go to



Claude right now and you start asking it a bunch of oil and gas terms, it's going to maybe know the definition of some of them. It's not really going to know how much is related to each other. I think in general, there's a sense, especially in Silicon Valley that oil and gas are like, oh man, if we assist people in oil and gas and there's a moral failing there. But the thing is, is that the people who are drilling oil out of the ground, those are 18-year-old kids, 20-year-old folks, people who could be people's grandparents in their 60s, right? There's that bimodal people who are towards the end of their career who's been doing this for 40 years and people who are very fresh out of high school, you want those people to be able to make it home regardless of if you want eventually things to move to wind or nuclear.

00:08:21 These are the jobs that people have. This is the things that's holding up a lot of the American economy. And so when you have people say, "Well, everybody should be adapting to AI." They want to adapt to AI too, but they can't. There becomes that issue of if I go into Claude and I say, "Hey, read this JSA, which is a job safety report and tell me about the wells in this area. And if there's been some past instances and it's going through the document and doesn't understand the document and says, everything looks good." And you're like, "Okay, I wrote everything right." They're like, "Yeah, you wrote everything right." And then somebody loses an arm or somebody falls into the wellbed and they're not able to continue to walk again, then it's a big problem. You want to have a system that not only knows how to read those documents and synthesize those documents, but also an AI model that's made with that in mind that says like, "Oh, okay, this is a job safety report."

00:09:17 I'm going to shoot it off to identify all of the rules in the municipality they live in. I'm going to give them suggestions. I'm going to push back." You can't have a sycophantic AI when you have people whose lives are on



the line. So when a person says, "Oh, I checked this valve, everything is fine." Oh no, actually looking at the data, this data is different than this other data, here's some possible things that could have happened. You need to go back and look at these other forms before you move forward. So that's been pretty cool work because it requires me to not only just go through and look at documents, which oil and gas has no shortage of, but it also requires me to go and talk to people who are actually in the field because there's these really cool documents where they might be talking about how big a well is and they might have a point in a map and it might say something like, "Okay, the point in the map of how big the well is, is you're going to smoke three marlboros and you're going to walk east and you want to smoke, you're going to smoke two more, you're going to go west." And it's like- Oh

- Jon Krohn: 00:10:22 My goodness.
- Jazmia Henry: 00:10:23 What does that even mean? I don't smoke, nobody around me smoke. How will I know that's an old school way of kind of judging the limp or something. It is
- Jon Krohn: 00:10:36 The most old school method
- Jazmia Henry: 00:10:39 Of measuring
- Jon Krohn: 00:10:39 I've ever heard. That
- Jazmia Henry: 00:10:41 Is
- Jon Krohn: 00:10:41 So funny.
- Jazmia Henry: 00:10:42 And so you want to have AI that's able to say like, "Oh, here's where that is. Here's that trajectory of how far that is. " And so yeah, I mean, that's the work that I do, but we have people on our team that do some amazing work that even builds around that, anything from identifying like, "Hey, this well that is an inactive well is close to a



school. And so you'll want to go ahead and do X, Y, and Z things to make sure it stays secure so that you're not having oil drip and lean out into these different areas." And then you have other folks that do work like leasing and titling work and being able to make sure that those documents go to the United States government appropriately so that everybody's compliant. So there's a lot of really great stuff that comes out of building that underlying foundational model and then everybody taking that model and building on top of it.

Jon Krohn: 00:11:40 I love it. Really cool work. To get into some more detail on it, you describe yourself as a full stack foundation model builder, which makes a lot of sense to me. So like end to end, all aspects of it and you're at the very cutting edge of machine learning research and language modeling research we're going to get into NERPS papers that you have for listeners who don't know NERPS. It's the most prestigious academic AI conference that there is. And so yeah, tell us what you mean by full stack foundation model builder. What is involved other than finding out how many Marlborough it takes to get to a rig, what does it involve?

Jazmia Henry: 00:12:19 Yeah. So really four steps. The first step is data curation. So this is the step that everybody hates. I actually yesterday sat down with my CTO and we had a very large whiteboard where we were just kind of going through and coming up with a multi-stage plan. We've been working on just data curation over the past six months. We have people who focus specifically on that. And what that means is I might get a textbook of a thousand pages and within that textbook, each paragraph might have a highlighted word and then around the highlighted word, it has a definition and then it has different type of topics and a topic might start off in chapter one and not be talked about again until chapter five and to chapter 20. And so how do you get AI to be able to look at that document and be able to synthesize of like, okay, yes, this



thing might have been brought up in paragraph one, but that doesn't mean that paragraph two talks about this thing.

00:13:23 This is not necessarily important. You need to skip over that when it comes to this and find the next time that's important in chapter six. You can do it different ways. You can come up with different graph databasing strategies and stuff, which are strategies that we use, but to enrich that data, you can also have somebody who's a subject matter expert who goes through and goes, "Oh yeah." So the title DV tool means this right here. And then somebody might say DV in this other chapter over here and then somebody might just say tool over here, but you know it means that when it's talking about hydrostatic pressure. So you want to have somebody who can help make sense of all of these things so that when we're creating that graph database, everything's appropriately attached. So that's the first big portion that constantly changes and we kind of make jokes sometimes that at the end of the day, our company is just as much of a data company if not more than an AI company.

00:14:22 A lot of it's just building data models.

