Making Magic with Gen AI: Capital One's Prem Natarajan

SAM RANSBOTHAM: Generative AI requires organizations to carefully balance product innovation, science, and engineering. On today's episode, a leader in the financial services industry shares his experience with these challenges.

PREM NATARAJAN: I'm Prem Natarajan from Capital One, and you're listening to Me, Myself, and AI.

SAM RANSBOTHAM: Welcome to *Me, Myself, and AI*, a podcast on artificial intelligence in business. Each episode, we introduce you to someone innovating with AI. I'm Sam Ransbotham, professor of analytics at Boston College. I'm also the AI and business strategy guest editor at *MIT Sloan Management Review*.

SHERVIN KHODABANDEH: And I'm Shervin Khodabandeh, senior partner with BCG and one of the leaders of our AI business. Together, *MIT SMR* and BCG have been researching and publishing on AI since 2017, interviewing hundreds of practitioners and surveying thousands of companies on what it takes to build and to deploy and scale AI capabilities and really transform the way organizations operate.

Hi, everyone. Today, Sam and I are talking with Prem Natarajan, chief scientist and head of enterprise AI at Capital One. Prem, thank you for joining our show today. Let's get started.

PREM NATARAJAN: Delighted to be here, Sam and Shervin.

SHERVIN KHODABANDEH: Describe your role at Capital One and the history of how you got there, please.

PREM NATARAJAN: My role at Capital One, if we just stick to the AI aspect of it, is to build upon. Capital One has a legacy of being a very tech-forward enterprise. It was the first bank — and, I think, today is one of the only major enterprises worldwide — that is all-in on a single public cloud. That kind of transformation takes both a deep belief in the power of technology and a willingness to mobilize the enterprise, if you will, around that kind of vision. It takes the vision and the willingness to execute, and the energy. And that, I think, puts us on a great footing to then harness the power of machine learning [ML], artificial intelligence, and all of that.

And Capital One is both integrating technology using technology as a transformative tool but also in machine learning. And so my role right now is to strengthen that kind of history, build upon that history of early adoption of a lot of technology. You know, we're in this kind of historical inflection point in AI with transformers and generative AI and all of that, and one way I see my role is to bring the power of all of this new technology to deliver value to the business, to deliver magical experiences, valuable experiences, everyday conveniences, to our 100 million-plus customers to help all of those folks.

SHERVIN KHODABANDEH: You said "inflection point," and I agree we're at an inflection point. Why do you think we're at an inflection point, though?

PREM NATARAJAN: This is not the first inflection point, but it does feel historical in some sense to me. In the past few decades, there have been a few such points, in colloquial AI history, if you will. People like to think of them as AI spring, followed by AI winter, followed by AI spring, followed by AI winter. And I feel each of those transition points between those eras is kind of an inflection point. Initially, it was all these expert systems and all of that, and then we said, "Oh, they don't really scale because they require so much human input."

Then the whole probabilistic set of things came in — Bayesian models. Then they later became, in some contexts, hidden Markov models and all of that for speech and language processing, etc. Those have all been inflection points where we said, "Oh, this thing." And even though sometimes people feel like AI has always been

promising, in many ways in my mind, the previous inflection points in AI history have actually become commoditized, which is the true sign of success.

Like, 20, 25 years ago, using speech recognition in standard industry practice, whether it was for an interactive voice response, seemed novel. Now all of us kind of expect it is there. And so once it is there all the time, we kind of don't think of it as AI, honestly. Like, "Oh, that's just speech recognition." But there was a time when it was the forefront of machine learning and AI.

And so now, this new inflection point, though, I'd say if we think about it as a stack — as a science stack, a capability stack — we go from being able to take phenomena and convert them into some representation, like a speech signal into the sequence of words, and the next step up is kind of interpreting some of that transduction into something meaningful, like maybe some level of semantic interpretation, etc. We keep moving up that slide. Right now, we're in this place where we've built all of these systems. They're showing tremendous capacity to adapt themselves to novel circumstances. But one new thing right now is that they're demonstrating behaviors that they were not necessarily explicitly trained for or designed for. And, if you're completely technical, they'll call it in-context learning. In the popular literature, we refer to them as "Oh, they respond to prompts. They follow instructions."

