

AI-GR Pod 41 04.10.26 Byron Crowe

[00:00:00] Airplanes are not flying cars, right? Airplanes are not flying cars. Even though airplanes came 20 years after automobiles. Right? 1903, Wright Flyer. 1885, Karl Benz invents the first automobile. And they look very similar at first, you know? Oh, forms of transportation. Oh, they both have engines. They require a trained operator.

They're economically transformational. They are harmful if they're misused. So, you might say that on the surface, oh, airplanes are just flying cars. Right? It's just a derivative of the old innovation. And, actually, that would be a huge category error. If we tried to control aviation the way we control automobiles, it would lead to broad dysfunction.

So, I think we're, we've been treating LLMs a bit, like flying cars, right? Just, you know, more software, right? So, just another consumer application that I download on the app store, and I log in on Google. But they're fundamentally different. Their capability is, is fundamentally different.[00:01:00]

Hi, and welcome to another episode of *NEJM AI Grand Rounds*. I'm Raj Manrai. I'm here with my co-host, Andy Beam, and today we're delighted to bring you our conversation with Dr. Byron Crowe. Byron is the Chief Medical Officer at Doctronic. He's been working in digital health for several years and he's done some interesting academic work, as well.

And he really took us back from his journey to becoming a physician, then, uh, launching an academic career, and his transition into being CMO of a digital health company. We asked him some tough questions about how you put guardrails on large language models. Where the sort of frontier is with human machine collaboration and even getting into autonomous AI that's able to do more than what we previously imagined.

Um, Andy, it was really interesting to hear his perspective and, also, I think, get some of his historically grounded takes on where we're going in the future as well. Yeah, there were a couple things that jumped out [00:02:00] at me in this interview. One is that we often talk about artificial intelligence and a previous guest we had on the way back, um, talked about autonomous AI or self-driving AI.

And I really think what Doctronic has been doing is much more in that autonomous medical AI where, you'll hear about this in the episode, but you

can get automatic prescription renewals by interfacing with their AI system. And again, like automating some medical care that can be automated today. And like you said, Raj, they've been at the forefront of two areas that often make folks nervous: IT security and having a thing that's used publicly by patients, and also at this sort of uncanny valley of regulation.

So, they operate in the state of Utah. Because the state of Utah allows them to do things like this and really kind of pushing forward regulatory thinking on how AI can be used in health care. So again, touching lots of things that we often don't get to touch on on this podcast. We tend to be LLM-centric a lot of the times, but this, I think was actually nice at pulling together lots of these [00:03:00] things that have to be pulled together to deliver patient care and to really innovate in a space that's hard to innovate in.

The *NEJM AI Grand Rounds* Podcast is brought to you by Microsoft, Viz.ai, Lyric, and Elevance Health. We thank them for their support.

And with that, we bring you our conversation with Byron Crowe on *AI Grand Rounds*. Byron, welcome to the podcast. We're excited to have you today. Andy, it's great to be here. Byron, let me also echo Andy, great to have you on *AI Grand Rounds*. So, this is a question that we always get started with. Could you tell us about the training procedure for your own neural network?

How did you get interested in artificial intelligence? And what data and experiences led you to where you are today? Wow. What was, what's my training database? So. What's your data? So, um, Raj, I, when I was coming up, I was really trying to think about this career in [00:04:00] medicine and trying to understand my own motivations.

And something that I learned very early on was that I think people go into medicine for two big reasons. They like science and they wanna help people. And that was true for me. But what I learned very quickly was I really liked systems thinking. Actually, I read Atul Gawande's book *Checklist Manifesto* very early on in my career and that, that was one of the tipping points for me to say, okay, I want to go invest in systems level change at scale.

And that's been the path for me for almost two decades now. I work in technology today, right? I work in digital health. But I see technology as a lever to improve systems. Some of my first work was really in quality and safety, and when you ask, well, all right, what's, what's your training data? It's that systems level mindset that really permeates how I view medicine, how I view the role of technology, how I view the role of humanity.

It, it all works together as part of one big system. And there's not just in these individual leavers [00:05:00] that we need to pull, right? It's not just one parameter, right? All the parameters have to work together to make this whole thing go. Do you have, like, an early memory of when you first got excited about medicine?

Was someone in your family, a doctor? Do you remember going to the doctors? Like, do you remember when you knew that maybe a career as being a physician is the right training path for you? Yeah. You know, I, I, I can remember very vividly, my favorite class in high school was AP Biology. Nice. I felt like it was this illumination into how biological systems worked and, you know, again, right, like science, you wanna help people.

That was the liking science part. Like, oh, I really like this topic, and I can see it in context. I was a biochemistry major as an undergrad, but I was also a history major, and I really came to enjoy and appreciate the intersection of biological systems and human systems and what that means. And medicine sits right at that intersection.

That's what we do all day, every day. And going through medicine, medical training [00:06:00] and beyond. I have only been reaffirmed in my choice year, year over year, over year. Awesome. So, uh, maybe we can start to go into sort of the next topic that we wanna talk about with, uh, some of your recent work. But maybe you can take us through your training arc just a little bit further.