Jon Krohn: 00:14:26 Something that we talk about on the show a lot. In fact, in a recent episode of the show, episode 993, we had two book authors on the show. They have a brand new book called *Architected Intelligence*. Their names are Jacob Miller and Jeremy Mumford. And one of the key takeaways from that episode, if not the main takeaway, is that successful AI deployments are about having the right data for your model. And so data engineering can be a bigger part of any successful AI deployment than the model itself. So agreeing with you 100%.

Jazmia Henry: 00:14:57 No, absolutely. And then there's that next step, right? What shape do you want the data to go into when you're using these data pipelines? I have built custom tokenizers



in our office so that when we have people who are pulling things from our ... So first off, starting off with bespoke tokenizers, but building on top of the embeddings models so that when we have people who are doing other processes, which might not necessarily have to do with AI specifically, but still they need that knowledge base from that graph, they're able to use that and be able to, "Oh, okay, well I have a document, a drilling document," and that drilling document is able to better be chunked and OCR works better because it's able to identify the information that it needs to pull from that document versus us actually going through and having to put bounding boxes across multiple things within a document we can better avoid that and on a text-based level versus having to always defer to a more expensive vision model.

00:16:08 So that's another portion there. And then there's the fun part, I guess, the part that everybody gets excited about, which is building the foundational model. There's multiple different ways that we do it. We do some continued pre-training. So just taking a model that we do benchmarks constantly and experiments constantly on the different types of models that are out there, both open source and the big box ones that you get behind an API. We compare which ones you're doing better at certain tasks. And then for those ones that are open source models that we feel like we can move forward with, do some continued pre-training specifically on oil and gas to make sure that we can extend that ability forward. But then there's also doing some reinforcement learning. A model doesn't understand the world around us and so being able to put physical validation around it by creating physics environments and saying, "Okay, go ahead and do these things.

00:17:10 I want you to use Python Rebel to create some code that's doing the mathematics." Yes, but it has to be defined within this space of physics. I can take a very long time to try to teach some model physics, but then it's going to



overfit to physics, especially if I have a smaller model and then somebody put oil and gas on it and everything kind of falls apart or I can just have the very basic rules of physics that are unchanging inside of an environment. And so doing that type of stuff, figuring out what things can we strip out from model training and instead put it in reinforcement learning is something that we do a lot of stuff with. And then the last part is inference, which is the I find most exciting, but a lot of people find less interesting. The way that you train a model is going to affect the way you have to serve a model.

00:17:59 And so if I do some really cool environment that's super amazing and makes the model do better than all the other models out there and I'm like, okay, I love this reinforcement learning thing, it's great. Well, reinforcement learning models are bursty. And so that's going to affect the compute level that we can use when we're doing inference. That's very pricey. So suddenly you have to have this inference thing on so that you can keep your uptime at 99.9% because you have SLAs you got to stick to, but you have this huge cluster that's supposed to be able to handle that burstiness. So being smart with how am I building the model and how can I serve it to the public so that they can get the best results from the model?

Jon Krohn: 00:18:41 This does sound like a fascinating role. One of the most fascinating roles I think I've ever heard of on this show. When you say end to end, really end to end, custom tokenizers, embedding tricks, reinforcement learning, you've got optical character recognition, so machine vision models, tons of text models that you've been describing there. Wow, really exciting place to be working. And then of course the inference aspects. So I actually have a question for you. I feel embarrassed that I don't know what it means for a reinforcement learning model to be bursty. What does that mean, bursty?



- Jazmia Henry: 00:19:19 Yeah. So the idea of burstiness is that, so reinforcement learning models, there's a lot of calculations that is kind of running at a given time and a lot of those calculations can be done on CPU. They don't have to be done on GPU. So the GPU is just sitting there idle waiting for the reward to come out and then from that reward know the trajectory, so the direction that it needs to go into at the next portion of training, which is great because it makes your model do very well, but it also means that you have just like, maybe I might have an eight billion parameter model and I might be able to get away with putting it on an A100 just one and then I can move forward. But if I have that environment that is running and crunching a bunch of numbers and then suddenly it spikes up by sending a whole bunch of information that wants to the GPU for the model to move to the next portion of training, that one A100 is going to fail.
- 00:20:16 Suddenly I need multiples. So that can end up being an issue.
- Jon Krohn: 00:20:22 I See.
- Jazmia Henry: 00:20:22 Yeah. So that's why they call it bursty because it's bursting suddenly with all of this activity that wasn't happening for a long time. And so that's something that a lot of ... When I was at Microsoft, when I first started there, I was working in reinforcement learning. And so before I moved over to AI and I find a lot of people now are starting in AI and then are moving to reinforcement learning. And so there's a lot of people starting to figure out like, oh, we can use reinforcement learning to make the model better. And then they're running into diversity issue and then they go like, oh my goodness, I'm getting all these seg fault issues, which is just like the GPU is just kind of overloaded, it's shutting off. And so trying to be smart with how I build things, but also being way more conservative with the way that I dedicate GPU space.



00:21:14 A lot of times people might go use 90% of the overhead on my model and on the KV cache, I have to scale that down so that I can have space for those reinforcement learning results to come up so the model can get better over time and do that self-adaptive stuff we want to have a model do when we're using reinforcement learning.

Jon Krohn: 00:21:36 I got you. That's really cool. So I would personally consider reinforcement learning to be a part of AI, but it's interesting how you describe them as distinct things. And I think when you're doing that, you're probably thinking of, correct me if I'm wrong, but when you're saying you move from reinforcement learning to AI and a lot of people go from AI to reinforcement learning, the AI you're describing, there's probably like large language models, right?