So that part, I think, kind of substantially lowers the bar for their use. You still have to do it in a responsible way. You still have to do it in a thoughtful way. But it lowers the bar for their use, where all of us can start using these in our pet projects and in our enterprisewide initiatives. So that's the inflection point I see: The developer experience has changed.

SHERVIN KHODABANDEH: Now when you say "these," you mean specifically large language models and the whole stack around that — is that right?

PREM NATARAJAN: Yeah, yeah. From a technical/technology perspective, [they are] transformers in terms of their manifestation as a capability — large language models and

generative AI — and I think it has the power to transform the developer experience. Like, your creativity's top and front and center, and you can use all of these resources relatively easily.

SHERVIN KHODABANDEH: Prem, you have a pretty interesting background. Maybe share a little bit about how you got started in technology and AI and your path to where you are today.

PREM NATARAJAN: Happy to. I should say the early part of the path is somewhat canonical for somebody of my background. I grew up in India, did my undergraduate education. I grew up in a multilingual community and society.

I grew up with in a four-language setting. My family is ethnically Tamil. I grew up in Maharashtra, where a lot of my immediate friends spoke Marathi. Hindi was one of the required subjects in class, and it was a reasonably cosmopolitan neighborhood I lived in, so there were people who came from other parts of India and spoke Hindi. And then English was the medium of instruction, so you know. [Laughs.] And so it was hard to not spot some aspects of language that are interesting. If you're just speaking Indo-European languages, you're used to certain word — subject-object — orders. But then you take something like Tamil, and it's not an Indo-European language — it's Dravidian — and so those orders are different.

So even at a surprisingly early time … I didn't understand there was actually a subject called linguistics. You just said, "I wonder why we say 'come here' in this language and then 'here come' in this other language." And there was some spark of curiosity built in early on in that way, and, again, not necessarily unique to me, but in my case, it kind of triggered some actions later on.

One of the summer internships I did during my graduate school was working on offline handwriting recognition, and that kind of reawakened, I think, my interest in language and its production in some form.

And so then I started working at this company called BBN Technologies, an MIT spinoff, [which built the] Arpanet. There was a lot of recent modern history in that place, and it had been a

pioneering place for speech and language research at the time, and so the next several years were just an incredible learning experience for me. So that was early. Then I expanded the set of things I was interested in. It led to computer vision, other areas. And all of that just happened to be a good thing for today's world, where our interest is AI. When we talk about it now, we talk about it in terms of multi-modality, reasoning, and things like that. I guess also wanting to constantly work on new problems while still maintaining some connection with old problems allowed me to increase the surface area of what I was doing.

And then I went to the University of Southern California as a faculty member and also as an administrator. I was vice dean in the School of Engineering, I was a faculty member in computer science, but I also was the head of the Information Sciences Institute. Then I went to Amazon, where I headed the Alexa AI organization; [I had] a fantastic learning opportunity there, too. I contributed and learned how to scale — do massive scaling. Then I wanted to go back to my original roots, where I was also building end-to-end solutions for end users in enterprise. And Capital One, now going back to the technology-forward lean — big investment, support from the very top for being at the frontier of technology — all of that just felt like an exciting place to just come and build.

SHERVIN KHODABANDEH: Yep. It's wonderful, and it does. I realize you might not be at liberty to talk about all the magic that's in the works, but are you seeing a future where that composition of the team that does AI is changing and evolving and maybe moving away from the hard-core data science a bit more toward other skill sets like engineering and prompt engineering and design and human-centered design and things like that?

PREM NATARAJAN: I think with every wave of technology, whether it's AI or something else, it's more of a rebalancing of the resources across the skill spectrum. When something new comes about, you need new skill sets in your enterprise. And then maybe it helps improve the productivity of certain other things, but then you need these new things, too. So basically, the overall enterprise is producing more through a

rebalancing of these things — so, you know, people learn new things, etc.

