

**SDS PODCAST
EPISODE 783:
GENERATIVE A.I.
FOR SOLAR POWER
INSTALLATION,
WITH NAVDEEP
MARTIN**



- Jon Krohn: 00:00:00 This is episode number 783 with Navdeep Martin, co-founder and CEO of Flypower. Today's episode is brought to you by AWS Cloud Computing Services.
- 00:00:16 Welcome to the Super Data Science Podcast, the most listened to podcast in the data science industry. Each week we bring you inspiring people and ideas to help you build a successful career in data science. I'm your host, John Krohn. Thanks for joining me today. And now let's make the complex simple.
- 00:00:35 Welcome back to the Super Data Science Podcast. Today we've got an interesting episode for you with Navdeep Martin. Navdeep is co-founder and CEO of Flypower, a generative AI startup dedicated to ensuring clean energy projects, particularly solar power projects, succeed. Previously, she held senior product leadership roles at VC-backed Bay Area AI startups, as well as for AI products at Comcast and the Washington Post. Before that, she was a software engineer for the CIA. She holds a degree in computer science from William & Mary and an MBA from the University of Virginia.
- 00:01:20 Today's episode will appeal to anyone who'd like to hear about the evolution of generative AI technologies in products and applications, including how you can make the best use of the various categories of generative AI technologies today, and how in particular AI is being used to overcome the social and regulatory hurdles associated with combating climate change. All right, you ready for this practical and inspiring episode? Let's go.
- 00:01:42 Navdeep, welcome to the Super Data Science Podcast. It's awesome to have you here. We met at the speaker dinner during Data Universe, which was in New York a couple of weeks ago at the time of recording. Fantastic first run for the conference here in New York. They had a big space, lots of people, great speakers like yourself, and it was a



joy to meet you at dinner. You're doing fascinating things with Flypower, which we're going to get to later in the episode, your company. But for now to start, where in the world are you calling in from?

Navdeep Martin: 00:02:20 I'm from the Bay Area.

Jon Krohn: 00:02:21 Nice. Very popular place to be with an AI startup. Yeah, welcome to the show and something that we're going to do that we don't usually do with guests, but because your background is so interesting, we are going to do the typical podcast thing that other podcasts do, which is going through your background and what brought you to what you're doing today to start off the show. So I understand that you got started off with machine learning and production at the Washington Post just after Jeff Bezos took over there.

Navdeep Martin: 00:02:52 Yes, that's right. I think it was just very fortunate timing for me. Bezos had just purchased the company. I had been with the Post for a while and he was looking for someone to lead an AI program. At this point in time, Amazon had just come out with their recommendation modules. It was very early on in AI's long-term path. So I got the chance to lead that AI program, building out some of their first products, their recommendation module at the bottom of an article, what else you should read, as well as their very first ad tech platform.

Jon Krohn: 00:03:28 Very cool. So are you able to go into a little bit of detail as to how those recommendation systems work technically?

Navdeep Martin: 00:03:36 Yeah, sure. So on the recommendation side, the first thing we had to do was figure out how to keep track of the topics that our users were reading. So our very first line of effort was really what are the topics that we want to have and track for the Washington Post. And at that time it was a lot of the... The interesting nuances about that

was we had to decide how up to date those models should be. So for instance, the Me Too movement was just coming out and becoming a thing that more and more articles were coming out on, and a very real problem was like, should we go be very timely and have Me Too as a topic? Because otherwise the topics were like campaign, politics, financing, school education. They were very broad in nature. So I guess the end of that story was no, we couldn't afford to be so timely, because topic models were very expensive to create. You have a whole data science team tagging, doing this entire effort. So unfortunately that meant we had to keep out the more up to date topics.

- Jon Krohn: 00:04:49 Keeping ML models up to date can be a labor intensive and compute intensive process. And also if you want to make sure that you're doing things accurately and ethically, just the review around all of those updates could be a big hurdle. Nice. And from the Washington Post, what happened after that?
- Navdeep Martin: 00:05:09 From there I went to Comcast and continued the recommendation side for Comcast. It was what should I watch? What TV and movie recommendations should we put on, say, your desktop TV, the mobile application, and that sort of thing.
- Jon Krohn: 00:05:24 And I guess in both of those companies as well as largely what you've been doing since you've been a product manager, but that is despite having a background in computer science.
- Navdeep Martin: 00:05:36 No, that's right. I graduated with a comp sci degree. I was an engineer for 10 years. I worked at the CIA, actually, right out of school in the basement of the buildings before. It's not as cool as it sounds. But yeah, worked as an engineer for all these different teams. And then when I

went to the post, that's when I switched my career over to product.

- Jon Krohn: 00:05:57 Nice. What was driving that change?
- Navdeep Martin: 00:06:01 Yeah, I was always more interested in the user. Why do they mean this? What's the business decision behind this? So I was always pushing my way to be in on those meetings and really got a really great mentor at the time was like, "You might want to consider product management. That's what they do." So it was a really nice transition for me because of course I understood the engineering side and still do, but was able to then really go on to really get focused on the business aspirations.
- Jon Krohn: 00:06:32 Very nice, all right, and then so after Comcast and doing the product manager stuff there, what happened? I think you had a big move in your life.
- Navdeep Martin: 00:06:38 That's right. I was in DC all of this time and I moved to San Francisco. I knew two people here. It was entirely life changing. But I moved for a startup, Primer AI, so got to continue that NLP journey.
- Jon Krohn: 00:06:55 Nice. What did you guys do there?
- Navdeep Martin: 00:06:56 Primer, we analyze news and social media, and the customers who wanted this were actually quite disparate from one another. You had national security agencies like CIA, Air Force, that sort of thing, trying to understand how America is being perceived overseas, trying to understand what's happening on the ground in a given location. And then on the flip side of that, brand managers were very interested in the same type of data, which is like, "What are people saying about my brand?" Walmart, for instance, was one of our customers. How do we address any people that might be not so happy with



their brand? What are our communication strategies around that?

