AI-GR Pod 28 03.17.25 Morgan Cheatham

[00:00:00] If you're a commercial insurer today and your average member is going to be with you for two years. Why do you care about paying for their molecular testing to prevent some disease that might manifest for that person in a 40-year time horizon? Why do you wanna be on the hook for the gene therapy or the siRNA therapy that would, in theory, cure that patient if you're not going to be responsible for the cost on the line?

These are the real questions we have to grapple with. Again, if people take away anything from this conversation that we have, I hope they hear me when I say: the science, the computation, the technological innovation is no longer the barrier. The tech is not the problem in health care and life sciences.

The problem is the business model, the economic model, the way care is paid for, and the incentive structure underlying that.

Welcome to another episode of *NEJM AI Grand Rounds*. Today we are delighted to bring you our [00:01:00] conversation with Morgan Cheatham. Andy, I find it really hard to know where to begin to describe Morgan because he's doing so many things. I have no idea. I think we did ask him at one point. Are there multiple Morgans?

He is a vice president at Bessemer Venture Partners, while also a medical student at Brown University. He's doing so many things. He's doing them so well. And it was just a really wonderful chance to just pick his brain about how he approaches everything and how he's thinking about both investment strategies while also navigating medical school and residency and all that.

Yeah. In addition to like doing stuff with us at *NEJM AI* and this podcast, the man has more than 24 hours in his day. And again, breaks my mental model. I think I, I think we discussed this in the episode for like what a med student can do and is capable of. By day, he's rounding on patients, doing research at the NIH, and by night is cutting a hundred million dollar deals at his role as a vice president at a venture [00:02:00] capital firm.

So truly just an extraordinarily talented young doctor, future leader, all of the above. I think also, you know, you and I have gotten to know him quite well. Just a great guy, like super nice, very, very easygoing. And yeah, it, it was great

to have him on the podcast and give him a chance to tell his story because it's such a remarkable one.

And I should say it's, it's a little bit funny to me. He also helps us edit these podcasts. I was gonna— He had, he had the opportunity— So, so Morgan— To also work on his own podcast. Thank you, Morgan. Morgan. Morgan. If you're listening, uh, thanks. Thanks, Morgan.

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 Morgan Cheatham. Morgan Cheatham, welcome to *AI Grand Rounds*. We're excited to have you here. Thank you, guys, so much for having me. Morgan, great to have you here. So, this is a question [00:03:00] that we always get started with. Could you please tell us about the training procedure for your own neural network?

How did you get interested in AI and what data and experiences led you to where you are today? I love this question. So, to give you guys a framework, I think today my training procedure looks a lot more like a Bayesian process than it does stochastic gradient descent, and I figured our listeners would appreciate that.

I think as humans we have this luxury where unlike most neural networks, which kinda discard past intentions in favor of converging on a single solution. We as humans can maintain these large prior distributions over complex environments. And we kind of constantly refine them over time. And so, what this looks like for me is a huge focus on information diet and curating that diet in a way that allows me to refine these priors in a way that I feel like I'm, I'm converging on something novel.

And so oftentimes that means spending time in subfields or in the interstitial spaces between fields. I realize it's a very amorphous answer. So maybe brass tacks my story. [00:04:00] Growing up right outside of Washington, DC, I actually always wanted to be a physician from as young as I can remember. And everything I did was in service of that goal.

I didn't falter. It's what I wanted to do. I had this moment. Did you have, sorry to interrupt, but did you have any physician role models early on? Not until high school. So, there are no physicians in my family. In the Washington, DC area it's actually, I think, a profession that's quite revered still. It's something that

I'd always looked up to. And being near NIH and all these amazing institutions certainly keeps you inspired.

But so, I actually did have this experience with a mentor in high school, which was kinda the first time I saw the collision between computation and medicine. And that was shadowing at a hospital called Inova. And at the time, interestingly, they were implementing Epic. So, they were going from paper to Epic.