Jazmia Henry: 00:21:57 Yeah. Yeah. So I'm using it in a sense of how people describe it. So it's funny because when I first entered the field, and again, all these things are underneath the umbrella of AI, but I graduated from graduate school, went and was working, doing some quantitative methods work over at a bank that was where I was in finance and we were using models that like Monte Carlo simulations, but a lot of people might not consider that AI. And then I went and started doing machine learning and building recommendation systems. And a lot of people still might not consider that AI. I was doing recommendation systems and attaching them to natural language processing tooling at the time. This is your glove embeddings work before we started moving into transformers as much. And then that reinforcement learning piece and then over to AI. So I think it's more of how people have defined AI in the past, even though we know that they're all underneath that same umbrella.

Jon Krohn: 00:22:57 Definitely. And something pretty exciting here that I can't believe I haven't mentioned about you, at least in this conversation. Maybe it'll be included in my intro when I



record that after you and I have this conversation, but you and I have something in common. I don't know if you know this, but we've both done PhDs at Oxford.

- Jazmia Henry: 00:23:15 Oh
- Jon Krohn: 00:23:16 Yeah.
- Jazmia Henry: 00:23:17 So cool.
- Jon Krohn: 00:23:17 And so when you're talking about grad school and you're being very, you're humble about the grad school that you've done because you did, correct me if I'm wrong on any of this, I'm basically getting it off of your LinkedIn profile, but you did a master's in political science and economics at Columbia
- Jazmia Henry: 00:23:38 And
- Jon Krohn: 00:23:40 Then you went on to do a PhD at Oxford in social data science and in that whole time focused on these kinds of AI things, generative AI, responsible AI, model explainability, ML ops. It's a very practical for a PhD, quite practical stuff that must be useful now in your role.
- Jazmia Henry: 00:23:59 Yeah, absolutely. Though I must say I did drop out of the PhD at OK, probably not drop out. I paused it.
- Jon Krohn: 00:24:07 I see.
- Jazmia Henry: 00:24:08 You have to do the pause. I paused it last year. I probably had one more year. I was still working on a dissertation. Who knows, probably by the end of the year I'll pick it back up. But yeah, I mean, a lot of the work at Oxford and also too in between there, I also got a technical fellowship at Stanford and a lot of my work at Oxford definitely informs, both was informed by my time of industry in industry, but also informs my time in industry. It's kind of hard for me to figure out where one begins and one ends.



- Jon Krohn: 00:24:45 Yeah,
- Jazmia Henry: 00:24:45 That's good. So yeah, it's been pretty cool.
- Jon Krohn: 00:24:49 Well, I wish I could do another one now because I really ... No, I'm not going to. I'm just saying I theoretically wish I could. I mean, I guess never say never. I have no plans to do a second one, but my point is that when you've been in industry already and then you go to do a PhD, you have so much appreciation for how these things will be used in the real world. Whereas I had just been living in a bubble of textbooks and academic papers until I went into industry and you don't appreciate how special that time is. And
- 00:25:25 I think if I had been in industry first, I probably, not to say that I didn't make good use of the PhD, published a dozen academic papers and so from that perspective of like academic success, it was successful, but if I had industry experience, I think genuinely 10X in terms of the impact and the synergies that I could have noticed using the time well on things that will matter in industry later on, maybe even making impactful strides in the real world while doing the PhD, as opposed to kind of just being in the bubble. So I think that you had a great path there and yeah, I also now see that if I'd clicked on the ellipsis on the dot, dot, dot on your bio, it does say left to focus on industry.
- Jazmia Henry: 00:26:18 Yeah. Oh, that's totally fine. I mean, I will say though, to your point, I think that it is an interesting time right now in AI. If I had went straight to my PhD after Columbia, if I went straight in, I think I would've just finished it because it's like also Oxford has a ... The European PhD and the US PhD is its own other, I'm sure, podcast where we could have a discussion on the differences of it, but three years versus six years is a very different commitment when it comes to ... But also it requires you to have a different, what's the word, understanding of



what you want to do immediately. Like you start Oxford and people already know this is my advisor, this is my focus. Whereas when you guys could get more time because of the timeline to do it. And so I can definitely see on one end if I would've done it immediately, I would've just checked it out because AI wasn't moving as fast at the time most of the natural language processing, even though the Attention is All You Need paper came out in 2017, most of the industry was still using natural language like the NLTK toolkit in Python pretty much up until ChatGPT came out.

00:27:44 And so you could still kind of get away with doing that and taking more time. I think it's because I started a PhD in 2023 and then all of these things are taking off at the same time. I had applied in 22 right before ChatGPT came out, started it and then that was on the run and I was at Microsoft so I was right in the middle of it. I think that that is a lot of why I was like, okay, if I want to do research that's cutting edge, right now it's an industry, but I'm sure the tide will change and we'll be right back to academia and I'll go ahead and finish the PhD and make my professor father very happy.

Jon Krohn: 00:28:33 Yeah, he'll appreciate it if you come back for sure. But regardless, am I correct in understanding based on your LinkedIn profile that basically throughout the PhD you were also working full-time at Microsoft and then later ISO That is wild to me, especially with that kind of European three-year program. I mean, yeah, you've got to be really productive. That's wild, Jasmine.

