



The Deloitte On Cloud Podcast

Gary Arora, Chief Architect of Cloud and AI Solutions at Deloitte

Title: **Balancing healing and technology: Northwell's Dr. Sanjeev Suratwala talks AI's potential**

Description: In this episode, Gary Arora and Northwell Health's Dr. Sanjeev Suratwala examine how AI is transforming health care by aligning innovation, clinical care, and business needs. They discuss AI's impact on workflows and efficiency, highlight the importance of governance, and address breaking down barriers like tech debt and bias. Gary and Dr. Suratwala also discuss how AI can help democratize expertise, build trust, and make care more equitable and accessible for everyone

Duration: **00:19:03**

Gary Arora:

Welcome back to On Cloud. I'm your host, Gary Arora, chief architect for cloud and AI solutions at Deloitte. I'm really excited about today's episode, because if you truly want to understand both the promise and the pressure of AI, you have to look at systems that simply can't afford to get it wrong. And nowhere that is more true than in health care, where lives are on the line. My guest today is Dr. Sanjeev Suratwala, an orthopedic spine surgeon and a member of the AI Scribe Governance Group at Northwell Health. It's one of the largest health care providers in the country. He is on the front lines of both medicine and AI adoption, helping shape how AI is governed, adopted, and trusted across clinical workflows. Sanjeev, thank you so much for joining the show.

Dr. Sanjeev Suratwala:

Gary, thanks so much for having me. I'm excited to dive in. This is a conversation that's both timely and personal for me.

Gary Arora:

So, let's get started. You operate in a world where healing patients and managing a business coexist, often uneasily. How do you see AI changing that balance between care and operational efficiencies, and do you really think technology can make both better at once?

Dr. Sanjeev Suratwala:

That tension between healing and business, every clinician feels that. On one side, you've got the mission to heal, to connect, to bring humanity into medicine; on the other side, you've got the system with budgets, margins, and workflows that need to keep the lights on. It's not a bad thing. It's just the reality that modern medicine operates within, but the friction between the two often leaves both patients and providers frustrated. The beauty of AI is that it might finally help bridge those worlds instead of forcing us to choose. A lot of the inefficiency in health care comes from waste and variability. People sometimes do the same thing 10 different ways. AI can standardize that by putting evidence-based practices right at everyone's fingertips, whether you're in a major hospital or a small clinic.

Think about diagnostics. An AI model trained on millions of imaging studies can flag subtle fractures or early signs of spinal degeneration, in my case, that a busy clinician might overlook on a hectic day. But now imagine that same capability available even in a lower-resource setting where there's no subspecialist. Suddenly a nurse or a clinician in a rural clinic without access to doctors can access the same level of diagnostic insight as a tertiary hospital. That's democratization of expertise, and it can transform how we deliver care across America and the world.

On the business side, AI helps predict and prevent problems before they snowball or bottlenecks, readmissions, staffing shortages, all the hidden inefficiencies that drain resources. Predictive models can anticipate these things that cost health systems millions, but don't really add value to patient care. By aligning these operational efficiencies with clinical outcomes, you start to see a model where the business of care supports the practice of care, and if you zoom out, that efficiency creates capacity when documentation is automated, when workflows are streamlined, clinicians can see more patients, including those who might not otherwise have access.

Telehealth, and now with AI-driven triage, could allow one specialist to extend their reach across an entire region, maintaining quality while scaling access. So, to me, AI isn't about choosing between healing and business. It's about harmonizing them, using intelligence to create a system where doing what's best for the patient is also what's most efficient for the system. When compassion and efficiency stop competing, everybody wins: the clinician, the organization, and most importantly, the patient who might otherwise fall through the cracks. That's where I see AI's real power, not in just in making medicine faster or cheaper, but in making it fairer.

Gary Arora:

No, I like how you put it. The balance that you're describing between the mission and the technology and the model, it feels like something every technology leader is wrestling with. I also liked your example of the diagnostic insights that can now be available without the high barriers to entry for things like resources and specialists, which not all hospitals may afford to have. That is the democratization of health care, as you put so well. This leads me to my next big question. There is so much hype about AI coming in the medicine space and doing crazy things. From where you sit in surgery and AI governance, what do you see actually working, what still feels like marketing, and what is real?

Dr. Sanjeev Suratwala:

We're definitely in that middle ground between hype and healing. There's a ton of excitement, and honestly, a lot of noise as well with this. What's working right now are very narrow practical tools, things that solve real pain points. AI scribes, for example, that take care of documentation, imaging algorithms that flag abnormalities, predictive models that spot post-operative complications early. Those make a measurable difference in real workflows. Where it still feels like marketing is when people start talking about AI replacing doctors or diagnosing everything overnight.

Medicine just doesn't work that way. There is context, there is judgment, there is empathy. That's still a very human quality. You can't feed a model a chest x-ray and expect it to understand the patient's life, goals, aspirations, and comorbidities. We can't meet patient expectations through just AI models. So, the best AI in health care is actually invisible. It fades into the background, doing the repetitive stuff so we can get back to what really matters, listening to patients and practicing medicine.