They weren't switching from some other system. And I think the chaos of that entire experience completely blew my mind. I mean, from the disgruntled comments about what the system was gonna do for medicine and papers flying and the chaos, I just, I knew this was a big moment and I knew I had to attach myself to it in some way and study [00:05:00] it and pursue it.

And so, when I went to Brown, ultimately, I created my own major as one does at Brown. Not in basket weaving, but in what we called computational decision sciences, which was an interdisciplinary concentration. Frankly, way too much for undergrad alone, where I wanted to interrogate how humans made decisions. And I wanted to do that through the lens of kind of classical mathematical microeconomics but update that with what we now understood in computer science and particularly computational neuroscience.

That experience is kind of where I started to go more deeply into AI and really understanding not only how these systems interact in a theoretic context, but ultimately how they would apply to medicine. Okay, so let's keep rolling that forward though, because I think, let's talk about what you did in medical school.

You did this interesting thing in between undergrad and medical school, too. So, let's hear a little bit about that, too. I will just maybe also flag that you're now the third person on the podcast who has had a meteoric start to their career. [00:06:00] Who has gone through this program at Brown. So, maybe there's something in the water there, but our own Zak Kohane has gone there.

I believe Atul Butte also went through the, yeah, and, and so now Morgan. So there, there's clearly something iconoclastic in the water at Brown that makes these physician scientists who are also like deeply computational and deeply disruptive. It is an interesting thread and those are obviously two of my heroes who I've been fortunate to, to meet and hang out with.

I think that's something about the freedom at the institution, frankly, where there's no institutional silos between departments that really allow people to

traverse. Whether it's medicine and computation or humanities in medicine, et cetera. But yeah, so I guess, I guess just to answer your question succinctly, when I was at Brown, I did have a physician mentor by the name of Graham Gardner, who was a physician, CEO of a company called Kyruus.

And what he was doing at the time was he was leveraging large data sets to try and match the right physician to the right patient at the right time. And of course, we can appreciate on this podcast, it's a problem we're still trying to solve today.

But it was really working with Graham and under his team that I learned how to [00:07:00] program and I actually began to gain real-world experience with these systems in a production environment. I will say that the other interesting formative experience I had after that was working at Goldman Sachs for a summer.

I was a very crappy banker, but I learned a lot about the ways that capital could scale technology and business in health care and life sciences. And I think that was a pretty eye-opening experience from the jump. And so, I guess the way I described the path was I'd almost seen too much. Like, I'd almost refined my priors a little too much, and I knew there was this expansive world beyond just practicing medicine that I wanted to explore.

And so, in a very serendipitous fashion, I bumped into a partner at Bessemer. I know folks sometimes like to hear these origin stories because it is very random to break into VC. Ultimately joined them as an analyst right outta undergrad. And I think it was, you know, supposed to be this two-year deferral from medical school that soon became four. The pandemic happened.

There was so much happening in technology and medicine and health care. And then we had our ChatGPT moment, and the rest is history. Now, in addition to AI, clinically, I've found a [00:08:00] home in clinical genomics, which is a field that is changing rapidly. And, like AI, seems to be benefiting a bit from a Jevons paradox in that as genomic technologies become more efficient, and accessible, and cheaper, their overall usage is increasing. And the field is also undergoing this interesting transition from an inherently diagnostic specialty to increasingly interventional by way of genetic medicines.

And so, this is where I've spent the last 18 months, clinically and from a research perspective. Yeah, when I first met you, I had a mental model for what medical students were, and they were these like very, I think cookie-cutter

biochemistry, you know, or go undergrad, taking the straight narrow path to medical school.

And I think I met you and you were a first-year medical student and a vice president at Bessemer, and I was just like, who is this guy? I would love to learn more about Morgan and how he got to where he was. You know what I love about your story so much is that you're like emblematic of this like new kind of medical student that like breaks the sort of [00:09:00] very cookie-cutter mold that medical students have been subjected to for a long time.