So, you, you completed medical school, you did residency, where were you a resident? And this was in internal medicine, right? That's right. I did internal medicine training. I trained out in Colorado. I moved to Colorado. I thought I was gonna live in Colorado for the rest of my life. University of Colorado has an amazing program, and that health system has done some really transformational work in advancing systems level thinking.

It was a big reason why I wanted to go out there and do my training. But while I was in Colorado, I became involved in digital health, and so I actually held a dual role while still in training. I was seeing patients, but I also worked at a digital health company as their associate medical director. And I did that for two years and really came to see, yeah, this [00:07:00] combination of direct frontline patient care and being able to make an impact through the application of technology was really what I wanted to lean into coming outta training.

After you finished up Colorado, take us to the next step. Yeah, so then I moved to Boston. Yeah. Uh, moved to Beth Israel and HMS. I started as an early career

faculty member, and in that role, that's where I had the pleasure of working with some incredible folks, some of the folks that we know very well.

Right. Dr. Adam Rodman of AI fame. And getting to know Adam and doing some of our early work, exploring the role of LLMs in clinical reasoning became the core focus of my research when I was in Boston. And really seeing the power of this new class of technology, what it was capable of, and projecting forward.

Okay, it's here now, right, GPT-3.5, kind of trying its best, but here's where it's gonna get to. And [00:08:00] this is something that really is gonna change medicine. That's been the story of the last three and a half years for me. Could I ask a question? We probably will come back to this, but um, given that you were like, uh, into digital health, into digital medicine, very forward-looking at early stages of your training, are you surprised by what's happening now?

I guess is the question like, do, given that you were kind of at the tip of the spear with digital health. I'm saying this 'cause we had Jonathan Chen on the podcast a while ago. Yeah. And he had the, the fundamental law of bioinformatics, which is AI plus MD is greater than either one of those. Um, I wonder if your like sense of what's possible with technology is different or consistent with what you thought it was at the beginning of your training.

Well, first off, Jonathan I also count among my friends and I love that what was it called? The central thesis of bioinformatics. The fundamental, the, the fundamental theorem of biomedical informatics or medical informatics, I think is what it is. Yeah. Is it a theorem or is it a, is it like a dilemma or a conjecture or—? It [00:09:00] sounds like a conjecture, but I think it's often stated as a theorem.

Nice. The nat—the natural law of bioinformatics. Yeah. Yes. Yeah. Yeah. Um, well, I, I, look, I think Jonathan's right and from my side, and this is not a critique of Jonathan, I see the world through a slightly different lens. And it's a lens that is informed by, honestly, more by my experience in history and the liberal arts than by pure technology.

And that's looking for and recognizing step changes in technology. And understanding and seeing what happens after that technology is brought into the world. With language models, you know, even if you go back to 2017 and say, wow, right, we have this new transformer technology, this is amazing. It has high potential.

Being able to fully foresee what would happen in the next decade, I think that's hard because it's unknown, right? That's an R and D question. It's not an engineering question. So, the pace of R and [00:10:00] D is often unknown, unknowable, but I think what happens after the technology is brought to light, recognizing, oh wow, a shift has occurred, and then adapting to it.

That becomes a really important framework that I like to use as I think about where I'm focusing my attention. Ten years ago, it was digital health. We're moving from really analog health systems to now we can do more with applications on a mobile device. With telemedicine, now we have medical AI. What can we do with this?

These new types of cognitive systems that we can point at some really important problems that we've never had the capability to do before other than the human mind. Yeah. Thanks for that. Maybe back to our regularly scheduled programming, Raj. Yeah. Sorry for the interruption. So, Byron, you took us up to a few years ago when you joined as a faculty member at Harvard Medical School and at the Beth Israel.

And you published this great paper with Adam and with Zahir [00:11:00] Kanjee on using I think GPT-4 for trying to solve these very hard cases in the *New England Journal of Medicine*, the Clinical Pathological Conferences. We've spoken with Adam about that paper. We've, we've talked about it a few times in some of the other episodes, and that was the first time that I think I got to know you 'cause I saw that paper and I was like, super cool.

I need to meet these three doctors and, uh, learn about what compelled them to do this. And then also to tell them—uh, and I should, I should tell you this as one of the co-authors, that that was awesome that you guys did that so fast and so early, and you know, now it's become an area that a lot of folks have focused on.

And not only using the CPCs, but you've inspired a lot of good work by very early on, uh, doing something like that and publishing that paper in *JAMA* couple years ago. So, you're having a great time, right? You're doing well. You're publishing good papers. You're sort of being a clinician, an internist at the B.I., what compelled you to leave that nice academic job that you had and go venture off onto, onto other pursuits?

Well, [00:12:00] well, Raj, uh, I'll just add one layer to that, which was alongside that work, I was also holding a role in industry. I was the chief medical officer at a different digital health company, and we were building out a

national network of chronic disease management tools that we brought to big insurance companies and employers, and they could make available to their members as a free benefit.