- Jon Krohn: 00:07:40 So automating brand analysis is the area?
- Navdeep Martin: 00:07:42 Automating brand analysis and giving them their action plan of what to do about it based off of the types of things they're seeing.
- Jon Krohn: 00:07:51 Nice. All right. And then after that, I guess you've decided to stay entirely in San Francisco, continuing to move into more and more senior roles, so senior product manager at Comcast and at Primer you became a director of product management, and then VP of product management at another Bay Area startup, Blackbird.
- Navdeep Martin: 00:08:10 Blackbird, yes. So Blackbird was also analyzing news and social media, but with a different angle. So they were trying to understand disinformation, the prime actors of disinformation being China, Russia, really understanding the narratives that are being spread from one social media outlet to another and being able to measure that.
- Jon Krohn: 00:08:33 Nice. And then you took your hand at going to a senior as you can be in a company and being co-founder and CEO. It looks like from your LinkedIn that you did try out a stealth startup for a while and then you moved on to being co-founder and CEO of what you're doing now at Flypower.
- Navdeep Martin: 00:08:53 Yeah, that's right. There was little bit of soul-searching as maybe my LinkedIn profile shows that in a way that I didn't realize. But yes, I wanted to go out on my own and really build my own company. I was ready to do that. This failed startup I had was actually around image analysis at scale and actually took some of the things OpenAI did and was trying to commercialize that at that time. I learned a ton from my failed startup, namely picking the

right co-founder and picking the right problem to solve. So we essentially came up with technology and then we were like, "Does anybody want this?" Versus what you should be doing is finding the problem you want to solve and develop your solution around that.

- Jon Krohn: 00:09:43 That is a classic startup mistake.
- Navdeep Martin: 00:09:46 That's right. So then I went over, and Flypower was initially an AI consulting company. I was trying to, I don't know, cash in on the craziness around... The crazy interest that I was seeing around AI at the time. Realized very quickly that I wasn't able to scale myself. I had a cannabis recruiting company talking to me one day, yet a management training consulting company another day. Every company I talked to was just completely different from one another, and that's where I was like, "I need to focus. Let me focus on one vertical." And that's where climate tech and Flypower came about.
- Jon Krohn: 00:10:29 Nice. And quickly before we get into deciding on that specific area, climate tech and what you're doing now with Flypower, how do you think that being a product leader makes you well-suited to being the CEO of a startup?
- Navdeep Martin: 00:10:42 Yeah, I mean, the old saying is the product manager is the CEO of their own little thing that they're in charge of. So I think I've always... Well, of course I was a product manager, it was nice to be seen as the CEO of your product line, but in so many ways that's true. As the product manager, you were doing everything you can to get this product out the door, so whether it's rolling up your sleeves to test the product one day versus talking to your customers in sales meetings, but learning and bringing that back to the engineers, or just being the scrum master. I mean, there's so many aspects of a



product manager that I think are very parallel to a CEO track.

- Jon Krohn: 00:11:23 Yeah, I think also when you're a CEO, a key part of that role, there's different things. Fundraising is very important, company culture is very important, but another thing that's key is having product market fit, being able to lead development of the product. So surely all of that product leadership experience, and no doubt your technical background in computer science also helpful for understanding the limits on what's possible and the speed at which that technical development can move.
- Navdeep Martin: 00:11:52 Absolutely, for sure.
- Jon Krohn: 00:11:53 Cool, all right. So in all those different roles that you've had over the years, Washington Post, Comcast, Primer, Flypower, there was also an evolution in technical capabilities and the kinds of underlying technologies. AI in particular has been moving very quickly. Tell us a bit about how the technology has changed over time and how that relates to the work that you do as a product manager and the kinds of products that you develop as a co-founder.
- Navdeep Martin: 00:12:20 Yeah, sure. I've actually been giving this a lot of thought because I have a lot of customers ask me, "Well, doesn't AI need a lot of data?" And they asked me a lot of things that I actually started to think about how to frame this for other people to understand. So if you think about it, initially we had human beings. Human beings were making all the decisions. Your IT team at a software company might've been there to fix your computer, help make sure your word processing software was up to date, and maybe to automate some complex computations. But primarily, for instance, at the Washington Post your IT team, where it's doing all these things, but basically the

news editors, the writers were the ones doing all the heavy lifting of actually producing the paper. So this of course is limited by the number of hours per day and the number of people that the Post could actually afford to hire in a given time period.

- 00:13:22 And then you got rule-based software came along. So rule-based software is what I would define as a series of if-then statements that execute until it reaches a conclusion. So using that Post analogy a little bit more, so the Post when this came out, they got really excited about this and they're like, "Oh wait, we don't have to fully write these articles out." So they actually started to use this in the sports section initially. So they're like, "Okay, from the game last night, we'll just fill out the teams and the score and it'll write the article itself." So it was like a Mad Libs type style format if you think about it. So this, though... This was the breakthrough at the time. We had Mad Libs-style if-then statements producing some sort of output.
- 00:14:12 Then you had classic AI, and classic AI is a term I use in my discussion with clients, and this refers to the series of AI technologies before a ChatGPT. And classic AI is what I used in most of my career. It gained widespread popularity when Amazon came out with their recommendation engine in the late '90s. But it was used long before that. Financial institutions, car companies, the US Postal service used it to determine where to send your mail. But once Amazon had a breakthrough with their recommendation engine, companies like the Washington Post followed soon after. So the recommendation engine that I worked on as well as the ad tech software was now possible with classic AI.
- 00:14:59 And then most recently we had generative AI enter the scene, and this shifted the paradigm not only by seeming to understand our text, but actually changed the



applications we could use this in. So you no longer needed a lot of data and your data could change. So it was in fact built with the expectation that your data is going to change. So for media organizations like The Post and others, they can now write an entire article with generative AI. We no longer need these attempts at doing so. It's possible to do so. And just note though, I have no knowledge that the Post is doing this, so I don't want to spread any rumors, but I do see journalism, I do see news organizations looking to do this. And the latest, I think it won't be long until we see this happening.