Gary Arora:

The best AI in health care is actually invisible. I really like that. You also mentioned scribes. Let's dig deeper there. When we talk about AI scribes or clinical assistants, it sounds like automation, but it's really augmentation, isn't it? So, what does a human-in-the-loop look like in a surgical or clinical setting, and what can we learn from health care into enterprises when teams are designing their AI assistant and AI copilots?

Dr. Sanjeev Suratwala: So, human-in-the-loop sounds pretty technical, but when you live in it, it's actually pretty simple. AI isn't really replacing the human; it's extending what we can do. When I'm in the OR, I'm constantly using technology that amplifies my precision: robotic systems, navigation tools, imaging guidance, but ultimately, it's my judgment that makes sense of all that information. AI is just becoming another one of those instruments.

We talked a little bit about the AI scribes. I'm not typing notes anymore. The AI scribe listens, drafts notes in real time, but I'm still reviewing, editing, and making sure that the story reflects the patient's reality. That's not really automation, that's collaboration. The machine's doing the repetitive work for me so I can focus on the nuance of taking care of patients, the empathy and decision-making that's still a very human quality that patients look for in health care. So, AI really shines when humans stay in control. It's not there to think for us, it's there to help us think better.

Gary Arora:

So, you help oversee AI governance at your hospital, and that's a model many enterprises are trying to build. What does an effective AI governance look like in such a high-stake environment, and how do you keep it from turning into a bureaucracy that would then just slow down any innovation?

Dr. Sanjeev Suratwala:

Well, that's a challenge. In health care, you can't just move fast and break things. The stakes are human lives, but if you move too slowly, you also risk denying patients the benefit of new technology. So, good governance is about enabling innovation safely, and at the same time not blocking it. For me, it starts really with clarity. Everyone, from data scientists to clinicians, should know who's responsible for what, how the model was trained, and what the validation steps were.

Then comes accountability. If something goes wrong, we need to trace it, not shrug and blame the algorithm. We should be able to trace it back to a decision, not some black box number. Finally, agility. Governance can't be static. We run pilots, gather real-world data, and adapt quickly, always keeping patient safety at the center. That's really how you balance innovation with responsibility.

Gary Arora: Clarity and accountability, tracing back the decisions to not just a black box. I mean, it's so relevant for enterprises, something enterprises are grappling with. Another thing that enterprises are grappling with is the tech debt and data silos. Every CTO knows the pain of legacy systems and fragmented data. I feel in health care, the stakes are higher. A lot of the providers historically have grown over multiple M&As and acquisitions. So, they just bring a whole lot of baggage from different systems that may not be talking to each other.

How big is this barrier in terms of data quality and interoperability? How big is that barrier to meaningful AI adoption? The only reason I'm asking is to see what you have figured out here that we can adopt in other industries outside of health care.

Dr. Sanjeev Suratwala:

That's a very salient point because even in my current health care system, we've grown by acquisition, and with acquisition comes legacy systems. So, there's always a challenge, eventually, of integrating and making everything work together. So, it's huge. This is probably the biggest one. Health care data lives everywhere, one system for labs, another for imaging, another for notes, another for billing, and they rarely talk to each other cleanly. There's also a challenge of connecting across multiple hospitals where patients might be getting care.

So, the problem really isn't just the tech, it's the mindset. Every organization wants to own its data when what we really need to do is share the stewardship of this data. If a patient's record could follow them seamlessly from one hospital to another, we would unlock an entirely new level of care, research, and AI training that can make everything better. The silver lining is that AI might actually help us clean up this mess. We're making progress using standardized frameworks across various health systems so that the data can talk across different systems.

So, ironically, while the poor data quality limits the AI, AI might actually be the tool that helps us overcome this limitation. It can harmonize messy data at a scale. So, the same technology that struggles because of poor data might also be the one that helps fix it, and there is cost associated with that. So, we're struggling with that as well. How to incorporate these and integrate these legacy systems and move on to newer technology and newer standards. Hopefully, we can use the AI tools and standardized frameworks to make this possible.

Gary Arora: You brought up costs that are associated with such a transformation, and that's a very valid point because this is not just a technical challenge, it's a financial one. It brings me to the economies of all of this. Health care runs on razor-thin margins. So, the innovation here has to earn its keep. How do you think about ROI for AI investments, not just in cost savings, but in outcomes and time and trust?

Dr. Sanjeev Suratwala:

Health care operates in its own little bubble, unfortunately, and sometimes it's a positive and sometimes it's not the best, but ROI in health care is really different from most other industries. It's not just about dollars, it's about value. If AI gives a clinician an extra hour a day, that's a real ROI. That's time for another patient or time to breathe between cases, and that time compounds. There's less burnout, more focus, better outcomes. So, those are a lot of intangible things that we can't quantify under regular ROI measurement on a spreadsheet.

Then, there's the outcomes piece. You get fewer complications, quicker recoveries, more consistent care. That's value that ripples across the whole health system and frankly benefits the patient who's at the center of health care, and there's also trust. That's the hardest to quantify, but maybe it's the most important. If AI systems make clinicians feel supported and patients feel safer, they'll actually get used to these systems, and that's when ROI becomes sustainable.