And that was really a great place to show the impact of scale. Right? We're putting, you know, we had, we had these programs accessible to tens of millions of people nationwide. And that was all well and good, but what I was learning and what we were all, what we were all learning. Right. What I, what I was seeing in my research with Adam, what you were seeing in your research, what a lot of us were seeing is the, that these language models were advancing, the capabilities were strong and we were moving toward readiness.

And again, this idea of a technological shift. Right. The shift is underway.

[00:13:00] It's this very dynamic period. And so, when the Doctronic opportunity was presented to me, you know, again, getting to know the team and seeing what the vision was, uh, it became very clear to me that this is an organization that can take the capabilities of medical AI and bring it into the real world in the way that I thought was the right way and, and still think is the right way.

And start to fulfill the promise of this technology. I think that's a great point to transition to, uh, what we wanted to talk about. Um, first we're sort of digging into some of your recent work. So, you're at Doctronic now, you're the Chief Medical Officer of Doctronic.

Maybe we could start off, could you just tell us what the sort of mission is, the vision of Doctronic, and then we've all been reading about this recent partnership with the State of Utah. So, maybe you can also tell us a little bit about that and how it got started and what you're, trying to show.

Yeah, happy to. So, Doctronic is a nationwide care delivery organization. We use AI to support care delivery in all 50 [00:14:00] states, and then we're building medical AI for the world. The care model we practice is called AI Native Care Delivery. It has four parts: AI Chat, which is general health Q and A, right? Come in, ask a general health question. AI Consult, which is an end-to-end patient encounter conducted by the AI that produces a progress note at the end, right?

It's more than intake. It takes a history and then reasons over that answer to produce a working differential end plan that then we use for our clinicians to see patients. That's the third layer, right? Human care either delivered virtually, which is what we do, or in person. You can take that progress note to an in-

person clinician and say, hey, here's, here's the results of my AI consult as a form of clinical decision support. The fourth layer of AI Native Care Delivery and the one that is the newest is Autonomous Care. Today, Autonomous Care is only available in the state of [00:15:00] Utah as part of a very specific program being undertaken, under state level practice and medicine authority.

But with time, Autonomous Care is the natural extension of these capabilities. What else can we do using these new systems that no longer require a human immediately present, making that decision but can be used by a technology system that has huge opportunity and is also complex and challenging.

And we're starting today in the state of Utah with a very specific program, refilling medicines with AI. Can you tell us a little bit more about that program in Utah? Yeah. How, how did you select the sort of population? How big is the program? Can I actually ask a, a much more basic question?

Yeah. Assume this is totally not the case, but just assume that I actually don't understand how prescriptions are refilled. Like what is the, like what is the challenge in automating this and what's the normal criteria for a prescription refill outside of AI? Yeah, [00:16:00] so Andy, here's how this works in most places today.

Most places today, if you need a refill of your medicine, right? You, you, you get a medicine prescribed for, let's, a year at a time. Mm-hmm. And at the end of that year, you need to renew your prescription, which means you have to ask your doctor for more. And there's a few ways to do that. You can call the office and talk to the front desk.

You can send a message through your patient portal. Sometimes people go to the office and talk to their doctor directly. But this is a high volume, high throughput activity, and even in a medium-sized health system, there's millions of these refill requests every year. And in order to process these with some level of efficiency and without overwhelming the health care system, these refills are protocolized.

You need a minimum amount of information to say, yep, this refill is appropriate, and then we can move on. Usually, it's done through this interaction with the portal and [00:17:00] the physician at the end of the day is going through their messages and saying, yep, refill, refill, refill. This adds up though, right?

The average primary care doctor may spend, depending on what numbers you look at, 10 to 30 minutes a day managing refill requests and at scale that becomes operationally burdensome and, and frustrating for patients, right? We've all encountered somebody who ran out. I ran outta my medicine. I couldn't get in touch with my doctor. So, it's a pain point for all stakeholders.

The challenges with refilling these medicines or renewing these prescriptions, those challenges are known, right? You make a protocol and, and install it in a clinic. But using an AI system to now participate in that decision where you can actually interact with a patient, you can ask 'em questions, you can take a medication history from them.

You can combine that with deterministic checks about what we know about their medication history and their medical history. Um, that actually creates an entirely new foundation upon which we can engage with patients. In something that [00:18:00] used to just be a very protocolized standard engagement without high levels of scrutiny, right?

You know, you're sending a portal message, usually not interacting with the doctor, and then the refill pops back into the pharmacy. Even if the use case isn't the most complex in the world the foundational infrastructure is entirely new to be able to do this. Yeah. Um, thanks to that, and I obviously was aware of all of that background, so thanks, uh, for explaining that to folks, uh, who weren't. It does seem like this natural place to interact with the patient.

I'm sure we will talk about this, or if you might already be thinking about this, but to collect other kinds of information? It's a place where the patient is activated to engage with the health care system because there's a need that they have, and it does seem like a very natural insertion point for AI as a result of that.