I would say, coming back to the thrust of your question, though, we are opening up a whole new set of possibilities in terms of what can be done, whether [it's] one of the most popular ways in which these are being used — these retrieval-augmented generation style uses. If you look at something like your favorite search tool today that uses generative AI, they're using some form of this. And those things allow us to become better, faster at things that you might do routinely, things that you might not enjoy doing. But when it comes to certain things around decision-making, etc., I think that end of data science still remains, in where you're bringing in your domain expertise to use these technologies to deliver more value in the domain. But I think the fact that these are more scalable, more adaptable, more capable of learning, able to consume massive amounts of context, makes that investment that much more valuable because you get so much more performance out of it.

SAM RANSBOTHAM: When you think about what something costs, if it costs A plus B, and the cost of part A goes down a lot, then the overall A plus B goes down a lot and you can do a lot more of it. Can you give us some examples of the kinds of things you're doing at Capital One? Earlier, you said "magical." What's something magical that you've got going on?

PREM NATARAJAN: What could be magical is things that anticipate my needs or things like that. But leaving aside kind of that speculative future … I was also just going to reflect: You know, was it Arthur C. Clarke who said, "Any technology that's sufficiently advanced feels like magic"? So it's also in that technical, science-fiction context that I was kind of saying "magical." But coming to this other question that you have about how are we doing it, I'll give you one high-level, kind of abstract conceptual thing and then something very specific as well.

At the abstract level, I think we see tremendous potential here to harness all of these advances in AI to deliver better experiences for our customers. Capital One has a whole portfolio of offerings for customers, and so we see a real opportunity to deliver continuously better

experiences for our customers. In that sense, I think what will happen is AI will become more and more central to how we deliver value for our customers, how we run our business, etc.

Now, on a specific example, let me talk about our fraud platform. We rebuilt this fraud platform from the ground up and basically to use ML at the center of that enterprise, and also to make it efficient so that we can make complex real-time ML decisions. Massive amounts of context are being consumed; massive amounts of data are being used. And in order for it to be really useful to our customers, those models have to kind of activate an outcome in the time it takes our customers to swipe their credit cards.

So it's a feat of science, but it's an even more impressive feat of engineering. I like the fraud example because I think it brings together the note of all the different disciplines we need to bring. I think the best work here will be at the intersection of folks with solid product vision who are envisioning the use cases, the folks with the science vision to translate that product vision into saying, "What is the invention that is required to enable that?" and then folks with the engineering heft to say, "I can do all of this. I can it reliably. It will work time in and time out, and it'll work in real time all the time. And you can count on this," etc. So it's just like something that exercises all the muscles of a complex multidisciplinary org.

SAM RANSBOTHAM: I like those three [elements — product, science, and engineering expertise]. I mean, if you think about any one of those three, if you didn't have it, it wouldn't be worth doing. You could have great science and great engineering, but if you've got the wrong idea, you're not going to go anywhere. But having all those three together, as you point out, [is a] big part of that.

Where's the challenge? Which one of those three is the hard one? All of the above?

PREM NATARAJAN: All of the above, but not all of the above in the same proportion in every instance. If you take something like fraud, which is a very highly developed use case in the sense that, just from a use case perspective or a

product perspective, we kind of understand the whole shape of this application well, there, the balance might be in the science and engineering. Like, you have to do, really, "What is the new invention I'm going to do?"

There might be other cases which, you know, obviously, as you might imagine, we're not ready to talk about yet, but where things don't exist, where you're envisioning the future — and this might fall in that bracket of "What's the new magic that will happen in the future?" There, I think all three of these have to be engaged, but there has to be a lot of product thinking upfront: "What's the use case? What's the impact it'll have? Who will benefit from it? What will be the business impact? What will be the customer impact?" So all of that. And it helps for the entire enterprise, over time, to start having more and more of a product mindset so everybody's thinking in that way. So I'd say, Sam, that it's all three, but not always all three in the same proportion in every instance.