Patients may actually seek care knowing that there's no ulterior motive and everybody is using standardized care across the whole country. So, for me, ROI is about what AI gives back. There's time, there's clarity, there's confidence in the system. When those improve, the business follows naturally. So, a progressive system recognizes these non-tangible benefits that AI can provide.

Gary Arora:

Non-tangible benefits, I like that. Trust and ROI are interesting because they directly connect to who gets served and who doesn't. As we expand access, we also have to confront the question of equity, not just who gets care, but how fairly that care is delivered. Let's talk about that. This is a big topic. Bias in data can translate directly into bias in care. So, when you're building AI systems trained on historical data and historical patterns, you might be designing systems that are replicating some of the inefficiencies. So, how would you and your governance team approach the fairness question in AI so that you can close the gaps in health care delivery? I know this is a topic that you're passionate about.

Dr. Sanjeev Suratwala:

That's one of the most important conversations in the AI right now. At the technical level, bias isn't just a concept, it can literally decide on the practical side who gets diagnosed, who gets care first. Our approach is multi-layered. First, it starts with representation, making sure that the data we train on actually reflects the patients we're serving. Using historical data from one community doesn't really serve or represent what see in New York.

Then, we continuously audit this model after deployment because bias can creep in over time as populations change. But the most powerful thing AI can do is highlight inequity. For instance, an algorithm might reveal that certain groups are consistently being referred later or receiving different treatment algorithms. That insight can drive interventions and policy change at the system level. So, yes, AI can carry bias, but it can also actually detect it and expose it and help us close all these gaps that have been going on for decades.

Gary Arora:

One of the biggest promises of AI is democratizing access, making quality care available regardless of geography or income. From where you sit, given the kind of tools that are available today, how realistic is this vision? Are there any early signs that you are looking at that make you feel more optimistic?

Dr. Sanjeev Suratwala:

It's already happening. Maybe not everywhere yet, but we're seeing real examples. AI can extend the reach of expertise. Imagine a small clinic in a rural area using AI-assisted imaging or triaging tools to get guidance from a tertiary center instantly. With AI-enabled imaging or triage tools, that clinic can now detect a condition early, consult virtually, and connect a person to the right care, faster. That's literally life-changing. There are AI tools right now reading chest x-rays for tuberculosis in parts of Africa or screening for diabetic retinopathy in India.

So, these are, again, life-changing applications in places where specialists are scarce. Here in the US, we're seeing telehealth combined with AI scribes and triage models that allow one specialist to reach patients across entire regions. That's huge for access. We need to make sure these tools are affordable and teach people how to use them. If we get that right, AI really could level the playing field in health care, especially for patients trying to get access to quality care.

Gary Arora:

Finally, if you look ahead, beyond the pilots and prototypes, I'm curious how you imagine this coming together. So, if you were to fast forward a couple of years, what would responsibly AI-powered hospital look like? What kind of mindsets or governance structures will separate the organizations that use AI wisely from those that just use it widely?

Dr. Sanjeev Suratwala:

I like the "wisely versus widely" because, I that's really the key. I think that the best AI hospital in the future won't feel like AI hospitals. We talked about technology being invisible. The technology will be everywhere, but in the background, quietly managing workflows, predicting needs, flagging risks, so that humans can actually focus entirely on patients. You'll see real-time decision support, documentation that happens automatically, analytics that guide everything from staffing to logistics.

But the real differentiator won't be how much AI they use; it's really how wisely they use it. The responsible hospitals will treat AI governance like patient safety: as a living, breathing discipline that's constantly monitored and improved. So, the future hospital won't necessarily be high-tech just for the sake of it, it'll be high-touch, powered by smart tech and deploying AI wisely.

Gary Arora:

Well, that's it for the episode. Thank you, Sanjeev, for sharing your lived insights. What I really enjoyed about this conversation is that health care as a domain forces you to ask the hardest questions about technology, not just what we can do, but should we, and how should we, and who benefits when we do.

Thanks for listening to On Cloud. If you enjoyed this episode, share it with your team and fellow leaders who are thinking about how to deploy AI responsibly, and stay tuned for more conversations. This is Gary Arora until next time. Yes, Sanjeev.

Dr. Sanjeev Suratwala:

Thanks for having me. It's been a great conversation. I would like to actually thank my sons, Jaden and Vian, for encouraging me in this podcast. They wanted to make sure their names were on the podcast somewhere. I was like, alright, I'll give a shout out to them at the end.

Operator:

This podcast is produced by Deloitte. The views and opinions expressed by podcast speakers and guests are solely their own and do not reflect the opinions of Deloitte. This podcast provides general information only and is not intended to constitute advice or services of any kind. For additional information about Deloitte, go to Deloitte.com/about.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor.

Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Visit the On Cloud library
www.deloitte.com/us/cloud-podcast

About Deloitte

As used in this podcast, "Deloitte" means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms. Copyright © 2025 Deloitte Development LLC. All rights reserved.