Jazmia Henry: 00:29:01 Yeah, you got to be productive. You have to have, and this is for anybody who might be listening to this and thinking about taking that route, you have to be open and honest with everybody. So I was on a Microsoft research team when I first started at Project Bonsai over at Microsoft and then I moved over to Microsoft AI. With my original manager, I had to have a conversation with him about



what this PhD was going to be, what types of things I would have to publish. Even during my application process, I applied with them knowing that I would be working full-time and got advisors who had other students who were doing the same thing. One of my advisors, Professor Chris Russell, one of his other people he was advising was at Google DeepMind. So having that understanding that we were going to work on this together, but then throughout the time having Having open and honest conversations both with he and my other advisor, Brent Middlestat, and being like, "Hey, here's what's going on in the industry.

00:30:08 Here's what I'm seeing. Here's what I'm able to do and here's what I'm not able to do. " And even with me deciding to take a step back and go back to industry, that was a group decision. That was me coming to them with a paper I had worked on and a patent that I had worked on and them being like, "Oh, great. Where is this patent taking you? " And I was like, "To a company." And they were like, "Oh, my gosh. If you're going to do a startup, you should probably consider because Professor Chris Russell, he's had his fair share of doing that type of work." So I think that you have to have the right type of team around you. If you don't, it doesn't matter how much you want to do it. It's going to be very hard to get something like that going.

Jon Krohn: 00:30:53 For sure. Well, congrats on getting a couple years into it and yeah, hopefully, I don't know. I mean, not hopefully. Whatever happens, happens, you're going to make the right decision. But yeah, one way or another, it's going to be interesting. And regardless of whether you formally finished the PhD or not, you're doing frontier AI research. Let's dig into some of that.

Jazmia Henry: 00:31:10 Sure.



- Jon Krohn: 00:31:11 At Collide, that oil and gas AI company that you're at right now, you are building rigs, not oil rigs, but R-I-G-G lowercase S. Rigs, it's a model name, I guess you can tell us if Rigs stands for something, but my understanding from our research is that it is a set of petroleum engineering foundation models for the oil and gas industries. And so I guess that ties into a lot of the stuff that you've been talking about already in this
- Jazmia Henry: 00:31:40 Episode,
- Jon Krohn: 00:31:41 But as part of rigs, you're tackling one of the most intractable challenges with large language models and that's evaluation.
- Jazmia Henry: 00:31:50 Yeah.
- Jon Krohn: 00:31:51 So a few weeks ago you dropped a new paper that identifies four systematic failures that make current evaluation frameworks structurally inadequate for assessing deployed agentic systems. As a consequence of those failures, reward hacking is a predictable consequence. So tell us about these systematic failures, what reward hacking is, why it's a problem, and then what your solution is.
- Jazmia Henry: 00:32:20 Yeah. So I kind of was touching on how we use reinforcement learning with our models. And so there is this very interesting thing that has driven me absolutely insane for a very long time in AI, which is regardless of what type of subgenre of AI that you're doing, which is that it takes an insane amount of time for you to figure out if your experiment actually worked. You spend time, okay, I'm going to tune these parameters. I'm going to take this pathway. I'm going to do all this great. Great. Wonderful. You begin moving forward with it. Let's say for ... I'm going to talk about specifically large language models here, but let's say you have your data, you decided your trunk size is for retrieval, you've done all



these things that's super fancy and amazing. You're like, okay, we're going to do QLORA because that's the greatest thing to do.

00:33:18 Great. You begin running training, you even do some checkpointing and stuff like that. Until your checkpoint is finished, you don't actually know how well your model is performing. So you have to wait for a checkpoint for eval. That could be anywhere from two hours if you're doing something very quick, fast, and dirty with a really small model, or if you're using a very large model, it can be days before you're able to actually sit here and go, "Okay, the model is moving in the right direction." Huge waste of time. And again, this isn't just large language models, this was even back when I was doing Monte Carlo simulations. It takes way too long. So my idea was, okay, what happens if we're already kind of spinning up these environments? What happens if we use that Monte Carlo simulation idea? Okay, I'm going to have a model that's training within these certain specs within this data.

00:34:12 And based off of the direction of trajectories that it's moving in within this reward model, because I can quickly pull out the calculations for the reward as it's going before it moves to a checkpoint, I'm going to do Monte Carlo simulations on those and see how well with these types of conditions it's going to be by the time it gets to the checkpoint. So that means that when I get my first burst or my first pass of reward that's coming out of that reinforcement learning model, which is probably about an hour or some odd of time, I can very quickly go, "Okay, based off of these trajectories, it's going to go in this direction versus the other direction." By the time I've had that do it a few times, I'll be able to say like, "Oh, it's moving in the wrong direction. We need to go ahead and chunk this experiment." And then I can run ablation tests as well.