You come from investment in capital formation. Other folks are coming from engineering and mathematics. I don't know exactly what explains it, because arguably the job at the end of this training procedure hasn't gotten better. You know, in many cases it has gotten worse, but it's been exciting to see the intellectual diversity of who's going into medical school now relative to like 10 or 15 years ago.

I truly don't think I'm unique. There are lots of Fs out here and I hope, I hope you have more in the pod, but I think there's this genuine desire from a lot of people in medicine, to your point, recognizing the current constraints of the system to have an impact at multiple levels. And so, I think when I look back on the last 10 years of my career, I think there's this very sacred one-to-one physician relationship that the canonical medical student or trainee pursues and studies and, and I think there's knowledge and that experience that cannot be replicated anywhere.

There's certainly privileges and opportunities that can't be replicated anywhere, but there's specific knowledge that you gain being responsible for someone's care. And then when you think about what's happening in technology and in the [00:10:00] capital markets, those are things that can take that kind of sacred one-to-one relationship and scale them in a non-linear fashion, whether it's thinking about AI, doctor constructs, or thinking about investing capital into a therapeutic that's going to cure a disease.

And so, I think people enjoy genuinely when you have this experience of caring for patients, being able to play at these multiple levels of the stack, if you will. Yeah, totally. So, I think that's a great opportunity to transition to some of the papers that you've worked on and the research that you've published. So, you, you spoke of the ChatGPT moment, and I think that you and some of your collaborators like really saw that this was a moment not only for AI but also how it was gonna impact medicine.

And you wrote what I think has become like a pretty seminal paper in the area, which was essentially a can ChatGPT past the USMLE? And so, there's a paper called "Performance of ChatGPT on the USMLE: potential for AI-assisted medical education using large language models." So, I'd love to hear a little bit about the context for that paper and sort of like how you put it together.

If I understand correctly, it's like a sprint to put it together. It was like nights and weekends kinds [00:11:00] of things. So, we, we'd love to hear some of those, like in the trenches stories, too. Sure. And I'll caveat this with saying my academic career, as you mentioned, has been unconventional, but it's also been quite random.

And I think this paper embodies that. So, I actually worked on this paper with one of our portfolio company CEOs at Bessemer, physician-scientist by the name of Jack Poe, who also was the person who introduced me to, to Zak. So, I'm ever grateful for that. And there was this moment after ChatGPT came out, we were all scrambling to figure out what this technology would be capable of in a medical context.

And, you know, I think I ground that in the reality that there are many, many teams working with GPT-2 and earlier versions of this. But the performance and the interface, I think invited a larger audience of folks to, to explore the capabilities. And so, Jack called me up and said, hey, you know, our team here at Ansible, we're, we're working on this project.

It's super time sensitive. Do you wanna be a part of it? And, when someone gives you an opportunity like that, I think you, you just ask them like, what's the next step? You don't even pause. The reality of my life at that moment though, is I [00:12:00] actually had a neuro shelf exam that Friday, and I think we were, we were meeting at the beginning of the week. So I was, I was supposed to be studying, but I was humbled to learn that the system that we were working with would've probably far eclipsed my performance on that exam.

Whether I studied more or not. In many ways when I reflect on this paper, I say it was both the best paper, best paper experience I had just in terms of, I think, the speed and the kind of like radical thinking that we proposed. I mean, we actually submitting it to some journals had ChatGPT listed as an author and were scolded by some of the larger, kind of more well-known journals. Ended up publishing in PLOS.

But I also say it was one of the worst. And the reason why is because the paper at the time positioned GPT as a potential learning tool, even though its performance kind of suggested there was more there. What we didn't grapple with is really the imperfection of the benchmark of the U.S. medical licensing exam.

And in many ways my, my worst fear is materialized, which is, as you both appreciate, the field kind of ran [00:13:00] away with this benchmark in a very obsessive, uh, obsessive approach. And so, I wish we had grappled more with the kind of shortcomings of the USMLE, both for human clinicians as well as for AI.