Raj, I don't know if you wanna pick up here with some of the other things that they're, um— Yeah. —building around this sort of like—. Yeah, totally. Uh, I mean, maybe Byron, you could just tell us a little bit more about the pilot that you're running with Utah. So, give us a sense of sort of how big [00:19:00] it is, how you selected patients, like what types of, are you narrowing the conditions under which patients can be sort of enrolled in the pilot?

Uh, how did you design it? And maybe just tell us a little bit more about it. Well, the first thing to recognize about the state of Utah is that this is set up as part of the state government's regulatory purview in something called a sandbox. And what that means is the state says, hey, if you want to try

something that would normally require a professional license, which obviously practicing medicine requires a professional license, we will suspend that requirement so you can try something under our supervision.

The sandbox approach. That's what we're doing in Utah. There are other states that are also exploring state sandboxes and sandboxes have been used for other technologies to bring them forward. And Utah had the vision for this, right? The key thing is the folks in the Utah state government, they [00:20:00] saw the potential of medical AI and said, let's create the space to bring this forward.

Then working with the state, now there's a scoping question. Right. The, the wrong move is to go in and say, all right, well now we have this opportunity to use AI. Let's pick the most complex, high consequence use case, right? That's not how any technology is brought into the world in a graduated way. There's a lot of similarities between this and, and self-driving cars, this is, all right, let's, start in the parking lot and, there's a human at the wheel.

Right? So, with— Or interestingly for self-driving cars, right? Like I think a lot of the— Yeah. —rollout of autonomy was on the highway. 'Cause, like, the highway is, like, so patterned and predictable. Right? And then the sort of final thing that Tesla users got was navigating these, sort of, weird city streets where there's not clean lines and hard turns.

Like think about the Boston cow paths, right? And, and, uh, that turn right by MGH and things like that. I was envisioning it in my head [00:21:00] as we speak the— Yeah, exactly. —the seven, the seven lanes colliding. I used to live right down there, so I know it too well. But exactly, like, and if you even go a little further back with autonomous cars, we start with the driver in the seat.

We don't start with no driver. And that's exactly what's happening in Utah. Every one of these—. So, so their supervision means that they have a licensed medical professional on their side. Who's overseeing this? Like what does that human supervision look like? Well, we have a nationwide clinic. Our doctors at Doctronic review every prescription in the first phase.

Okay. Even, even the ones that are coming sort of autonomously out of the AI system, there's a Doctronic doctor who's looking it over. That's exactly right. Got it. The Doctronic doctor reviews the prescription and if there's any concern they can—. They can intervene. They can intervene and they can request that the patient have a video visit with them.

Got it. Uh, okay. Very, very helpful. And then my understanding from reading some of the, the articles about this is that [00:22:00] you started with only renewals, not new prescriptions. Is that right? That's exactly right. These are patients who, their doctor has prescribed them the medicine and now they need to renew it, but it's not a de novo start.

And that's an important pause point because these medicines, there's a difference between the risk of starting a new medicine and the risk of renewing a medicine. Right. Let's take, let's take something that is on our Utah formulary, right? Apixaban. Um, Apixaban's a blood thinner. If you were prescribed Apixaban, you have a risk of bleeding.

It can cause a death or serious harm. At the same time, it has a real benefit, right? We can prevent strokes or other blood clots. But that decision of whether to start an anticoagulant is very different than whether to renew the prescription for an existing anticoagulant. Right? Right. Where you don't want people to have a gap in their medicine.

Actually, the harm is by, uh, [00:23:00] by a barrier to renewing it. Recognizing that it, it can be nuanced, but recognizing that distinction is really important and an important part of our choice to say we're starting with renewals, not de novo prescriptions. That totally makes sense to me. I, um, I, I guess maybe my, my following question is, are there certain drugs or medications, prescriptions, or even certain conditions that the patients have that you think are ineligible for this program?

Like, something that is either higher risk or even for the renewal, you would want the patient to have a conversation with a human physician before renewing the medication. Cause like, I'm, I'm thinking about just to sort of, even within renewals, if there's a spectrum of risk from your perspective.

Well, to that point, let's first talk about some of the things that aren't on the list. Controlled substances are not part of this program. Antibiotics are not part of this program. Uh, even some other categories like injectables, right? Yeah. Not part of this program. [00:24:00] And so, from that standpoint, we carved out a whole category where there is a higher level of consequence and scrutiny, and we moved that to a future state.

Right. And this is along the lines of what you're saying too, which is sort of, you're very carefully trying to choose what you start with from the perspective of where the sort of risk is. Exactly. And then from there, yes, there's within any one of these medicines. On this formulary of over 190 medicines.

Within any of these, there are triggers and escalation to speak with a doctor. In this case, it would be one of our doctors, right? Do that video visit and clarify accordingly. Got it. And then how does a patient get enrolled or become eligible to use the service within the sandbox? Like, if they're on one of these medications, do they get a notice?

Uh, if they're randomly selected, like how do they become sort of, I dunno if the right word is enrolled, but how do they become eligible to be part of this pilot in Utah? We, this is gonna sound wild, we tell them [00:25:00] about it. Okay. So, anyone in Utah can, can come to our website and they can go through the system and use the program.