I'll give you an example in the case of fraud. If you get deeper into this thing … you know, the world is constantly changing, and so you build a something that models fraud, etc. But you know the world adapts itself to everything. That's the amazing thing with humans: We learn and we say, "Oh, these patterns no longer give me the outcomes I wanted. I have to adapt how I do this stuff" — which means the models also have to be updated regularly. When we first started, it used to take months to update models. We now do it in days.

When we think about from it from a product perspective, sometimes it's about the high-level application fraud. But then if you're really thinking deeply down and say, "What are my developer partners experiencing? What is the friction they're experiencing on a day-to-day basis that might stand in the way of me, as a company, getting maximum benefit from this technology? Oh, look how much effort it takes to update these models. Ah, can I bend that curve?" Well, it turns out some of these new technologies help us bend that curve too because they lend themselves to more easy updating and things like that because they learn so much more effectively.

SHERVIN KHODABANDEH: These are all the ingredients of magic, right — I mean, what you're talking about? What's going on in the industry overall, your own unique background with language being a pretty enduring thread.

PREM NATARAJAN: Yes.

SHERVIN KHODABANDEH: And then, of course, the mix of both the theoretical and the practical and the entrepreneurial, and then the support over there. That does sound like magic. And so I guess the question I have is, give a glimpse of the kinds of things that make it magical. Can you give us a teaser on any of this?

PREM NATARAJAN: I can try. As a lead-in to the answer to that question … look, I think these are still kind of very early days, because generative AI, as I said, has been showing properties that are not entirely designed. The fact that you give it an instruction and it follows it —

SHERVIN KHODABANDEH: Mm-hmm.

PREM NATARAJAN: Things like that. So I think it behooves us to be responsible — especially responsible, thoughtful — in doing some of these things because once something can do things that you haven't designed, you really want to think through things.

So imagine that you're reading … I don't know, you have a 30-page paper to read, right? Well, how about something [where] you say, "Can you please give me a half-page summary of the key points?" Now, part of it might be that half-page summary may miss some of the key things, but just as somebody who wants to keep up to date with a lot of stuff that's going on, right? To me, that feels like a pretty magical thing — like, just as somebody who wants to keep abreast of a lot of stuff that's going on. The next step could be, "Hey, here's a collection of four papers that I've been looking at. Can you give me a summary of what's different between these four?" Oh, now that's really getting exciting, right? Because it's one thing to read a paper in depth and understand what's there, but now to read four papers to kind of zoom in on what's different between them. So I think in those kind of

relatively open use cases are the seeds of the kinds of things that you could do …

SHERVIN KHODABANDEH: That's right.

PREM NATARAJAN: … for other kinds of applications where those kinds of things are particularly-valuable but give you a sense of the power of this technology. Now, in order to do it right, remember how I said, "Can you tell me what's different between these?" I could give it the four papers and say, "Can you summarize these four papers?" and it might actually just focus on what's common between them because it might think that's what's important. Or it might make the statistical inference — let me not say "think"; I don't want to anthropomorphize this too much.

SAM RANSBOTHAM: [Laughs.]

PREM NATARAJAN: But if I give it a specific "Tell me what's different," so now I have to develop some level of understanding of how to instruct this technology to get the output I want. It's doable, and that's where gaining familiarity with these things becomes useful.

SHERVIN KHODABANDEH: Yeah. And that's sort of what I wanted to sort of tease out of you because I also feel like it's a technology, but it's different in many ways, whereas if I think about some of those AI models that you were talking about, you know, with fraud, AI was a tool, and it would dramatically improve the efficacy and speed and accuracy of predictions. It feels like here, generative AI is more than a tool. It maybe is a coworker.

PREM NATARAJAN: Yes.

SHERVIN KHODABANDEH: You're sort of having a conversation with it, and like a new coworker, you're training it and you're seeing unintended consequences, just like you would … I know you're not trying to anthropomorph- … anthro- … OK.

PREM NATARAJAN: … -pomorphize it, yeah.