And I think now in the last kind of couple years, more folks have started to grapple with this in a real way. As we saw most recently with the CRAFT-MD paper. Just to calibrate, right. I believe Raj has some comments on this. Yeah. Yeah. So just to calibrate, just to give your paper more credit, I think, than, than even you're giving it now, Morgan.

This is back in the beginning of 2023, right? So now I think we all take for granted that these models are very good at standardized tests. But it's very easy to say that now, after we've had a flurry of two years, or three years, or four years of results on NLP models, AI models doing so well on standardized tests.

I also think at the same time, you know, this is actually even before GPT-4, right? This is the [00:14:00] original ChatGPT that was released, if I recall, now I'm gonna get the years wrong, this is November 2022, right? That's right, yeah. This is, yeah, so this is like, you know, right at NeurIPS 2022. I think Andy and I were there at that one together and it just took over the whole conference.

That's the model that you're using. This is even before they released GPT-4, you know, I think a month or two after you guys published this paper. And so, I do think it's worth just remembering where we were at that moment, and I'll go back even a few years earlier than that. I was spending a lot of time, uh, with this guy, with Andy Beam here, my cohost, just talking about this idea. And at the time people would pillory it the other way, which is that this is an impossibly hard task to solve, which is to have an AI model take a U.S. medical licensing exam. And it's very far away.

And I remember Andy getting comments like that all the time and, okay, wow, you guys are [00:15:00] really aspirational, or really visionary, ambitious, right? But I think we should contextualize where we were and what these sorts of

milestones led us to now believe. And then I think as you're saying, I do think lots of folks are pointing out. I mean, I think there's another great example of this, the AMIE paper from Google, right? So, this is now almost a year ago, I think, the pre-print came out and they really very carefully studied conversational interaction with AI models.

And then there was CRAFT-MD from a couple weeks ago. I think a lot of folks are now correctly pointing out that there are limitations for these benchmarks. But at the same time, I don't want to rewrite the past and say that it was always obvious that we could do this. This is an easy task, and I think it's, we're also in danger of saying that these tasks mean absolutely nothing. So, they don't mean that, I don't think they mean, I think we're having a good conversation right now [00:16:00] as a society around this.

They don't mean that they're sufficient, that someone could practice medicine, or an AI can practice medicine. I don't think most people are saying that, but I think we're going to the pendulum swinging the other way where we're saying they mean absolutely nothing about what these models can do, and they are completely meaningless.

And I think that might suffer from the same sorts of sort of intellectual impulses that would conclude in the opposite direction for someone to say that the benchmark means everything. So, it's just not, we don't need to be so black and white about what it means, what it doesn't mean. May maybe I'll hop in there. I can see, yeah. Yeah, and just add like a little color.

So, like in 2016 or 2017, I had been going around giving this little talk that saying the USMLE should be a benchmark for medical AI. I gave a talk at like at 2017 at GTC, NVIDIA GTC, like literally saying we should have this be a benchmark for medical AI. I was obviously working on it. That was a hot team.

Yeah, I was obviously working on it myself, but I was like a couple hundred billion parameters short of a scale of model that you need to actually pass it. But I got those. But they're probably in your brain. They're probably in your brain. Yeah. [00:17:00] Yeah, yeah. Um, but like the criticisms were then as they are now, and that like Raj said, some people were like, this is insane.

Computer could never do it. So, it's, it's also like I have the same sensation that Raj has where it's what was once impossible is now trivial. And like I have now, like, lived that experience many times in AI over the past like seven years. But that was, I think one of the first times I had that experience.

But two, like people would say, well who cares if they can pass the USMLE? There's a whole side story. Actually, I got in trouble by the NBME for saying that AI was gonna pass the USLME. They don't seem to care about that anymore, but we'll leave that for another day. I'd love to see those emails at some point.