- Jon Krohn: 00:15:53 This episode of Super Data Science is brought to you by AWS Trainium and Inferentia, the ideal accelerators for generative AI. AWS Trainium and Inferentia chips are purpose-built by AWS to train and deploy large-scale models. Whether you are building with large language models or latent diffusion models, you no longer have to choose between optimizing performance or lowering costs. Learn more about how you can save up to 50% on training costs and up to 40% on inference costs with these high performance accelerators. We have all the links for getting started right away in the show notes. Awesome. Now back to our show.
- 00:16:33 Nice, so what would you say is the differentiator between what you describe as classic AI and generative AI? Maybe from a bit more of a technical perspective, what kinds of techniques would fall into classic AI? You mentioned the example there of recommendation systems, recommendation engines. More broadly what kinds of techniques might be associated with classic AI that are different from the generative AI that we're seeing more recently?
- Navdeep Martin: 00:16:57 Yeah, I mean I think behind the scenes, if you think about classic AI models, you're training a model to do something with texts that it's never seen before. So in the

Post's example, topics, we had to come up with those topics, and literally, if you imagine... Whenever I explain this to people, I tell them, imagine a spreadsheet of articles that the Post might've written, and imagine tagging that spreadsheet. So for every article you're tagging what topic that's about, and you have the ontology of topics you can choose from. So you can't pick something that's not there. Imagine doing that for thousands and thousands of records. That is what classic AI brought us. And on the unsupervised machine learning side, of course there was... Hey, there's some semblance of trying to understand text, but it was rudimentary, it basically, I don't know, clustered similar texts to one another, so you could do things like understand the topics of Google news or even just for sure was being used to predict disease outcomes and that sort of thing.

00:18:06 But I always say with generative AI, now what this has given you is... Recall I said your data can change as much as it needs to, so you can now deal with changing data, and you don't need a bunch of data to train a new machine learning model. So this now increases all of the use cases that I couldn't even do because my data was changing a lot, which is what we're doing at Flypower actually is like, okay, now I can understand all of these revolving permitting codes and understand community sentiment in a way that I couldn't before. This generative AI now allows you to do that.

Jon Krohn: 00:18:46 Awesome, yeah, so to really quickly recap on the classical AI situation there. So particularly over all of this, whether we're talking about the rules-based systems that you were describing for generating sports articles at the Washington Post, whether we were talking about what you described as classical AI, which involved supervised classification, say to add tags automatically to articles or suggest tags for articles, or unsupervised models which

are clustering related articles together, with any of those approaches... And also, I guess I should quickly say that for listeners who aren't familiar, the key difference between supervised and unsupervised learning approaches, is that with supervised learning approaches, you have some labels. So you would, say, have a whole bunch of newspaper articles and you have them labeled as sports articles, classified articles, world news articles, and so on, politics.

00:19:38 So you could have all of these labels that up until recently you would expect a human would have to label all the samples and then you could train a system based on some new newspaper articles to show up, okay, automatically they should be classified as politics. With an unsupervised learning approach you have, say, all of the newspaper articles, but you don't have some label as to what it is. But you can nevertheless use language to, in a label-free way, categorize articles with similar language together. And when you do that, you might specify with some kinds of approaches like K-means clustering as an example of an unsupervised learning approach... I don't know why I said classified. Something to do with just saying classified newspaper articles.

Navdeep Martin: 00:20:27 You said it earlier.

Jon Krohn: 00:20:28 Unclassified, top secret modeling techniques. K-means clustering is not top secret. It is open source. But it is an approach to unsupervised learning. And with K-means clustering you might dictate, as the practitioner, you might say, "I want to have 10 categories emerge from this analysis, or 5 categories." So you're shaping how many clusters end up, but then say if you pick 5 or you pick 10, you'll end up with these 5 or 10 clusters of newspaper articles, and then you as the evaluator, you can look at, "Okay, you know what, this is clearly a cluster of classified ads and this over here is clearly a cluster of

sports articles." And then you can use that to create your classification model without having to go through all the labor of doing the labeling like we were describing in the supervised learning approach earlier. You can also, interestingly, with unsupervised learning approaches, you can end up discovering some structure in your data that you might not have thought of. So it can be interesting for exploratory data analysis in general. Anyway, I am talking way too much for you being the guest in the episode.

- Navdeep Martin: 00:21:42 No, you explain that much better than me, so thank you.
- Jon Krohn: 00:21:47 So where I was getting to with all of that that I just said is that with the classical AI approaches, everything that you've been talking about is applying those two natural language processing problems. And same thing for your rule-based example that you gave. This was still a natural language processing example. So everything has to do with handling large chunks of natural language, in this case, probably always newspaper articles. Now at Flypower, you're taking advantage of generative AI specifically for addressing climate change. So tell us more about how you're doing that and what generative AI allows you to do relative to what you were describing as classical AI approaches or rules-based approaches.
- Navdeep Martin: 00:22:32 Yeah, sure. So I'll first hit you with a stat. Did you know that 70%, seven zero, 70% of solar projects fail to get off the ground?
- Jon Krohn: 00:22:46 It's shocking. I mean, I do know that stat because you also used it on me at the Data Universe dinner. But it is a shocking statistic. It is wild because we have so much... There's so much groundswell around people wanting solar panel projects, wind power projects, any new power creation that is going to allow us to have a better climate for our children and our grandchildren. So it's such an important thing. We see things like the current federal

administration in the US spending tons of capital on supporting green initiatives. So it's shocking to me that 70% of projects that people try to invest in the beginning... I guess it doesn't necessarily mean that 70% of the capital is wasted because I guess it would be an initial amount of capital, not the whole project capital that it ends up being wasted. But nevertheless, and maybe you have some stats on this as well, a huge amount of wasted capital that could be avoided.

- Navdeep Martin: 00:23:57 No, yeah. So on average it's \$2 million get wasted on a project that never happened. It can go up to 4 million. So people of course throwing good money after bad when things seem to go awry. So Flypower is there to address that. So the top reasons from our research for failures is due to community resistance and the intricate permitting and zoning process. So in the community resistance side, you have whole pockets of people that self-organize to reject solar in that area. And interesting, if you actually drill down into the reasons, there are all sorts of reasons. There's not in my backyard type folks, but there's also, "Hey, this is going to ruin the character of our farming community."
- 00:24:46 And then there's down to people just worried about the land and being able to use it for farming after, say, the solar developers have left town. There's a ton more too. It's quite interesting just seeing this all... Also from my past experience of can I analyze social media to understand community sentiment for brands for a national security agency? So now I'm actually applying that here and finding some really great synergies with just my past toolkit of doing this for other people.
- Jon Krohn: 00:25:21 Okay, so let me try to explain this back to you and then you can help me elucidate how AI is helpful in this. So one of the key issues with solar projects, or I'm guessing

probably any kind of... I imagine there's this thing with wind power projects as well.