SHERVIN KHODABANDEH: But I think there's something very, very unique here.

PREM NATARAJAN: Yeah. There is something unique. It's something uniquely exciting.

SAM RANSBOTHAM: Let's transition. Prem, we have a segment where we'll ask you some rapid-fire questions, so just give us the first thing that comes up in your mind. What's the biggest misconception that people have about AI right now?

PREM NATARAJAN: The notion that AGI [artificial general intelligence] is around the corner. And I think it's a misconception that works at two levels. One is that there is the belief that there's a shared understanding of what is AGI.

SAM RANSBOTHAM: Mm-hmm.

PREM NATARAJAN: And kind of the fear around it … I think it's appropriate to be very guarded and very concerned and very thoughtful about how, as a society, we should respond to it, etc., but I think as humans, we've overcome so many things. So I just feel like we will prevail here, too.

SAM RANSBOTHAM: Prem's talking about artificial general intelligence here versus the more narrow definition of AI that shows up in so many of the things that we're doing.

PREM NATARAJAN: Yeah.

SAM RANSBOTHAM: What do you see as the biggest opportunity for AI right now?

PREM NATARAJAN: The biggest opportunity for AI is to, in my mind, democratize access to a lot of services resources, etc., across the entire spread of our social spectrum.

SAM RANSBOTHAM: What was the first career that you wanted?

PREM NATARAJAN: Oh, what was the first career that I wanted? That's interesting, right? It was actually a lawyer.

SAM RANSBOTHAM: That's consistent with language.

PREM NATARAJAN: Consistent with language. I also felt like many of the role models at that stage in my life that I looked at, especially in the Indian context where I grew up around the time of the independence movement, a disproportionate number of them were lawyers or teachers.

SAM RANSBOTHAM: So where are we using too much AI? Where are we making this hammer fit all the screws?

PREM NATARAJAN: I don't know that a particular pattern stands out to me, but I'll say this: I think when something is working really, really well, like, for example, my tap: You know, if it's working well, and then I can do a touch — you know, nowadays you have these touch taps — that's awesome. I think if you get to a point we're saying, "Tap, turn on water tap," that feels like, you know … So I'd say there are things where I think as humans, we can say, "Is AI improving, reducing the friction I'm experiencing in my life, making things easier, or does it just feel like, you know, strange?"

SHERVIN KHODABANDEH: Technology for the sake of technology.

PREM NATARAJAN: Yeah.

SAM RANSBOTHAM: So, what's one thing you wish AI could do now that it cannot?

PREM NATARAJAN: Oh, that's an easy one. I wish it could make me an awesome singer. I love singing; I love music. I don't have a singing voice, and so if AI could transform —

SHERVIN KHODABANDEH: I think it can.

SAM RANSBOTHAM: Yeah, I think we might be there.

PREM NATARAJAN: I would like to have an AI attachment to me that I go out on a karaoke and I sing, and everybody's like, "Man, this guy is belting it out."

SAM RANSBOTHAM: [Laughs.] Good. I really think this framework that you mentioned about

the idea of product combining with science combining with engineering and how all those pieces fall together and are necessary but have balance, different balance in different situations, that alone will probably resonate with lots of our listeners. Thank you for taking the time to talk with us. We've enjoyed having you. Thanks.

PREM NATARAJAN: Thank you.

SAM RANSBOTHAM: Thanks for listening. On our next episode — our final episode of Season 8 — Shervin and I chat with Mark Surman, executive director of the Mozilla Foundation. I'm excited about this one. Please join us.

ALLISON RYDER: Thanks for listening to *Me, Myself, and AI*. We believe, like you, that the conversation about AI implementation doesn't start and stop with this podcast. That's why we've created a group on LinkedIn specifically for listeners like you. It's called AI for Leaders, and if you join us, you can chat with show creators and hosts, ask your own questions, share your insights, and gain access to valuable resources about AI implementation from *MIT SMR* and **BCG**. You can access it by visiting mitsmr.com/AIforLeaders. We'll put that link in the show notes, and we hope to see you there.