- 00:35:06 So let's say I have an environment that's doing the calculations on a return, I can go ahead and say, "What happens if I have multiple different small variants of this experiment and I change this one small thing? So I'm going to weight the data more heavily towards equations for this set, or I'll weigh the data more heavily towards oil and gas." And I can do the same thing of I can run multiple experiments at the same time, do simulations on all of them, see which one is going to be the best and then move forward with the winner. And so that basic idea was like, okay, this is great because it means that we don't have to wait for days. I think the worst part, especially when you're at a startup So back when I was at Morgan Stanley, back when I was at Microsoft, you just ate that time.
- 00:35:57 You had a company that was big enough to where you can eat that time, you can eat that compute cost, whatever, who cares? Microsoft was charging itself to pay itself for the compute. It didn't matter. But when you're at a startup, you have to be scrappy, you have to identify what can actually work. And also we might have somebody who sells to our customer on Monday, "Hey, Riggs performed well at X, Y, and Z thing. By Friday, they want to go ahead and demo it. I don't have time to wait because if the demo doesn't work out, we might not hit our numbers, we might not hit our funding rounds." So that was where that came from speaking about the failures though. So when you have reinforcement learning, what ends up happening is models are very smart in that they are able to identify the least that it has to be able to do in order to get a good reward because over time it's seeing the same things over and over again.
- 00:36:57 It starts to go like, "Okay, fine. I'll just do X, Y, and Z." And that's another reason why these types of things, these way of evaluating can be so powerful because in that first hour it hasn't seen the same thing over and over again. And so when I do a simulation on how well it's



going to do by the checkpoint, if it doesn't look like it's going to be good, but then it comes out in a checkpoint and it did amazing, I already know what reward had because it went ahead and identified of like, "Oh, she's asking me these questions that talk about math and I don't have to know the definition of something if I just guess at the equation or oh, we're doing retrieval and I'm going to be judged based off of my retrieval at K and she's going to do three and five and judge me based off of how much I do get in three and five." So I'm just going to guess the five most common equations that are used in this set, one of them would be right and then I can move it to the next bit and then suddenly if I give it way more equations, it falls apart.

00:37:58 So that's another issue. And then another one of the failures that comes with award hacking is the data that you're using to train is different than the way people are actually going to use the model. So I'm going to ask questions at eval like there's something called the petroleum engineering exam and at eval, I might have the model test itself on how well I can do at the petroleum engineering exam. Okay, given that you have this well and you have these amount of miles and you have this tool, what's the hydrostatic pressure? Okay, great. People are not going to ask a perfectly formulated created by academics question with multiple choice question and answer about oil and gas. They're going to come in and go, "Hey, tell me about this well." And it needs to be able to tell them about that well and then they're going to go, "What's the decline curve?" Might even say, "What's the decline curve?" Some people just say decline curve.

00:38:59 Some people just say decline. So being able to make sure that the model as well as able to converge to that, that's again where those simulations can come in at because now suddenly I can begin injecting questions into the environment and see how well it does on those questions at the moment. So the model might be doing very well to



be able to answer those SP&E questions and I'm like, "Okay, great, this is exciting, but it might not do well at being able to understand the underlying text and definitions and how things are related. So being able to answer those questions before it gets to the end so that way I can make the adjustments as it goes on.

- Jon Krohn: 00:39:44 Nice. Yeah. So to kind of sum up this reward hacking issue, it's this situation with reinforcement learning models in particular where they are ... The way that you program the reinforcement learning model to learn is by having some kind of reward function where a simple example that people can easily visualize is if you're training it to be really good at Tetris, then you can use the point score in Tetris as the reward because the higher the point score in the game, the better they're performing. But that's easy in Tetris, which is just a video game in the real world when the reward is more complex, you gave lots of examples there, but like for a self-driving car, if you were using reinforcement learning to ... So is it every meter that you travel without an accident, you get some reward and then if you hit a pedestrian, you get like negative reward.
- 00:40:41 And so coming up with what the reward function is, is really complicated in the real world and you can end up with, you might think that you've crafted a really good reward function, but the reward hacking happens when, as you described, the reinforcement learning model figures out some way of getting to like that Tetris high score or that self-driving car high score, getting a really high score in a way that you didn't imagine it would, and that is different than how you wanted it to behave.
- Jazmia Henry: 00:41:12 Yes. And that happens a lot with a reinforcement learning algorithm called GRPO, which is just like a group policy. I don't want my model to just know one thing, which is like maybe math, right? I want it to know, be able to choose the right equation. I want it to be able to exist within the



fiscal constraints. I want it to answer a question about oil and gas well, want it to retrieve well, but in the end, I want the final answer that it has to be correct and maybe the model might learn that, oh, I actually don't have to ... I don't get as big of a reward if I know the math equation correct, for example. So I'll just not learn the math equation correct. I'll just make sure that I get the final answer correct by looking at some keyword that's in a sentence and know that every time there's keywords in a sentence, the final answer is likely going to be one of these four things.

00:42:06 And so when you have that issue, you want to make sure that your model isn't over optimizing just to get receivable award. And then by the time it gets into someone's hands, it's completely useless because it's just trying to get the final right answer. And we kind of see this a little bit with different models on different varying levels that use reinforcement learning.

00:42:33 People have issues with a lot of models where they're asking a question and it will just like, " Hey, you're right, this is actually the correct answer. "And you're like, " No, actually that's not the right answer. "And it's like, " You're right, my bad. "And you're like, " Okay, I need you to be grounded and essentially that's what that paper is for. It's try to have a level of groundedness that's deterministic within it. And also too, another big portion is handling how much it takes to train these models with reinforcement learning. Many times you have to have a reference model, which is a whole other conversation that we can have. And so by having these deterministic environments, you don't have to have a reference model. You can simply take the weights of the model and how it was performing at a given time, use that as a simulations and then be able to compare how well your model is doing, which means lower compute, don't have that big old issue as big of an issue with burstiness and able to actually create a model that works well on a MacBook as



opposed to having to always go for the much larger GPU cluster.