Navdeep Martin: 00:25:40 For sure, yeah.

Jon Krohn: 00:25:41 But so with sustainable climate projects, you end up having people like NIMBYs, not in my backyard people, and you were also describing farmers who say, "If you put a solar plant here, or you put wind farms here, you're going to either ruin the character of this area, it maybe impact the value of my home," in the NIMBY case, or in the farmer case that, "We need this land to be farming. That's our livelihood." So somehow you're able to leverage maybe public data and use AI to somehow inform better decisions about maybe where you could be developing this project, where you'll encounter less community resistance or maybe provide you with ways of addressing the community resistance in a constructive way.

Navdeep Martin: 00:26:31 Yeah, you're on it. So a few things. I think what generative AI... And I did a small consulting project right before I went down this path of analyzing websites, analyzing websites that are going to change on occasion, and how do you think about that? How do you scrape a website, how do you understand its contents? So having just done that, so that's a generative AI problem, and that's not a classic AI. Generative AI is helpful because I don't need a bunch of data to train to go do this. I'm going to do this one time, I'm going to keep track of this information and then ask questions of it. So I need to do it just for this one single use case, and then I'm moving on to my other features.

Jon Krohn: 00:27:20 So you're there, you're collecting data from the web, I guess, or other resources, and then you're fine-tuning a large language model so that you can ask it questions or maybe using RAG, retrieval augmented generation, to be able to pull back relevant documents from a large

repository of documents and then have the generative AI model on top from those relevant documents that were brought back, feed all of those relevant documents into the context window of a generative AI model so that it can answer questions in natural language.

- Navdeep Martin: 00:27:54 No, that's exactly right. So yeah, picture for a solar project, what that means is, "Okay, I want to understand what the permitting codes are in that area to help elucidate that process. I want to understand what's being said on social media about solar in this particular area." The local news outlets is another really great source of information. There's also energy news outlets that also talk about this stuff in these areas. And then some of the more really nuances very particular to solar is you can look at approved projects and failed projects and derive a lot of insight from that. Even the approved projects have a long story because they might've taken two years to get approved, and by the way, they had to do all these concessions to get approved.
- 00:28:42 So if you take that as a baseline, the approved projects and like, "Okay, there was dissent in the following areas, but this is how this particular developer was able to get this project approved." We now have a community action plan, so to speak, that we can provide when... Hey, if you see residents that are concerned with the heavy construction going across their small town roads, here's how this other developer addressed that successfully, and here are basically your talking points as you go into that community.
- Jon Krohn: 00:29:20 Data science and machine learning jobs increasingly demand cloud skills, with over 30% of job postings listing cloud skills as a requirement today, and that percentage set to continue growing. Thankfully, Kirill and Hadelin, who have taught machine learning to millions of students, have now launched CloudWolf to efficiently

provide you with the essential cloud computing skills. With CloudWolf, commit just 30 minutes a day for 30 days, and you can obtain your official AWS certification badge. Secure your career's future, join now at cloudwolf.com/SDS for a whopping 30% membership discount. Again, that's cloudwolf.com/SDS to start your cloud journey today.

00:30:02 Very interesting. Do you also help out with helping your customers ask the right questions? So it sounds like you've developed these databases. I can totally understand why these would be helpful. So information, natural language about approved and failed projects in their area, about permits and zoning in their area, about social media posts in their area, suggesting the sentiment around solar projects. But then I also suspect that as you develop more and more specialization in this, there's probably some kinds of questions that you see as being typical or typically important for your customers.

Navdeep Martin: 00:30:40 When picking out a property or...

Jon Krohn: 00:30:42 Yeah, exactly. So when a customer of yours is thinking about a particular solar project in a particular area, I'm guessing that the service that you offer is more than just saying, "Hey, we've trained this large language model, or we can do RAG over these documents that are relevant to developing solar power projects in your area." But I suspect that more than just providing that tool, there's probably some level of providing them with guidance on what kinds of questions they should be asking the tool.

Navdeep Martin: 00:31:12 So we're not doing a Q&A RAG... We're not having them question the data, because we know what they need out of that. And that's a note that not all users want that, and certainly not in this industry. This is an industry that's newer to AI, so I need to actually give them... So they're getting a report. They're getting a-



- Jon Krohn: 00:31:35 I see.
- Navdeep Martin: 00:31:35 For sure, yeah. And then we do have, on our roadmap, should we add a Q&A chat bot here? That's a question we need to see if customers are even asking for. But yeah, initial V1 is that report.
- Jon Krohn: 00:31:49 I got you. So right now for this initial V1, a customer comes to Flypower, comes to you and says, "I want to build roughly in this region. I know that there's a demand for new energy roughly in this area. What neighborhoods specifically or what plots of land should I be looking at potentially acquiring for my solar power project?" And then you can create reports.
- Navdeep Martin: 00:32:15 Yeah, literally picture Zillow. So you type in Preble County, Ohio, for instance, and Preble County, maybe we've ranked it as an orange. They've had a lot of issues. But here's a town that's 200 miles away that's a green or a yellow. By the way Ohio has all sorts of... They might all be yellow there. But anyway, being able to provide them with heat map of here's what rating we're giving the area you're looking at, but here are some other areas that might be more amenable to solar. Looking at community sentiment, but looking at environmental impacts, looking at cost of energy in the area. It might be really expensive, that they might have a really high electric bill, or maybe they're subject to the pollution of a... What is it called? Typical energy plant is what others call it. So having green energy in the area, there's more than one way to motivate the local residents to actually want it versus fighting it.
- Jon Krohn: 00:33:18 Gotcha, gotcha, gotcha. Okay, cool. So you actually have built at Flypower a tool that allows somebody to maybe self-service, go over different regions within a given state, and be able to see... So you're aggregating multiple different kinds of factors using some proprietary blend

that you've probably developed using your experience to mine these kinds of documents like approved and failed projects, like permits, zoning, social media information, laws, and you're able to combine all of that into some kind of overall proprietary score that gives a heat map by region of red places, places you don't probably want to try to have a climate project, and yellow, and then places where there's green where you're like, "This seems like a great spot to try."