We verify their identity, we verify their medication history, their active med list. They go through the medical interview. And then the determination around eligibility. And at any point if someone wants to talk to one of our doctors, they can and they can get their refill. Can I hop in and ask a question about the value prop here?

So, this to me seems like a valuable service but what I have learned is that doesn't mean that it's a valuable business because health care has these weird incentives and weird structures. So, could you walk us through who the primary customer, so I know I understand who the primary user of Doctronic is, but who's actually on the other side paying for that user interaction, and what is their perceived value proposition of this?

Well, the first place to start, Andy, is Doctronic is a clinic. We don't charge for our technology for patients. We [00:26:00] bill insurance, or you can pay cash. This is how every medical clinic in the country operates, right? You, you go anywhere you pay with insurance or you can pay cash. Um, and that's what we do too. So, in, it's not dissimilar from what we're used to.

With the refill service, we initially set the price at \$4. Right now we've dropped that and it's currently free for most people. With time, the cost will come down even further. So, it, I mean, it used to almost be like a single fee per use, but if it's free, then what is the sort of like, overall, how, how did you drive that cost down to zero, I guess would be the, the next question?

Well, the cost isn't zero. We're making it available so folks can use it. Um, we don't want there to be a barrier to adoption, and, and the, the monetization is a, a secondary consideration, right? The first place is to say, all right, let's, let's just get [00:27:00] folks using the system and take it from there, right?

We can figure out the, the business, quote unquote aspects that's a later stage consideration, right? We're our bui—, our, the other, you know, nuts and bolts part of our business, our clinic is, um, is, is working well and doing, doing just fine. So, Byron, I want to transition a little bit. I mean, I think that's incredibly helpful.

I feel like I have a much better picture, of sort, of how Doctronic operates in general and then also how you designed this Utah pilot. I wanna transition into this topic that I think about a lot in the context of my own research that Andy and I talk about a lot, and that I think any company that's embracing LLMs has to be very vigilant and very aware of which is the sort of guardrails that you put on these systems.

I know you've seen this that we talked a little bit beforehand. There was this recent, I guess it was a security company or cybersecurity company of sorts, Mindgard, that did an audit of your system and they wrote up this article where they have some [00:28:00] examples of, uh, how they were able to kind of prompt, uh, hack your system in a way where they conclude, and I want you to sort of, you know, give us your responses, some of the takeaways that they were able to exploit the system.

Prompt of Doctronic's system and turn it into a doctor who spreads conspiracy theories about vaccines, recommends methamphetamine as a treatment for social withdrawal, various other things. What's your sort of response to that?

And maybe also tell us about whether what they found, what is it all sort of relevant to what you're doing with the pilot or with anything else? Because my sense is that it's sort of divorced from the pilot. They're talking about the chatbot that's on the website.

Uh, but maybe you can, just tell us about that a little bit and just use that to tell us about how you think about guardrails a little bit more generally. Yeah. Well, security is the most important thing in our business. And to your point, this is not the Utah pilot. This is like the chatbot [00:29:00] on the website.

This is the chatbot on the website. And it was an exploit. Some people call it jailbreaking. Yeah. Uh, but it, it's, it's basically you trick the chatbot into saying outlandish things, which is exactly what happened. And, hey, it's an opportunity for us to be able to harden our system even further.

That experience is not something that any good faith user of Doctronic is exposed to. Right. You have to know how to, to do these particular tricks to get

the chatbot to say wild stuff. Now, as we go further, we're gonna look every health system and clinic in the country experiences attempts at fraud and abuse, right?

Every ED in the country knows that there are people who try and, and go to the, that emergency room and get doctors to prescribe them controlled substances. There are, there are times when people try and commit insurance fraud. Um, there's all types of ways to, to put forward fraud and abuse, and we're gonna see new ways of doing that with AI. In this [00:30:00] case, hey, the chatbot said some, some wild stuff for one user.

And we take that very seriously. It's not a systemic problem, but there is an industry-wide way to think about this, which is we all need to look and say, look, there's gonna be new vulnerabilities that we need to get ahead of and be really vigilant about. What are the guardrails that you place on the system that's being used in the Utah pilot?

How do you guardrail that, that system? In the system, the guardrails in cybersecurity are the same level that we have for our clinic and our technology, right? There's a HIPAA compliant backend. Things that happen in the chat interface are very different than the security and the level of compliance that happens with the backend, right?

So, that's an important piece to differentiate. You can try and trick the chatbot into saying some wild stuff, but that doesn't flow into the backend system and that definitely doesn't make it into patient care. And so, from our standpoint, [00:31:00] everything we do is treated as HIPAA compliant by default, even when our chats we don't technically have to meet that standard.

If someone comes in and chats anonymously, but we do it anyway because we wanna make it very clear that if you come to Doctronic, you chat with our system. We treat that as if you are sharing medical information, even if even when we don't have to. Do you, just as a sort of follow to that, do you do any, for the system that is being used in the Utah pilot, is there like input sanitization or screening for what patients can interact with when they're just talking to the chatbot?

