DR: [00:00] AI is moving so incredibly fast. Companies keep announcing new models that have upgrades and features and new capabilities. Where are we on that innovation curve?

Neerav: [00:10] I really do think we're going to see rapid capacity and capability increases. The chips in the research are just every year getting two to three X improvements because of the pace we're on. So if you multiply three X times three X, that's nine X. If the models are improving, give or take 10 X every year or two, we feel very plausible that in just five to six years, the models are going to be a thousand times more powerful than they are today.

Announcer: [00:41] This is Telco in 20, a podcast that helps telco execs achieve a competitive advantage with AI and the public cloud. It is hosted by Danielle Rios, also known as DR. Today we're talking to Neerav Kingsland, the head of business development at Anthropic.

DR: [00:59] Hi guys, I'm DR. Man, is there drama in the AI world? The public first became aware of it about a year ago when Sam Altman was unexpectedly fired from his own company, OpenAI. Since those events unfolded, there's been a ton of high-profile turnover at the company, like Greg Brockman's sabbatical, the departure of co-founder Ilya Sutskever, and most recently, Mira Murati, the CTO of OpenAI who resigned at the end of September. But maybe the beef over at OpenAI started before last year. In 2021, a group of key researchers left OpenAI to start their own company. They left over a difference of opinion about AI safety and ethics and set out with a goal to build a new AI model that users could trust.

[01:47] Today, that company is known as Anthropic and they make Claude, one of the best AI language models available on the market. At the center of their business is what they call constitutional AI, an approach to creating AI systems with built-in guardrails designed to make them more trustworthy as they advance. Anthropic is working closely with telcos, and last year announced a $100 million deal with South Korean operator SK telecom. So today we're going to talk to Anthropic about why safety is at the center of their company. What is it about the telco industry that Anthropic sees as an exciting opportunity and how SK Telecom is using Anthropic to become an AI-first telco?

So let's take 20. Neerav Kingsland is head of business development at Anthropic. Hi Neerav, welcome to Telco in 20.

Neerav: [02:38] Hi, DR. I'm thrilled to be here.

DR: [02:39] Oh my God, this is going to be so exciting. We're talking with Anthropic and this whole industry is moving so fast. Literally every day I wake up and there's crazy news going on, and so, so excited to talk to you today. You're the head of business development for Anthropic and Anthropic made news in the telco industry with the signing of a $100 million dollar deal with SK Telecom in South Korea. I'm sure you guys have looked at tons of different industries to apply AI. What makes Anthropic excited to bring AI to telco?

Neerav: [03:09] When we looked at the Venn diagram of what AI can do right now and what industries need, telco really was in the sweet spot. A lot of the core functions and cost drivers of telco, like customer service, which is obviously a huge one and we can talk about other use cases, are things that Anthropic's models can do right now.

DR: [03:31] Yeah.

Neerav: [03:32] And some of the use cases are a couple of years away and we'll be building towards them. But telco is one of those industries where we thought AI could add immense value as we were launching new models.

DR: [03:42] Yeah, no, I'm so excited about that. I recently visited the Anthropic offices in San Francisco. I got the chance to meet one of your co-founders, Daniela Amodei, and she really spent a lot of time impressing upon me Anthropic's focus on safety, which I think really sets you guys apart from the other vendors in the space. And so why do you guys make this such a big focus and why is it so important for the future of AI?

Neerav: [04:09] Yeah, we think that AI could be the most powerful technology humanity has ever invented. And so with that power comes responsibility of making it safe, and that's why we founded it Anthropic. In terms of what that means for enterprises, when you go out there with AI, increasingly AI is going to represent your brand. It's going to be talking to your customers, and it's going to need to reflect your values. So

normally the way us and our peers historically had created models is we train them on the corpus of the internet. They'd come out, they wouldn't be the safest models in the world, the internet can be a dangerous place. And so we'd all hire thousands of human contractors to give the models feedback. And there you were subject to human error. The people who were shaping the values were spread all over the world. You didn't really know who they were.

DR: [04:57] Yeah.

Neerav: [04:57] And so we created something called constitutional AI where we wrote a constitution for our model Claude, and it's on our website and you can see the values that guide it. And what we were able to do is train one AI to give feedback to another AI based on whether it was following that constitution. And so that allowed us to have much more direct control over that feedback process and to iterate very quickly. So that's just one example, but constitutional AI was a major breakthrough that we think will be good for society and also great for enterprise use cases.