- Navdeep Martin: 00:34:12 Yeah, for sure. And then we're summarizing that. So if someone can read through and then they can click into the links of where do we get this information? From my experience, no one wants to trust the AI and they shouldn't. They should go to make sure that what we're saying is accurate. So we show our work, we show reference to each source where we got this information from and from a UI perspective, there's some nice ways you can do that to make it digestible so it's easy to understand, you only click on it if you need it.
- Jon Krohn: 00:34:44 Very nice. I love it. This makes perfect sense. So the generative AI tools are being used in a pretty fully automatic way, behind the scenes, to populate these reports, to bring back references so that people can dig deeper themselves. That all sounds really cool, really helpful. Do you have a sense... I know it's relatively early days for you at Flypower. Do you have some sense of how these tools, maybe specifically your tool, are reducing that 70% project failure rate?
- Navdeep Martin: 00:35:20 We are early on. We're looking for our first pilot customer, so that's where we are on our path. We are seeing really strong interests, though, in this. No one's doing this. No one's applying this community sentiment piece to this. Interestingly, on the permitting side, we have just as many customers interested in that. So we're a little bit at a crossroads, which one do we do first? And it'll honestly



be that first pilot customer who's ready to go down that path with us is probably going to be the decision maker.

Jon Krohn: 00:35:56 Makes perfect sense. I mean, at the time of recording, your company was founded a year ago, so to have a product built and that you can be selling I think is amazing, to be already at that point, to be looking for those pilot customers, and hopefully things like this podcast are helpful for finding those initial first customers. I look forward to seeing how the journey comes along. Awesome, Navdeep. This is really inspiring. I love what you're doing with Flypower. Beyond Flypower and this particular use case that we've gone over, what are other ways, maybe to get the brain cells of our listeners going on other ways that they can be thinking about AI to tackle climate change, which is probably something of concern to many of our listeners, and I think a lot of our listeners also want to be making a positive social impact with the data science that they do and the companies that they start, the products that they build. So to get those juices flowing, to get those neurons firing, what are some other examples of ways that AI can be used to tackle climate change?

Navdeep Martin: 00:36:56 Yeah, absolutely. Prior to Flypower I worked with a disaster resilience nonprofit. So what they were doing, their name was IBTS, and they were basically advising municipalities on how to prepare for a natural disaster. So in climate tech, you have the prevention companies, which is where Flypower sits. Can we prevent climate change from happening? And then you have the basically the climate change is going to happen, how do we deal with this? So IBTS was on this other side of how are we going to be ready when there's a natural disaster?

00:37:40 So municipalities that they'd work with would be towns or cities that wanted to be ready for a natural disaster in their area. So they would engage with IBTS. The IBTS

team would take 40 hours on average to produce their initial assessment of that area. So let's say it's Fairfax, Virginia, they want to understand, "Hey, what do I need to do to be ready for natural disaster?" There'd be two people producing this report. They'd be doing extensive research all over the web to be able to produce this. So their finished product was actually 26 pages of content conducted by these two individuals.

- Jon Krohn: 00:38:18 Very cool. So this was IBTS. What does that stand for?
- Navdeep Martin: 00:38:22 IBTS stands for Institute for Building Technology and Safety.
- Jon Krohn: 00:38:26 Nice. So while Flypower is doing its best, its utmost to avoid climate change happening as much as we can, there's other folks out there IBTS that understand that the world has already warmed, we're already seeing disastrous effects, some regions are more impacted than others, and IBTS is looking to be able to come up with solutions to prepare regions that are likely to be affected. So my apologies here, but the AI here is that there's a generative AI element in what they're doing?
- Navdeep Martin: 00:39:04 So IBTS, as a part of that research effort, they would go to the county's website first off, so orlando.gov, whatever the county dot gov typically. And those websites are massive. So what they would be trying to do is ascertain whether they had taken the best practices themselves to be ready for that natural disaster. So in addition to their website, they'd go to these other ones, like CDC and FEMA, and really trying to understand that particular area.
- 00:39:39 So they came to me and said, "Hey, can you make this easier for us? We want to try out generative AI. We think this is a really good use case." So we went down that path and said, "Okay, well give us your finished reports." And they didn't have that many, actually. They had maybe

10... I think it was less than 10 finished reports. So again, this is great use of generative AI. I don't need a ton of finished reports. I just need to generally understand the questions. So their reports were questions and then here's the answer that they had produced from these websites. So yeah, taking that and being able to do that for a county it's never seen before.

- Jon Krohn: 00:40:16 Gotcha, gotcha, gotcha. So this is actually... There are quite a few parallels between what you're doing at Flypower and what companies like IBTS are doing. What's different is the underlying data that they're using for their generative AI systems and for the kinds of reports that they craft.
- Navdeep Martin: 00:40:32 Yeah, that's right. I think it's like a... This example is symbolic of what many other companies are going through right now, which is how do I store this data so that when asked a question, it picks the right thing on the retrieval augmentation side, but also what questions should I be asking? And that's the fun part. But yeah, how do you tweak this so that it's not hallucinating, that it's not coming up with the wrong stuff?
- Jon Krohn: 00:41:03 Nice. Makes perfect sense. Do you happen to have any technical tips for listeners on how we can avoid hallucinating with our RAG systems or what some of your best practices are?
- Navdeep Martin: 00:41:13 Yeah, I have some ones that I really love, actually. So there's the one where you're framing it, telling the LLM that you're the expert, you've spent 20 years in sustainability, answer the following questions. But there's another one that I just discovered at Flypower that I really like, which is answer the following questions. If you don't have an answer, leave it blank. Leave this blank. And interestingly, just saying that, we've been only getting right answers. Can you believe that? We've solved

hallucination. Or it looks that way. So our next step is to make sure we didn't miss anything, it's not missing facts.