- Jon Krohn: 00:43:46 Yeah. So great explanation there. Jasmina, thank you. You mentioned grounding in there a few times, but I don't think you mentioned specifically that in the paper your framework for mitigating this reward hacking is called grounded continuous evaluation, GCE, right?
- Jazmia Henry: 00:44:00 Yeah. Yes.
- Jon Krohn: 00:44:01 Yeah. Yeah. It sounds pretty cool. So
- Jazmia Henry: 00:44:04 Using
- Jon Krohn: 00:44:06 Simulation-based fine-tuning to make everything happen, get around these reward hacking issues and get all those kinds of intermediate steps that you were describing as well, not just the final answer.
- Jazmia Henry: 00:44:16 Yes, exactly and injecting some determinism. It's so interesting that I think everyone has just kind of assumed that the non-deterministic nature of large language models is just a foregone conclusion. It's like, "Oh, not deterministic. Who knows why decided that? " And I think that it's cool that there is a level of non-determinism. I mean, even in this conversation we as humans, we are existing both on a non-deterministic and a deterministic frame. The non-deterministic frame is that I might not know exactly what word you might say next,
- Jon Krohn: 00:44:56 But I do
- Jazmia Henry: 00:44:56 Know- I don't even know what
- Jon Krohn: 00:44:57 Words I'm going to say
- Jazmia Henry: 00:44:58 Next. Yeah.



00:45:01 But I do know that in the context of this conversation that there's going to be certain guardrails, there's going to be certain things that you aren't going to do. You're not going to jump up and down and start doing a bunch of jumping jacks because- Well, I'm not standing like you are. But identifying that there's certain contextual cues that kind of inform us on different rules of behaviors that we just kind of exist within and we've all decided to trust each other with these things and anyone who violates that we're like, "Whoa, that's a violation of a norm." And so here's me pulling in my undergrad philosophy major here, but we need to begin to inject those levels of norms and determinism in our models. And by having that GCE, I try my best to explain things without using so much jargon because not everybody comes from those same backgrounds, but that's really where that idea comes from of I want the model to have a sense of this is the space where you have to exist and anything that deviates from that is not going to receive award.

Jon Krohn: 00:46:14 Really great analogy there, bringing into human deterministic or non-deterministic norms as well. I love that. Let's move on to another paper that ... Or actually, I'm not a hundred percent sure it was a paper. Definitely it was a NERPS presentation. You can fill me in here. So it was a 2023 NERPS, you demonstrated that scaling bad data makes models worse, not- Yes,

Jazmia Henry: 00:46:37 That was a paper. Tell

Jon Krohn: 00:46:37 Us about that. Okay,

Jazmia Henry: 00:46:39 Great. Yeah, that was a paper that I worked on. So this is where I might sound a little bit cheeky and arrogant here. When I made that paper, everyone thought that I was absolutely ridiculous. So there's this calculation that came out of OpenAI. It's a paper written by Kaplan et al. Essentially, it's this idea that is the underlying rule that people still kind of exist by, but less so, which is that if



you have data, compute and parameters, you can linearly extend the capabilities of a model as much as you just keep on adding on to it. There's some people who pushed back and said, "Well, data's most important. Those people quickly lost." And then there's some people who said, "Well, no, parameters more important." Those people also have been to loss over time. Compute has still remained that important thing where, and we're seeing it now, we have a compute shortage happening right now where people go like, "Okay, just throw more compute at it."

00:47:42 " If you just add more data and more compute, then the model is going to be even more amazing. And back in 2023, I was kind of looking at it from a perspective that comes from economics. So there's a law diminishing returns in economics at a certain point money is great, for example. We live in a country that is very much fueled by money, but if our government is going through something like a recession or something like that and they print a bajillion dollars, then suddenly the value of money goes down. There's a level of printing that has to happen, but there comes at a cap at which suddenly the laws of supply and demand end up becoming an issue. And so that was the idea of the paper. The idea was data is great, bad data is not great for a couple of reasons.

00:48:33 One of them, if you have data that hasn't been deduplicated, which a lot of these models in order to have just that much share data, not only are they not deduplicated, but it's nearly impossible to deduplicate them because they're huge, huge amounts of data. It begins to make the model get more forgetful over time, less intelligent over time, begin to overfit over time. It's so interesting because in large language models, we stopped talking about overfitting, even though that was one of the basics of machine learning. And then the same thing from just a ... I spoke a little bit about compute, but focus more on data, but also from a compute perspective, right? You can't compute your way out of bad data. So it doesn't



matter how many GPU clusters I can have. I can have a machine with 10,000 GPU clusters and all of it has really bad data and the model is going to perform worse versus if I have less compute and I have less data, but it's better data, it's more quality data.

00:49:40 So yeah, those were really great paper at the time. Again, a lot of people were skeptical. I got accepted into a workshop, top 5% paper, which is pretty exciting because I just assumed everyone was going to boo and

00:49:59 Ended up we're kind of starting to see that happening and at play now where you consistently see a 70 billion or 120 billion parameter model be able to outperform a much larger one trillion plus parameter models.

Jon Krohn: 00:50:15 Sure. Yeah. And I did find that paper while you were speaking. I'm embarrassed that I wasn't sure whether ... Well, because it did also have a poster presentation, right?