DR: [05:31] Yeah, and I think you guys published this constitution, as you put it. It's readable. People can go access it on your website to really learn about how you guys are guiding the model.

Neerav: [05:41] Exactly. We pulled from diverse sources. We pulled from the UN Declaration of Human Rights. We actually pulled from some of Apple's trust and safety language. And we were just trying to cobble together a set of rules that could guide the model to act ethically. And then we wanted to make sure those rules were transparent.

DR: [05:58] Yeah. The industry body of the GSMA also has published their AI ethics. It'd be interesting as you start to fine tune telco-specific models to even add GSMA's view and couple it with a model like Anthropic that already has a good human ethics model on top of it.

Neerav: [06:16] Yeah, exactly. And you can imagine different industries are going to have different rules, finance, law, telco.

DR: [06:21] Yeah.

Neerav: [06:22] So we definitely want to have a base constitution that hopefully gets some of the core rules we can all agree on, but a lot of room for differentiation and specialization when it comes to how you want the model to behave.

DR: [06:34] Yeah, that's awesome. So Amazon and Anthropic announced a strategic collaboration around a year ago, where Amazon invested $4 billion in Anthropic. What are the priorities of this partnership and what are you guys looking to get out of it?

Neerav: [06:46] Amazon's been an incredible partner, both in terms of capital and expertise and in going to market together. Amazon has built an amazing reputation of being one of the safest and most efficient cloud providers out there. You talked earlier about Anthropic differentiating itself on safety. We're also really trying to differentiate ourselves by being enterprise first. And we want to give a safe model on a safe platform, and who better than Amazon to partner with to get out to enterprises? And so it's just been an amazing partnership in putting our models on their clouds where their customers can access Claude directly through AWS.

DR: [07:28] Yeah, and I think that's really interesting because that goes through Amazon's Bedrock product. And what I think is a real differentiator for this is that there's a lot of capability for the enterprise to not share their training data with the model. And I think that's very attuned to what enterprises really care about.

Neerav: [07:47] Yeah.

DR: [07:47] When the ChatGPT models were coming out, some employees of different companies were loading information and working with the AI, but it was getting into training. Microsoft had to turn off access for their employees for a little period of time. Right out of the gate, Amazon's Bedrock, this is one of the key features, making sure their proprietary information is not being released into the general public and certainly not to their competitors.

Neerav: [08:11] Exactly. Yeah, from day one we just do not train on any customer's data and Bedrock ensures that could never happen. And so Amazon's built a great managed platform with Bedrock

where you can access our models, and other models, in an extremely safe and secure environment where your data is safe.

DR: [08:30] That's a nice double layer that's built-in. I didn't realize that, and that's super amazing. But I think there's another thing that's really interesting with this partnership with you and AWS. There's been a little bit of debate out there in the news of where do you run AI workloads. NVIDIA certainly is a strong proponent about building it in a private cloud and on prem. There's been a couple of other players out there, VMware, I guess it's Broadcom, saying AI will be done in a private cloud, and I actually disagree with that. I think it's really important to get access to new chips. I think chips are super expensive. These workloads have a lot of scale, but you don't need them running at scale all the time. And partnering Anthropic's Claude with AWS's compute, chips, scale. Do you guys think about that partnership from that angle, of AI workloads will run in the public cloud?

Neerav: [09:21] Definitely. Right now, I think it's fairly widely considered that the three best models in the world are Anthropic's, OpenAI's, and Google's. And for the largest models, the only way to access those three are on the cloud. So you never know how things will unfold and how that might change, but right now, if you want to work with the best models in the world, you're going to be on the cloud.

DR: [09:43] Yeah. So as I mentioned previously, Anthropic announced this groundbreaking partnership with SK Telecom, and at MWC in Barcelona, there's this great video interview with SK Telecom saying that they want to become, "An AI-first telco." And so how are you guys working with SK Telecom to help them achieve their goal of becoming an AI-first telco?

Neerav: [10:06] Yeah, we talked about how telco was just an incredible industry to start applying AI right now. And when we went out looking for partners, we were just wowed by SKT's ambition. It's one thing in 2024 to say you're going to use AI. It's another thing to really devote resources and shift your company towards it. And we felt that from SKT. So we started building with them from day one. Our researchers and engineers were working with their researchers and engineers to build a telco-specific model, and we worked together to build out a really ambitious

roadmap. So we knew we wanted to start with customer service, which is where we saw the low-hanging fruit.