Yeah. Uh, but they would say like, if, even if a computer can do this, like who cares? This isn't gonna help them be a good doctor. And I would just, like, what Raj said, that this test is a necessary but not sufficient condition for humans to be a doctor. And we should view it the same way for AI. Like if it can't do this, then we wouldn't trust it to practice medicine.

But it's, again, it's a necessary condition, but not a sufficient condition in [00:18:00] that there's all these other things that we need AI to be able to do well. And I, and I still think that there's signal in how much better it does relative to a person. It's clearly testing some type of medical IQ or test taking ability.

So, I'm also likewise similar to Raj, like I wouldn't want to dismiss these out of hand. I think there's also signal in the relative ranking of the models on this benchmark, even if we don't understand what the absolute calibration means for clinical practice or for being helpful in clinical practice. I also think, yeah, I think I, I'm just hesitant to like discourse extreme one way or the other, right?

It's meaningless or it's super meaningful. It's just one piece of information and I do think the general, I think, move recently towards more realistic clinical benchmarks. I do think that has been in the water for some time. I also think that's not something that we're realizing two weeks ago, I think years ago.

I think the folks who were also putting forth the USMLE as a benchmark were saying the same thing. And I think just out of respect for intellectual forebears, I think if you look decades [00:19:00] earlier, and you look at the way people are saying things about medical education and training and teaching doctors and residents, they've been saying the same thing, too, which is that this is a piece of information. But we need realistic clinical appraisals, realistic clinical evaluations, and we need different types of evaluations.

And so, I think this actually goes back much further than even the last year. This is a decades old debate in psychometrics and medical education. And really is worth kind of bringing that history in because there's also lots of limitations for the non-standardized tests as well. And I think we should, we should keep those in mind while we're integrating them into this conversation.

One thing I love to talk about, just in truth, and I, I agree with a lot of what you all have said is just the experience of going through medical school during this period right of my second year, seeing this technology released into the wild and frankly grappling with the reality that the benchmarks I was being assessed on my shelf exams.

The USMLE did not explore what it would be like to be a [00:20:00] physician in that era. And so, I think that there's kind of two separate conversations. One is knowing that once you're a board-certified physician, most of your recertifications exams are actually open book. And many of my colleagues are using OpenEvidence, ChatGPT, Claude, to pass those exams now.

How do we refactor the human-centered evaluations to consider AI? And then there's the separate question, which is an important one, which is how do we think about AI benchmarks that grapple with the context of a clinical environment? And I think the recent CRAFT-MD paper is interesting in that it simulates this, you know, interaction between the patient and the physician.

And I'm a big information theory fan, and I think one thing that is lost in that framework is the reality of the role of the physician as this information gatherer, as this person who in seven minutes in a primary care setting, is forced to build a trusted relationship with someone and to ask questions in both an efficient, empathetic, and approachable way that they're going to maximize the channel capacity of that interaction.

And the study, I'm hoping someone will run. So, this is a [00:21:00] request for study for our listeners, is to actually run the experiment between a voice AI agent, which is we appreciate our increasingly kind of sympathetic and human-like performing that medical interview with a patient autonomously.

Comparing the diagnostic performance of that AI with a human physician on the information that they respectively gathered. So, the human physician conducting their interview as they typically would, and then assessing the downstream diagnostic performance based on the information gathered. And also saying, did they gather the same information in the same way?

And what were the key differences in how that information was gathered? How the patient responded, right? So, I think there's certainly limitations to simulating the patient response through AI because. When you have a living, breathing human being in front of you, sometimes you're more willing to share.

I will say there is some evidence on that from this AMIE paper from Google where they do this like turn-by-turn AI, there's a human acting, a human actor on the other end who's like role-playing as a patient who has a disease. And I only know this because my grad student was on this, and he just offended.

So, I had a [00:22:00] good jog of my memory on this, they look at if you give the information collected by the human physician to the AI system, it essentially doesn't change the diagnostic accuracy. So at least in this study, the AI and the human were eliciting equivalent amounts of information. The AI was able to get to the correct diagnosis more reliably with that same amount of information.