Jazmia Henry: 00:50:23 Yes. So

Jon Krohn: 00:50:24 It's

00:50:24 Both. And so I think that's where I got confused. So the paper title is Scaling Laws or the Laws of Diminishing Returns: How Scaling Law Quality Data Degrades Model Performance. And so I'll have that for people in the show notes as well as your earlier ... Well, sorry, the paper we were discussing earlier that came out only a couple weeks ago called Beyond Static Snapshots. So I'll have the archive link to that one available as well. Jazmia, this has been a fascinating technical episode, but I do also want to cover a couple of other topics quickly before we wrap up. So you are a champion for linguistic equity.

Jazmia Henry: 00:51:05 Yes.



- Jon Krohn: 00:51:06 So dedicated to ensuring that the future of AI includes the voices of the marginalized through your pioneering work in creating the African American vernacular English datasets,
- Jazmia Henry: 00:51:15 So
- Jon Krohn: 00:51:16 AAVE. Tell us about that dataset and why it's so important.
- Jazmia Henry: 00:51:20 Yeah. So we talked a little bit about my major. So my background when I was in college was a philosophy, English and poli psy and a focus in polysci was on economics. I'm an economics nerd, but the idea came from kind of the way I was raised. My dad is an Afro-Latino Caribbean man and my mom is a West African Black American woman. My mom, Low County, South Carolina, Gullah accent, she just says all types of stuff that I'm like, "I don't think that's fully English." And my dad spoke me different languages when he was growing up and has an accent and I grew up in Atlantic, Georgia and they have their own vernacular way of saying things, way of doing things. And one of the things that I kept on picking up was, and it started off with me doing my thesis in undergrad on this, which is that African American vernacular English is like this tie that binds between the African languages and between English.
- 00:52:37 And if languages can have a tie that bind to each other to where I don't necessarily have to understand exactly, I don't have to understand tweet to kind of get a sense of what somebody is saying because the way that they might say something. And so yeah, I went and was working on that at Stanford because I was like, I think that there can be a way for us to create a corpus that is able to tie these two things together. I ended up not moving forward with it though because I began realizing that linguistic equity, yes, it's about representation, which is very important, but there's also a beauty and actually I was told this by



one of my peers that works in a lot of African language work that actually worked on a model called Afrobert, which has recently been taken down. But there's a beauty in people being able to own and have ownership of their language.

00:54:33 So yes, we want to have languages and systems and linguistic equity from a sense of language models being able to communicate across different people, but also having an understanding that there are certain individuals who feel unsafe by the powers that be. And so giving those powers access to be able to mimic their language, which gives them a sense of feeling safe, but then being able to use it for things that might be nefarious isn't the best thing to do. And so giving people ownership of their language, acknowledgement, but also ownership and being able to say, I am a guest in an AAVE home, even though I identify, obviously I'm African American, but identifying that with my parents being from different places, there is a certain level of guestness that I have to acknowledge and allow people to have ownership over their own thing. If they decide they want to come out with an AAVE language bot, somebody who's 80, somebody who comes from being an American decision of a slave, then that's a tooling that they can use.

00:55:43 I've had some people reach out to me to do that, but I can provide the bridge. That doesn't mean that I need to be the person who crosses it, if that makes sense.

Jon Krohn: 00:55:52 I see. Yeah, that does make sense. Yeah, it's a nuance that I didn't appreciate as I was asking the question

Jazmia Henry: 00:55:58 And

Jon Krohn: 00:55:59 Am now enlightened on

Jazmia Henry: 00:56:00 It. I didn't appreciate it as I did the research. I was like ...



- Jon Krohn: 00:56:04 So in that most recent response, you did mention your family background and so I just want to have ... I've got one last question for you before ... Well, technically I have three questions left for you, but I ask my final two questions. I ask every guest those same two and they're short, quick ones. So this is the last one that's for you
- 00:56:22 Specifically and it goes back even further into your background. So we've talked about, you did your undergrad at Tulane, we didn't specifically mention Tulane, but great university that you did your undergrad at, master's at Columbia. You did the technical research at Stanford HAI Lab, which is a super famous lab that's world leading. The Oxford PhD work, you've worked at Microsoft, you've been a data strategist at Morgan Stanley, a head of machine learning at the Motley Fool. And now you're at Collide at the frontier building end-to-end models for the oil and gas industry and publishing in NERPs. I mean, you're doing such amazing things all over the place. And so I don't know. I don't know if you have kind of insights into how you became like this, that you're able to achieve so much, like what drives you and do you have any tips for our listeners who would like to be competing at the frontier in so many ways like you have over your career?
- Jazmia Henry: 00:57:36 Sure. I laugh because I'm like, oh, I don't know if it's always a great thing to
- 00:57:45 Be like this. I think I've just always been a kid. I already talked about my parents a bit. Atlanta is what they call a chocolate city. It's a very black city where people come from all over the world there, but my specific type of mixture is a more rare combination. And so my sisters and I, we just kind of existed within our own kind of pod and then I have cousins who, like all of my uncles married interracially and so all of my cousins, we all have different racial identities. And so I kind of grew up knowing what it kind of felt like to look like everybody but



not feel like everybody and kind of be excluded from some stuff and accepted in other stuff and try to navigate that world, but also navigate that world with the understanding that we have a country that has multiple layers of privilege.

00:58:42 And in some places I had an insane amount of privilege that were different than other people who looked like me and in other places I had less privilege than people who looked like me. And so I've always wanted to answer why that is, what does that mean and how can I use that experience for positivity as opposed to walking around and being one of those people who's just trying to create these strange intra-racial, I guess that would be called race wars of like, well, Caribbean versus all that silliness or without being somebody who felt like I had to try to overindex by being somebody that I wasn't, or just saying like, "Okay, well, I'll just completely advocate and I'll just only be around kids that were like white kids or whatever." Having a beauty of the fact that I could exist in all spaces and what would that look like?