What inputs can, can even go into the LLM or is there, and what I've also wondered about is, sort of, and I think you're alluding to this, there's the kind of chat side of this, right, which is talking to the patient, and then there's the system that actually does the renewal, if you will, right? And, um—, Mm-hmm. I, I would imagine that the system that is actually sort of responsible for the

renewal is not the same chatbot, but it's uh, kind of [00:32:00] separated and has much tighter guardrails and maybe even as sort of a rule-based system that can only operate within a couple lanes or something like that.

And that's just me speculating about how you might design this so that even if the chat goes haywire, if you will, there's, Adam and I joke, there's sort of a separation of church and state, right? There's a thing that can actually prescribe you medications even if the, like, it has a failsafe in a way that it can only operate within a few lanes versus just saying something, generating tokens that can go in any space of token space, right?

Or of language space, which can get wild. Tell us what you can tell us, but like, is there some, some type of that? How do you think about the security more generally? This gets back to systems thinking. Every system that has the potential for error, needs layers of defense, and layers of checks, and layers of redundancy.

This is how every high-reliability organization designs to avoid errors getting through. And we take the same approach. So, [00:33:00] you're exactly right. There's ways to make sure that quality is high, the inputs are accurate, things that are unexpected can get surfaced and escalated.

And then ultimately, we have human, multiple levels of, of human systems as well, right? We have the doctor who the system can escalate to, and then the pharmacist at the pharmacy counter who, uh, is another place to check. And, and I'll just add one more, which is the patients themselves. Patients who are coming to the system, you have to be on the medicine to begin with.

You have to come in already on the prescription and it's not de novo new prescriptions. And so, that also becomes a layer of, by design, to say, all right, we, we can automatically limit the the number of folks who are coming in to, or, or I should say limit not the number, but the, the scope of complexity that we're dealing with in this first iteration.

Um, and then grow in [00:34:00] capabilities as complexity increases. Cool. Thanks for that. So, I think some folks may be thinking now like the opposite side of the coin for security is regulation. So, could you walk us through how the FDA views Doctronic and things like Doctronic? As someone who's not the most steeped in regulatory philosophy, how does the FDA view what you're doing and what are the regulatory implications of this?

I think specifically this automatic prescription renewal capability that you've developed. The current regulatory framework for this in Utah is under state level practice of medicine authority. And Andy, I obviously can't speak for the FDA. The FDA has commented on this in our political article. Uh, and I'll, you know, folks can feel free to look to that as their, um, initial point of view.

What I will say is there's a discussion happening, and it's an important one that we're having nationally about where this [00:35:00] new class of technology best lives when it comes to oversight. And just to make that very real, we're doing this in Utah, under the purview of the state. Not in all 50 states, all at once. In one state, under one program. And without that program, it can't expand, right?

If more states don't, don't create sandboxes, this isn't going to be able to, to expand other places. And so that discussion and that dialogue that's happening, in real time right now is healthy and we'll get answers soon on where all actors are gonna fall on this from state to federal government and beyond.

And just to play that back, super helpful. Uh, the state of Utah specifically has created a regulatory environment that allows you to advance what's possible with AI-mediated patient interaction, and it's yet to be determined whether or not the FDA will adopt a [00:36:00] similar posture at the federal level.

Is that like a correct characterization or did I miss, would you like to restate that? I know that these things are always tricky but just wanted to play that back to make sure I understood that. What you just described at the state level is accurate. And what happens at the federal level is, look, I don't have a crystal ball.

So again, that's part of this broader national conversation about what comes next. And I do think there's a systems level angle here that when a hospital changes its workflow, right? That's not something that we regulate at the federal level, right? When a clinic redesigns its care pathways and protocols, that's not something that goes through those pathways.

So that's, that's practicing medicine and that's making changes at the systems level. And so, from that standpoint right there, when you think about it from that systems lens, that's part of how we can start to think, okay, does this fit into that bucket? Does this fit into a different bucket? Right? It's part of the, the reason we're having this dialogue is [00:37:00] there's elements of this that are in between.

And we're seeing for the first time, and we're needing to clarify. Every technology transition is really blurry. I mean, Andy, I, I come back to this. Airplanes are not flying cars, right? Airplanes are not flying cars. Even though airplanes came 20 years after automobiles, right? 1903, Wright flyer, 1885, Karl Benz invents the first automobile. And they look very similar at first.

You know, oh, forms of transportation. Oh, they both have engines. They require a trained operator. They're economically transformational. They're harmful if they're misused. So, you might say on the surface, oh, airplanes are just flying cars. Right? It's just a derivative of the old innovation. And actually that would be a huge category error.

If we tried to control aviation the way we control automobiles, it would lead to broad disfunction. So, I think we're, we've been treating LLMs a bit like flying cars, right? Mm-hmm. Just, you know, more software, [00:38:00] right? So, just another consumer application that I download on the app store, and I log in on Google, but they're fundamentally different.