00:41:57 And then there's the concept of just having metadata about your data. So if your data source is going to continue to accrue over time, like in the Flypower example, understanding the permitting keeps changing the community sentiment obviously is an ongoing thing, making sure that you are structuring your data, so your data storage, you're actually keeping track of what should go where, I guess is the best way to say it at a high level, but so that you can build this thing over time and accept that new and incoming data.

Jon Krohn: 00:42:32 Eager to learn about large language models and generative AI but don't know where to start? Check out my comprehensive two hour training, which is available in its entirety on YouTube. Yep, that means not only is it totally free, but it's ad free as well. It's a pure educational resource. In the training we introduce deep learning transformer architectures and how these enable the extraordinary capabilities of state-of-the-art LLMs. And it isn't just theory. My hands-on code demos, which feature the Hugging Face and PyTorch Lightning Python libraries guide you through the entire lifecycle of LLM development, from training to real world deployment. Check out my Generative AI with large language models hands-on training today on YouTube. We've got a link for you in the show notes.

00:43:15 Those are some very cool prompting tricks. I like those. So the first one is telling the LLM that it's an expert in an area of 20 years of experience, and then the second one is that if you don't know an answer, just don't try to answer, leave it blank. That's cool. That seems to have stopped a lot of hallucinating. That is something that these models seem to try to do is that they try to... The RLHF, at least in the last year or two that we see in the leading systems,

there's some kind of priming there. I guess humans have historically liked getting an answer, so there's this impulse that's baked into by the reinforcement learning from human feedback to try to just spit something out, even if it's not confident about an answer. So yeah, that makes a lot of sense. I'd love that prompt idea. I will need to be using that myself. I don't know, maybe this is proprietary and you can't go into it, but do you have experience calling different LLM providers and do you have a reference as to which one you go with for your use case?

- Navdeep Martin: 00:44:15 Yeah, at the moment I think they're all roughly very similar to one another. I don't know, maybe you have a better insight to that.
- Jon Krohn: 00:44:25 Yeah, so I guess among the leading proprietary providers, the ones that I end up talking about the most are obviously OpenAI. They were the first mover. So GPT-4 is very much state-of-the-art from OpenAI, and one presumes that they're working on a GPT-5 which could end up blowing a lot of things out of the water, but in the meantime, Claude 3 Opus from Anthropic, that is my personal favorite now for generating text. I have a ChatGPT subscription for using... There's a lot of bells and whistles in there. Particularly I love the code interpreters, so what they now call advanced data analysis. And I actually did a whole episode on that on this podcast in the past, so looking that up quickly, it was episode number 708, which has a 20 minute...
- 00:45:18 It's like a 20 minute episode or 25 minutes on hacks for data scientists in particular about using the ChatGPT code interpreter now called Advanced Data Analysis. It's incredible at being able to execute code in the browser. So similar to using a Colab Notebook or Jupyter Notebook, ChatGPT is executing code, but it's also writing that code for you as well. So you can upload a CSV file, say, or

upload a document, and ChatGPT will just start asking, "Oh, this looks like a CSV file. It looks like these are the columns. Would you like some exploratory data analysis or would you like a supervised machine learning model or..." So it guides you with natural language, and this is a cringey word to use, but democratizing data science for a broader group of people, because when ChatGPT comes across errors, instead of stopping, it tries to fix the error and usually does.

- Navdeep Martin: 00:46:16 So I'll give you the perspective from the... I haven't programmed in 10 years, so let me tell you this story. So you're in it every day. I haven't programmed in 10 years and I'm like, "Let me get... For IBTS, let me get this code running on my own workstation. I want to see it, I want to be able to demo it," that sort of thing. And basically this was four or five months ago, so maybe ChatGPT has gotten better, but I was like, "Help me..." I installed Jupyter Notebook, all of these things, Python. And basically it kept giving me wrong information and by the time I'd gone through... I realized it's hard to back out of when you've already set your class... What's it called? The path.
- Jon Krohn: 00:47:01 Like your execution path?
- Navdeep Martin: 00:47:02 Yeah, the execution path. Basically if you do it wrong, it's just hard to back out of all of the things I'd already done. And then I started to say ChatGPT... And I had followed my own advice, I guess. I was like, "I haven't programmed in 10 years," because it was assuming I knew all of these things that I didn't. It already assumed I was an engineer. So yeah, I guess the takeaway there is to prompt it with the fact that you don't already have these things set up because probably the people what it was trained on, it was people-



- Jon Krohn: 00:47:35 Stack Overflow and so... Typically engineers asking questions and engineers giving answers. So that's one of the things that all of these big proprietary providers, LLM providers are great at is taking knowledge that was provided in one context, let's say like Stack Overflow for a lot of this coding advice, but being able to convert that into some other style. So to an adult who hasn't written code in 10 years, or to a 5 year old, you can have... Or in the style of a pirate or Snoop Dogg if necessary.
- Navdeep Martin: 00:48:12 No, for sure. And that's the learning curve of using it for our users as you think about it. They don't know what... And that's why I'm not designing an AI bot for any of my users because they don't know what questions to ask and they have no patience. Some people... Because looking at my group of friends, just for instance, out of your group of friends, how many people have taken to generative AI? I mean, we're in a tech circle, so probably more than average, but if you talk to someone in marketing or a doctor in other fields that don't do this every day, it's not exactly intuitive, which is interesting. It's interesting. Not everyone has gone on the wave.
- Jon Krohn: 00:48:55 Absolutely. I had a mind blowing experience recently. I was on a conference on a cruise ship, which was a cool thing to do. It was called Summit at Sea. I think I have talked about this a couple of times on air. Oh yeah, I did because a recent guest, Sol Rashidi, in fact, she might have been... She's one of the most recent guests on the show. So her and I met at Summit at Sea. So I've already talked about Summit at Sea a little bit, but the context of that doesn't matter so much. But what happened on this cruise ship, so you're meeting people from all kinds of different walks of life. There are quite a few people there in AI, like so, "Wow, cool. You're an amazing AI content creator, writer, experienced at tons of Fortune 100 companies building AI systems. I'd love to have you on my

podcast." But you're also meeting very accomplished people from other walks of life.