00:59:42 So that's why I did my undergraduate the way that I did, but that's also why I took the jobs that I ended up taking because all of them in some way were kind of helping me answer questions that I had for myself. When I first left College, I was working for the Clinton campaign and it was like obviously I'm working in politics so that's answering some questions that I have about what I can do to be able to help people. But then I went to graduate school and graduated and the thing that paid student loans was going to be in finance. But then learned very quickly that finances is a lot of the reasons why people have privilege versus not in this country and learning that a lot of the things that are separating people might appear racial in some aspects because they are, but also are class interplays.



01:00:43 And Atlanta definitely gave me a masterclass in that. And so me just trying to learn how to navigate what that means to be petite bizois versus being somebody who might be politariate or in some other space. And then Motley Fool, same thing. It was an instinction of that. This idea of a company that's like, we want to give the average person the same access that everybody gets when they have a financial advisor and me being like, great, I'll do that work. And then we go into Microsoft and working at first for a reinforcement learning team that was actually an industrial manufacturing. We were building reinforcement learning agents that were powering robotic arms for companies like Shell and PepsiCo and Abbott Nutrition when they had a Similac shortage and doing that work and going, wow, this is really cool. I can actually help single mothers or just any mothers who are having a hard time accessing Similac during COVID because I can create a robotic arm that can do that work so that people can stay at home and be able to avoid being too close to each other during those COVID things.

01:01:55 So I think my career has always been guided by that passion of trying to figure out where I exist and how I can exist in a way that can give whatever privilege that I have while also protecting other people who don't have it, while also acknowledging that I don't have all of the privileges in the world, but trying to find ways to express to people who do ways that we can all work together.

Jon Krohn: 01:02:22 You brought together in that final answer, economics, social aspects, philosophy. In a few minutes, you managed to capture a lot of your influences and a lot of the most important factors driving so much that's happening in our world today or at any given point. So I'm looking forward to seeing where your journey takes you next and how you continue to answer questions about the world for yourself and for us through your publications. Really fascinating. So for people who are



also excited about what you're up to, what's the best way to follow you in your work?

- Jazmia Henry: 01:02:59 Yeah. Best way is LinkedIn. Occasionally I post articles or medium or towards data science, but I'm most consistent at LinkedIn, so that's the best place.
- Jon Krohn: 01:03:12 And you probably also crossbook when you do something on Medium or towards data science, you probably post about it on LinkedIn as well. I do. Great. So one stop shop. Fantastic. And then I apologize that I was supposed to warn you about this last question before we started recording. I've forgotten a couple of times with guests recently. I always ask my guests for a book recommendation. And yeah, it doesn't need to be something in our space. It can be in our space. It can be fiction, nonfiction, whatever you want. Do you have anything for us?
- Jazmia Henry: 01:03:44 Yeah. Value Systems Design by Batya Friedman. She wasn't the only one who wrote it, but she's the head author of it. That book has guided all of my decisions in AI regardless of what function of AI I was working in. Essentially the idea of the book is how do you make sure that you're building systems that are rooted in the values that are important to us as other humans because the technology that we create is simply an extinction of human values. And so it's absolutely amazing work. I think everybody should read it.
- Jon Krohn: 01:04:23 Nice. Thank you for that recommendation, Jazmia. And thank you for this whole episode, taking time out of your busy day to enlighten my listeners on so many topics. Really enjoy this conversation, Jazmia, and hope to have you on the show again in the future.
- Jazmia Henry: 01:04:39 Yeah, it'd be awesome.



- Jon Krohn: 01:04:42 Very interesting episode indeed with the brilliant Jasmina Henry in it. She covered her work at Collide, where she builds foundation models for the petroleum industry in which 90% of documents are unstructured and AI's accuracy can be the difference between a routine shift and someone losing a limb. She described her full stack foundation model building as having four stages, first curating and structuring data, then building bespoke tokenizers and embeddings, training the model through continued pre-training and reinforcement learning is the third step and then finally optimizing for inference at scale. She also talked about how reinforcement learning models are bursty because they idle the GPU during reward calculation and then dump enormous loads on it all at once. And finally that reward hacking happens when a model discovers the laziest path to a high reward rather than actually learning the task we wanted it to. 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 Jazmia's social media profiles as well as my own at [superdatascience.com/995](http://superdatascience.com/995).
- 01:05:47 Thanks, of course, 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 and our founder Kirill Eremenko. Thanks to all of them for producing another super episode for us today for enabling that super team to create this free podcast for you. We are deeply grateful to our sponsors. You can support 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 [Jonchron.com/podcast](http://Jonchron.com/podcast). Otherwise, you can help us out by sharing this episode with folks who would like to learn about end-to-end foundation model training, reinforcement learning, and so on. Review the episode on your favorite podcasting platform or on YouTube,



# SuperDataScience

subscribe if you're not already subscriber. But most importantly, I just hope you'll keep on tuning in.

01:06:37 I'm so grateful to have you listening and I hope I can continue to make episodes you'd love for years and years to come. Till next time, keep on rocking it out there and I'm looking forward to enjoying another round of the SuperDataScience Podcast with you very soon.