DR: [10:44] Totally.

Neerav: [10:44] But we wanted to get all the way up to a virtual agent, and it might take a year or two, but we felt we had a partner that we could build with over the long haul. And to date, it has just been great. They've been an awesome partner. They've maintained their ambition and we're thrilled to see what they'll do.

DR: [10:59] Yeah, I totally agree. It's one thing to say, "Hey, yeah, we're using AI too." It's another thing to be AI-first. We're trying to do it at our company as well, Totogi. People always ask me, "How many people do you have?" And I'm like, "Less than whatever I said before," because we're really trying to transform the work to use more AI. And so you just naturally won't need as many people as you used to.

Neerav: [11:20] Yeah, I think it's really important to be open and direct that we're going to go through a big transformation. And I'm highly confident that at the end of that, we'll be a more productive, wealthy, integrated society, but these transitions can be tough. And so it's up for companies, policymakers, governments to work together to make sure we get the best and avoid the worst.

DR: [11:43] Yeah, absolutely. And so you mentioned that you started with SK Telecom with some customer support and service use cases, but you have some bigger ambitions and more ideas. What are some of those big future forward AI ideas that you think will have an impact, not just for SK Telecom, but just more broadly in the telco industry?

Neerav: [12:02] Yeah, I think the next big wave will be in personalized marketing. So Claude can look at your customer data and write a tailored email to you in a second, and it can do that a million times over, and that was just something that wasn't feasible when humans were writing them. And then I think we're going to get out in the field. When you think about workers working on physical towers, maintenance, and so forth, with multimodal models, Claude will be able to use vision to look at the actual piece of equipment. It'll be connected to all the guides and

manuals. In an industry where training can be a big cost, you often see a lot of turnover. I think we'll bring a lot of efficiency to the physical side of the work.

[12:41] And then networking, routing traffic, all of that is a huge cost driver in telco. I think we'll see advances in AI there. And then I think in the next couple of years, the core knowledge work parts of large telco companies, finance, strategy, those will be integrated with AI. We've launched Claude for enterprise, which is a way to start getting Claude in knowledge workers' hands in their day-to-day. And I think that'll only scale and evolve where the AI will be able to do more of the knowledge work.

DR: [13:11] At Totogi, we're actually using generative AI to do personalized marketing, and actually we use machine learning to come up with the right offer and pair these two technologies together.

Neerav: [13:22] That's great.

DR: [13:22] And typically telcos will see maybe a 2% revenue improvement. We're seeing with our customers 10% revenue improvement, which is five X.

Neerav: [13:33] Yeah, that's spectacular.

DR: [13:33] And I'm pushing my team, can we get to 20?

Neerav: [13:35] Yeah.

DR: [13:36] Because I do think telcos really struggle with a lot of different data and a lot of different databases mostly trapped in siloed systems. And when you start to bring all this data together from the customer service side, from the CRM and the customer interactions from the network, you really get this very rich insight of what the customer needs and how they're using the network. And there's tons of revenue opportunity that we haven't even scratched the surface of and it's super exciting.

Neerav: [14:05] Yeah. I think the ceiling is high here.

DR: [14:07] Yeah, it's going to be so amazing. And so keeping with that theme, AI is moving so incredibly fast. Companies keep

announcing new models that have upgrades and features and new capabilities. Where are we on that innovation curve? Is there still a lot of room for improvement? What do you think?

Neerav: [14:23] I really do think at least the next three to five years, the improvements are already baked into what we're doing. And so we're going to see rapid capacity and capability increases.

DR: [14:33] Right.

Neerav: [14:33] For as complicated as large language models are, there's really only three inputs into them. There's the quality and quantity of your chips, so your hardware, there's the quality of your research and code, that's your software, and then there's the quality and quantity of data you have.

DR: [14:48] Yep.

Neerav: [14:48] And two of those, the chips and the research, are just every year getting two to three X improvements because of the pace we're on. So if you multiply three X times three X, that's nine X. If the models are improving, give or take 10 X every year or two, we feel very plausible that in just five to six years, the models are going to be a thousand times more powerful than they are today.

DR: [15:12] That's insane.

Neerav: [15:13] So what we're seeing now is amazing. And even if we froze the technology, there'd be a lot to do, but I definitely don't think the technology is frozen and we're going to continue to see rapid growth.