Their capability is fundamentally different, right? Yeah. I think that airplanes are not flying. I think very well said. That makes total sense to me. And right-sizing regulation has always been the end goal. I think that often we try to think about deregulation versus massive regulation, but right-sized regulation, I think, is absolutely the goal.

So, how about, Byron here we transfer to the lightning round. Are you up for some lightning round fun. I can't wait. I've been waiting all week for the lightning round.

Nice. Alright, so—. I heard, I heard that you guys collaborated or maybe colluded with Adam Rodman. Oh yeah. Did Adam give you a heads up already? Nice. I do see, I do see a question in my script that says courtesy of Adam, Adam Rodman. So, I do think that your intel is correct, so. Alright. Great.

[00:39:00] Alright, so the first one, it sounds like you're familiar with the rules, so we'll hop right into it. Um, what was your first job? My first job ever. My first job ever was as a part-time clerk at a QuickTrip gas station in Atlanta, Georgia, in high school. Oh, the QT. Nice. The old QT. The QT, yeah. Amazing. Alright.

This one is the one that you had some, I don't know if Adam told you the question, but this one is courtesy of Adam Rodman. When you're seeing a patient and you're stumped about something, what tool do you use first?

Okay, I'll call Adam. Adam is the tool that you use first. Alright. And what is the, what is the AI tool that you use first? I, um. Or let's say computational more broadly. Yeah. Yeah, yeah. I, um, look, I was, I was using ChatGPT, and now, now I use Claude and, and I'll tell you why I use Claude because Claude will shoot me straight.

It's like, [00:40:00] hey, I'm not sure that's a good idea. I think you should do this as, as opposed to being overly, I appreciate your honesty or disagreeing with whatever you're sort of bossing towards. Okay. Yeah. ChatGPT is just, wants to be a hype— Yeah. —a hype bot and yeah, I, I find that less helpful when you're know, you're making critically, and you need to critically appraise your own thinking.

Apparently, the personality's a little bit different with the newer model, but I haven't played with it enough. But one of the most surreal moments I've had with a LLM is I got in a fight with Claude about a plot hole in a book I was reading, and it kept arguing that I was wrong and I was like, no.

Anyway. So, I do think Claude pushes back in a way that GPT does not. Andy, I think we need to have you interview Claude on the podcast. What a, can we just pause for a second? What a line: I got in a fight with Claude. I know the whole, like the goalposts have, like, been blasted into outer space now. Guys, I literally left the conversation angry because I was like, it's just being so hardheaded.

I was, like, so anthropomorphizing the interaction. It was insane. Okay, so we're gonna recover from that and we'll ask you the next question. So, this is one that if you've listened to the podcast, you're [00:41:00] probably prepared for, but we ask it. The heterogeneity in response is always interesting. So, will AI in medicine be more driven by clinicians or AI researchers?

AI in medicine will be driven by clinicians, and I firmly believe that. And part of why I firmly believe that is because I live it every day at Doctronic. And I see that medical AI needs physicians at the helm in order to work. Period. Full stop. It is not more software. And we needed the AI technologists to develop large language models.

They did an amazing thing. And now clinicians working alongside technologists can take it and bring it to life within our profession for the benefit of human health. Byron, what is the most important skill residents should be cultivating today? Excellent judgment. Clinical reasoning [00:42:00] is now abundant. You can type a presenting complaint into Claude and you can get a great differential,

but judgment under uncertainty remains something that human beings are going to need to continue to exercise.

And. Yeah. Fundamentally, the doctor is a moral agent, not a technocrat, right? Our job is really to help determine what is right and what isn't when it comes to matters of human health, not just make diagnoses and management decisions. So, we should lean into that, get better at great judgment. Hmm, great answer.

Okay, next question. Byron, if you could have dinner with one person, living or dead, who would it be? I keep trying to get dinner with Raj, but he's always busy. He's always busy. Oh man. Sorry. I was outta town last week or whatever, two weeks ago. I missed you. Yeah. You know who I've been thinking about lately that I would really love to get dinner with? Guy named John Bumstead.

John Bumstead was the [00:43:00] physician who treated Anne Miller. Anne was the first patient in America to be cured by penicillin. Wow. That was at Yale Hospital in 1942. Wow. And Dr. Bumstead treated her the next day. Her fever of 107 degrees had broken and the penicillin transformation was underway. I see a lot of parallels between that and the situation we find ourselves in today with medical AI.

I see why you and Rodman are friends, like, the history of medicine, shared enthusiasm for that just comes through a lot, so, yeah. All right. This is our last question. Also courtesy of Adam. Do you think that Doctronic's AI will eventually be able to do everything that you can do? I think he means as a clinician, as a practicing clinician. Yeah.

Doctronic is very talented. But these are artificial systems. Cognitive systems are artificial and intelligent, but they don't replace humans and. [00:44:00] They will never have the full context that a human mind and body can have in society. So, I think the answer has to be no. Even though I appreciate the spirit of the question.