00:49:48 So at a brunch I was sitting across the table from a woman who is an international surgeon. So she is concerned not only with her patients in, I think it was Texas if I remember correctly, or she works at a hospital, but she also does policy work internationally. So specifically in Africa. And I don't remember, she had some particular specialization, I don't remember what it was, it could have been anesthesiology or radiology or... It doesn't matter. She said to me, "Oh, that's really cool that you're in AI. Can you think of any ways that I could be using AI to be doing policy better for my surgery in Africa?" And I said, "Well, I don't know much about that space, so I'm sure there's tons of ways, but that's a perfect question for a generative AI system like ChatGPT, which I'm sure you're familiar with."

00:50:49 And she was like, "What? Tell me about this." And she was like, "You can just ask questions to an interface?" She wouldn't have said the word interface. But you could [inaudible 00:51:03], and just have your questions answered. So there in real time sitting across from a brunch, open the ChatGPT app on my phone, click over to GPT-4 and ask an extremely simple question. I did no effort at massaging my prompt, just flat out, "What are some ways of applying AI in this surgical area for policy reasons in Africa?" And it gave an enumerated list of a wide range of wonderful ideas, and I showed it to her and she was like, "I can't believe this. This is incredible. You have to share this with me. Here's my email address. Please send me screenshots of this." I'm like, you really should try spending 20 bucks on this because it can change your life."

Navdeep Martin: 00:51:49 No, it's interesting, because you have both sets of people then. For me, as I talk to solar developers, ChatGPT is the

best demo tool ever, because I now have solar developers who also aren't in AI. It's barely technologies even. It's just moving over. A lot of things are still done with fax machines and ink signatures. I mean, this is what this industry is. So the solar developers, though, a bunch of them, they get it. They get exactly what I'm doing because of ChatGPT. So I literally don't have to demo anything. I can just describe how the thing works, and because ChatGPT is out there, they're interested and they understand it. But on the other flip side, there's people who you literally need to demo it for them. And I've had similar stories like that too. You download it for your parents, you download it for whoever hasn't used it, and you just have to show them, and then they're hooked. But somehow there's a bridge that you need to make for those type of people.

- Jon Krohn: 00:52:51 And it's wild to me that this can be worldly super well-educated people like a surgeon who's advising on policy all around the world. It's amazing to me that you can be... So I'm just really keying on how you were saying amongst my group of friends, I feel like everyone knows generative AI and uses generative AI, and it seems I end up being... Because I'm in this bubble hosting a data science podcast and have been talking about generative AI for years and meeting with the best people in the world, obviously I'm very much in a bubble where it just seems to me like everyone must know about these things, and it's wild that there are very well-educated worldly people out there for whom it is still new.
- Navdeep Martin: 00:53:30 No, absolutely. So yeah, as we build these products for our users who may not be in technology is something absolutely to think about and remember and realize there is that learning curve that will still very much be there.
- Jon Krohn: 00:53:44 So we've just taken a very long tangent away from your question about favorite proprietary LLM providers. So

yeah, Claude 3 Opus is my recommendation for at least at the time of recording, the state-of-the-art in terms of getting natural language results back. I still also have a ChatGPT subscription for things like the code interpreter, other bells and whistles that Open AI have built into their apps, and they've been able to get that first mover advantage on. And then I do also use Google, use their Gemini, particularly 1.5 Pro, because it has a million token context windows. So anytime that I want to put in a huge amount of information, you can put in hours of audio, an hour of video. So there have been circumstances where that's been really helpful to me. You can put in 10 novels worth of information.

00:54:36 So sometimes that ends up being super helpful. And one of the craziest things for me is that over that enormous context there isn't even a load time. I mean, obviously you need to wait for your file to upload to the Google Cloud, but then you ask a question, and I was expecting I'd have to make a cup of coffee or something and wait for the answer. I don't even know why I said that because I don't drink coffee. But that idea. And you instantly get answers back, it somehow operates overall that context instantaneously. And I guess you could refer back to our transformer episodes with Kirill Eremenko as to how that's possible. So episodes number 747 and 759, talking about the magic of Transformers. I also did do a whole episode on Gemini 1.5 Pro and its million token context window in episode number 762. Anyway, there's my overview.

00:55:31 Also, we've done tons of episodes on this show about open source LLMs, and it sounds like at this stage with where you are Flypower, it doesn't make sense. It's the kind of thing... I think a lot of similar to your situation at your preceding startup, I think open source LLMs are a place where you can have a lot of wasted development time in a startup because the data science team or an engineering

team are excited about open source LLMs and having their own, but the cost of training actually doesn't need to be that much because you can use things like parameter efficient fine tuning, which we've also had episodes on the show about, so specifically episode number 674 is a good introduction to parameter efficient fine tuning, or PEFT for short. So it can be the case that you can be fine-tuning pretty large LLMs with tens of billions of neurons on a single GPU for potentially hundreds of dollars, thousands of dollars. But it's the inference time costs of having a GPU up and running.

00:56:37 Whereas these proprietary providers, the OpenAI API, the Cohere API, you don't have any of those cost requirements. So in the vast, vast, vast majority of cases, it's a waste. I don't know, maybe I'm exaggerating by saying vast, vast, vast. I'm sure there are a lot of use cases out there where open source LLMs make sense, but there's, I suspect, more cases than not where you're better off, especially early stages, early stage of a startup or early stage of a product idea, way better to be using one of these off-the-shelf proprietary LLMs and just calling their API than trying to fine-tune an open source LLM yourself.

Navdeep Martin: 00:57:17 Yeah, absolutely. I would think of... If I needed to fine-tune something for myself, very specific use cases like law or medical where you have a bunch of data that it's very specific to a vertical. But yes, for the type of thing I'm doing at Flypower, the time to development is decreased significantly, and now the business owners like myself and my co-founder can immediately provide value. What's also interesting I think is the back and forth with your engineer is quite different. Back in the day I was looking at those Excel spreadsheets of tagging data and that sort of thing, and then testing out the results. So now it's basically I'm actually helping up coming with the



prompts. It feels more like you're part of the process versus the person who reviews it after the fact.