DR: [15:23] And what I found is that you've got to start the transformation now, even though maybe the model isn't totally ready, can do like 80% of the work, or even only 30% of the work, but start the work now because those models come out and it's a step function change. It's not a linear incremental change. And you need to get your organization in position, people, processes, thinking, so that you can take advantage of the advancements as they come.

Neerav: [15:51] Yeah, this is not something you want to get left behind on.

DR: [15:54] Yeah.

Neerav: [15:55] This is a real technology that's powerful and we're already seeing it. So we surely encourage our customers to just start using it, find a use case, start prompting the model, start learning the skills of building evaluations, start learning about fine-tuning, and us or others can help you through that. But couldn't agree more that starting is the best way to get going.

DR: [16:16] Totally. Like I said at the beginning, nothing this big has ever happened and it's stunning. But you actually have a past life as a writer. You've written two fiction novels, which is amazing. Congratulations.

Neerav: [16:29] Thanks. Well, all of maybe 20 people have read them, but I did take a hack at writing the great American novel.

DR: [16:35] Well, besides this scintillating podcast, I write a blog and it's focused on the move to the cloud and about how AI is impacting telco. But I've got to be honest, it's like a slog to consistently write high-quality blogs. And so what are some tips you can give me on how I can improve my writing process?

Neerav: [16:55] All right, well, I actually used to write a blog too, so I feel your pain. Writing is thinking, and so I think it's awesome that you do it. There's no better way to get clarity in your thoughts.

DR: [17:04] Absolutely.

Neerav: [17:04] I do think using Claude is an amazing tool. I use it to write.

DR: [17:08] Yeah.

Neerav: [17:09] When I was writing a blog, the thing I always tried to do was make sure I had an idea in there. Part of the joy of writing for me was I want to work through a hard problem in public and maybe I'll get it right, maybe I won't, but I'll be part of the dialogue.

DR: [17:24] Yeah.

Neerav: [17:24] And so I always thought sitting down to write a blog post, I was sitting down to try to solve a problem, and writing

was my vehicle to do that, and then I was going to show my work and that was motivating for me.

DR: [17:34] That's a really good way to think of it.

Neerav: [17:36] Yeah.

DR: [17:36] Well, Neerav, this is such an amazing conversation about Anthropic and all this work you guys are doing in telco, and so I really appreciate your time. Thanks so much for coming onto the podcast.

Neerav: [17:46] It was great to be here. Thank you.

DR: [17:51] Stick around. We end each podcast with a "Telco in 20" takeaway. I've got two minutes to tell you something you need to know. AI demands massive computational power and state-of-the-art chips, and Neerav just said that chip performance is improving two to three X every year. This is one of the big reasons I believe telcos need to be using the public cloud for their AI workloads. Operators need to be able to access ready-to-use tools on the latest chips, on demand, and at a low, low cost. But a few weeks ago, Broadcom CEO, Hock Tan, was heard advocating for a return to on-premise infrastructure because of AI, promising cost savings, and faster deployments with VMware. Say what? It's time to clear up why the public cloud is where AI innovation will thrive.

[18:41] Did you know that the hyperscalers are investing more than a $100 billion dollars in AI infrastructure alone and offering access to foundational models and chips with no upfront costs? This saves you the trouble of acquiring hard-to-get servers from NVIDIA, amortizing expensive chip investments you're putting on-premise, and eliminating the risk that they'll be obsolete before you're done writing them off. And because you can move workloads to whatever chip you want in the public cloud, you can always be on the latest and greatest hardware, ensuring your AI workloads are processing at blazing speed and efficiency. The bottom line, the public cloud offers unmatched scalability, flexibility, and access to the latest AI tools and chips. I know what I'd do if I were you, I'd take my AI workloads and I'd run to the public cloud.

[19:29] Want to explore this topic more? Catch me on TelecomTV's AI Native Telco Summit, running October 15th through the 16th. I'm taking part in the Benefits of Becoming AI Native panel where I'll be sharing insights about how telcos can fully embrace AI to reshape their organization. After you watch the panel, you're going to want to talk more. Let's connect on LinkedIn or X at TelcoDR, and don't forget to tune into more amazing "Telco in 20" episodes. Like and follow, and leave us a five-star review. Sign up for my must-read email newsletter on TelcoDR.com, and check out our awesome YouTube channel for more telco insights. Later nerds.