Could you say more about why they'll never have full context? 'Cause I often think that like an LLM keeps way more in context than like my meager lizard brain is capable of keeping in context at any given time. Well, Andy, it's an important nuance. I mean, I, I don't diminish these, the power of these systems.

I agree with you. I think this is about as a human being we uniquely experience the world and that will, by definition, never be equivalent to a machine. And where there's overlap, great. We can use the machine to our advantage. Where

there remains space and differentiation, we should lean into that because those are the things that, where we bring unique value as people.

But do you not think that [00:45:00] this is, we're going off script now with the lightning round. Okay. But this is a, it was a fascinating answer, and I think it goes along with one of your responses to, uh, one of our other questions, which is, what should residents be cultivating as one of their top skills?

Do you not think that future LLMs or some would say even current LLMs are capable of eliciting values of the patient, eliciting their goals, and then piping those into a system that can also help reflect the patient's values in their approach to care. Raj is doing the thing, Raj is doing the thing where he restates one of the titles of his papers.

Just to be clear. Is that a thing? Do you not? Is that a thing? I do, do you not agree with me? Exactly. By the tone of your voice? Um, Raj, I, I think that is absolutely true. I think these systems. It can do a lot for helping us make value-aided decisions and there's a limit to how far they can go, where, what's available to the LLM and what's available to the [00:46:00] human, uh, diverge.

Hmm. Fascinating. Alright, so this is our last, uh, actually that's the lightning round. Byron. So fantastic. I think it was a very, very interesting set of answers. Uh, you did very well in the lightning round. We have actually just one sort of big picture question that we wanna kind of zoom out with. This actually goes along with some of the discussion we were just having.

Uh, you know, things have moved so fast, right? I have to just pause and reflect on where we are in 2026, what we take for granted now. So, Andy just told us about his fight with an AI about a book and the AI's obstinance and hardheadedness that it was right. To be clear, there's a very big plot hole in the Hyperion Book series for those who are listening.

And I'll get Claude, too. Oh my God. Uh, full Andy on full display here. And so, this is a, say, like, that is not even, like, even surprise us that Andy can be talking to a computer model [00:47:00] and it can be arguing with him and doing so coherently, and maybe wrongly in this instance, but having a sort of full-on argument, right?

As you said, we have AI models now that are demonstrating a lot of capability in diagnosis, right? To the point that now we take for granted results, of course, like the professional competency exams, like those aren't even interesting to folks now. Oh, we got a hundred percent on the next incarnation of the USMLE.

And they can solve these CPCs. They can solve lots of clinical vignettes. That has moved so fast, right, from, I can remember these conversations I would have with Andy when we were postdocs and when I was a grad student. And it's just moved so fast over the last couple of years where the goalposts are now on the capabilities of computational models for what seemed to be very hard medical tasks.

So, I'm gonna ask you to do something hard, which is now reflecting on the progress of the last couple years, can you project where you think we're gonna be in five years? Let's say around 20, 30, maybe a year after that? [00:48:00] Can I make it even and let's, let's say something, like, even harder. Yeah. Uh, ChatGPT-3.5, the bad one, came out in November of 2022.

Uh, four. ChatGPT-4 came out in March of 2023. So, this would be like, uh, 1.5 times the amount that has passed— The interval? —between the interval. Yeah, the interval between then and now. So, uh, go wild. Where are we going? I wanna start with what I think is in front of us, which is a choice. We have a choice as a society to make.

The evidence is in. It's clear that medical AI is a transformational opportunity to improve human health, and if we invest in the next decade, we can bring it to scale globally. The scarce resource of cognition is now abundant, but [00:49:00] just like penicillin was the first of many that would come in the category of antibiotics.

The cognitive systems we have today are the first wave of what we can bring into existence in the next decade if we choose. But it is a choice and whether we want to choose urgency and bring these systems to scale, or whether we want to accept delay is a societal values choice. And the decisions we make now are going to determine whether we have it in the next decade.

Or in 20 years, or in 30 years, right. The harms of inaction are real, and speed is not incompatible with responsibility and with rigor. But it is a choice. Is it too uncertain to predict or do you think in five years the practice of medicine, the life of a resident is gonna look very different than it looks today or, or not really?

It. This [00:50:00] gets to your point, Raj. It already looks different. Yeah. I mean, the majority of doctors in the United States now use AI in daily practice, and that's an objective statistic, like 40% plus maybe more adoption of OpenEvidence. You know, we're all using it for our, we know it makes it us better. And, um, yeah, it, it, if it's, it's, it already looks different today, it's

gonna look different in five years and how different it looks and for whom ist he choice in front of us.

Alright, I think it's a good, good note to end on. Byron, thanks so much for coming on *AI Grand Rounds*. Thanks a lot. Yeah, we enjoyed it. Hey, it's my pleasure guys. Thank you for having me. Of course.

This copyrighted podcast from the Massachusetts Medical Society may not be reproduced, distributed, or used for commercial purposes without prior written permission of the Massachusetts Medical Society.

For information on reusing NEJM Group podcasts, please visit the permissions and licensing page at the NEJM [00:51:00] website.