- Jon Krohn: 00:58:16 Totally, yeah, it's a lot of fun. Also it allows hands-on practitioners, data scientists, machine learning engineers, software developers to... It's interesting how... The way that you just described it, executives, co-founders, a CEO like yourself, you're able to get a little bit in the weeds on prompts, and simultaneously, even for individual contributors, they are moving in your direction. So it's this evening the playing field in terms of data analysis capabilities, in terms of AI capabilities, because that individual contributor, they don't need to be bogged down as much anymore with writing each line of code. They can rely on these kinds of systems to be doing a lot of the drudgery of the code, and they can be thinking more about product and users and system architecture, and that's going to happen more and more and more and more where it doesn't matter so much to be expert at Python or whatever programming language, it's more about thinking about software systems and data science systems as the software developer than about typing out each character of the code.
- Navdeep Martin: 00:59:29 Yeah, I think absolutely. I think the architecture of the system now is increasingly important as far as who do you hire first? Back in the day it was the data scientists and then you hope they can put together the back end and so forth. And now it's the back end person. Actually on my end it's like, all right, I need someone who can take all of these different outputs because I can help on the configuring and the testing and getting it working to the output that I know it needs to be, but I need someone to put this all together and have it figure out... A very real question for us is how often do we go get this data from these various websites and do we store this data in-house or do we continue to get it and is that going to be really expensive over time? So there are real architecture

questions that I have very first off that we need someone with that type of background to solve.

- Jon Krohn: 01:00:24 Totally. Well said. All right, this has been a fantastic conversation. I've really enjoyed this, particularly near the end of the conversation here, getting deep in the weeds on AI and product development with AI, and throughout the episode it's been great to get your perspective on how we can be using AI for lots of product use cases, but particularly in tackling climate change and associated issues. Before I let you go, I always ask my guests for a book recommendation. Do you happen to have a book recommendation for us?
- Navdeep Martin: 01:00:56 I'll give you my most recent book. I don't know if it's the best book ever, but I've been reading this book called Slicing Pie. So I'm in the process where I've created this LLC. I'm starting to have one co-founder, if not two, join. And I'm trying to address a very real question on how do you think about equity? Do we need to convert to a C-Corp? Are we going to get funding? When do we convert to that C-Corp? Can we stay as the LLC? And what Slicing Pie actually does is it helps you think about the equity question, and more specifically, it actually says, "Oh, 50/50 is absolutely wrong and here's why," and it gives you a very easy to read and understand way of how you basically value everybody's contributions in a way that you're not left in a situation where someone departs from the company or maybe one person's working part-time, but 50/50 for the part-time person doesn't sound like a good option. So really helping you think clearly on that and getting your co-founders on the same page.
- Jon Krohn: 01:02:03 Very cool. Sounds like a practical book for our co-founders out there and our future prospective co-founders. Fantastic. Thank you so much, Navdeep, for this great episode. For people who want to hear more of

your thoughts after the episode, where are places that they can follow you?

- Navdeep Martin: 01:02:23 My LinkedIn profile is probably the best one, so it's /NavdeepMartin, but maybe we can post a link somewhere.
- Jon Krohn: 01:02:32 Of course, we'll have that in the show notes. Absolutely. Awesome. All right, thank you so much for taking the time today, this has been a great episode, and yeah, maybe we'll catch up with you again in a few years and see how the Flypower journey is coming along and all the impact that you're making and how that 70% figure has come down.
- Navdeep Martin: 01:02:47 Yes. Hopefully I'll have some stats for you by then. Thank you very much. I appreciate you having me.
- Jon Krohn: 01:02:59 Nice. I hope you found that conversation with Navdeep as inspiring and informative as I did. In today's episode Navdeep filled us in on the evolution of AI technology, from the rules-based generative AI they used to use to generate sports articles at the Washington Post, to the unsupervised and supervised approaches used in recommendation engines, to the modern LLM-based generative AI systems we have today. She also talked about how Flypower leverages retrieval augmented generation, RAG, to create highly localized reports of how likely solar projects in a given region are to succeed based on things like historically approved or failed projects, permits, zoning, and social media chatter. She also talked about how in addition to AI applications out there hoping to prevent the most disastrous effects of climate change, there are also opportunities in preparing for the negative effects that we are likely to encounter. And she filled this in on her prompting tricks, including asking LLMs to act as though they have 20 years of experience with the



question that we're asking them, and to not respond if it doesn't know an answer.

- 01:03:58 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 Navdeep's social media profiles, as well as my own at superdatascience.com/783. And if you live in the New York area and would like to engage with me in person, this Friday on May 17th I'll be hosting a panel live at the New York R Conference that will feature iconic open source community members, Drew Conway, JD Long, Sumaya Kalra, and Jared Lander. There's always pizza and beers afterwards, so we can catch up over at cold one then. Would love to meet you or meet you again. Huge names like Hadley Wickham, Andrew Gellman, Hilary Mason, Wes McKinney, and Sean Taylor will be there too. It's such a crazy lineup. The New York R Conference is one not to miss, regardless of what programming language you do data science in.
- 01:04:46 Thanks to my colleagues at Nebula for supporting me while I create content like this Super Data Science episode for you. And thanks of course to Ivana, Mario, Natalie, Serg, Sylvia, Zara, and Kirill on the Super Data Science team for producing another practical and inspiring episode for us today.
- 01:05:01 For enabling that super team to create this free podcast for you we are deeply grateful to our sponsors. You can support this show by checking out our sponsor's links, which are in the show notes. And if you yourself are interested in sponsoring an episode, you can get the details on how by making your way jonkrohn.com/podcast.
- 01:05:20 Otherwise, share with people who might like to hear about it. Review if you really like the episode. Subscribe of



course, if you are not a subscriber to the show already. But most importantly, just keep on listening. I'm so grateful to have you listening and I hope I can continue to make episodes you love for years and years to come. Until next time, keep on rocking it out there and I'm looking forward to enjoying another round of the Super Data Science Podcast with you very soon.