These are simple case presentations though, so like the information elicitation is like relatively simple and straightforward. So, it probably would be more interesting to look at more complex cases. Absolutely. I totally agree, Andy, and I think, something that you both just said, which is critical but can be lost in a lot of these types of studies, is the importance of a human baseline and an evaluation of how humans work with AI. And how AI compares to humans.

Completely agree that the human element here is key. The chatbot construct with simulated patients, I think is directionally interesting, but of course limited. And if we wanna understand how well AI can interact with human patients, it would be interesting to see the AMIE study [00:23:00] rerun with a purely voice AI construct.

For some patients, I would imagine voice being a more intimate material for sharing special information or intimate information. Others, I think would prefer text, and I would suspect this will stratify along the lines of age, but the ability to personalize the paradigm of interaction, I think is the beauty of our new medium, and we have to build with that in mind.

And importantly, I think what this research also shows is that we need to revisit the fundamental clinician-patient interaction and begin to unbundle the components of the medical appointment that we've classically understood now with an AI first mindset. I think we are learning a lot and some of it is surprising us, right?

There's been a lot of good soundbites about how humans and AI will just do better together, right? If we insert the two together, you combine them, they'll be better than either alone. I think the picture's becoming a little bit more nuanced right now around where there are areas where humans and AI can collaborate effectively, [00:24:00] where humans do certain things better, where AI does certain things better.

There's a good essay that one of my colleagues. Benod Raj Picard just wrote in the *New York Times*, I think literally a day ago it came out that it goes into this where maybe it is the case that there are parts of the workflow that are autonomously conducted by AI and by humans separately. But I think it's very, very important to not look at one of these systems in isolation.

Right, like the LLMs alone, or as you're saying, Morgan, the sort of simulation of a patient or of a doctor by an LLM. It's something that I think needs a lot of validation and I think it's very important to integrate actual studies of humans, both in the use of the technology, but also in the appraisal of the outputs of the models themselves.

Maybe actually that's a good transition point, Andy. Lemme ask the one, the one, one last question before, before we transition. So, I mean, you kind of hinted about this, but like step one and step two are these watershed moments for med students. It dictates a lot of what the rest of your career is gonna look like.

[00:25:00] Was there any sense of remorse, mixed feelings, hesitation that ChatGPT passed step one and later with flying colors, before you, yourself, the future doctor to be, passed it? Did you feel like John Henry. Honestly, I was happy to see it. My zone of genius is not testing. I've always been like a learner who learns through experiential dynamics, and so I was happy for ChatGPT.

I was happy that I got it done and kind of moved on with my life. I think that's a very centered way to look at things. Yeah, I think that's, that is, yeah, I totally agree. I think a lot of folks did not look at these results that same way. Right? And it was a little bit of, I mean, I was talking to my doctor friends and family members when some of these results started coming out, including your paper Morgan, and I think there was a sense of existential concern. But then I think it quickly morphed into, well, the tests really don't mean that much.

And they, as we've discussed, they are not the end all, be all by no means. Alright, so Morgan, we wanna jump, we've actually already talked about this a little bit. But [00:26:00] we want to just spend a little bit of time. I think one of the reasons we invited you on is because both Andy and I are amazed at your ability to balance many, many things. Like, you know, we all talk about multitasking.

I think we're all busy, we all do lots of things, but you really are, balancing several different, amazing, amazing activities as part of your work, both within medical school, of course, and then also academically and then in VC world.

And so, I guess my question is, are there two Morgans or more than two Morgans?

You know, medical school is already really hard, and how do you possibly manage to get everything done? What is your productivity hack? I, honestly, I'm not the most productivity tool hacker person you'll meet. I'll be honest. I think a lot of the reason I was able to do this was because of my team at Bessemer and their willingness to support me through this.

The scheduling was crazy. Like, going into the OR at four in the morning some days and stepping out in the afternoon when our cases were [00:27:00] done to hop on a call with a founder taking a board meeting at night. I mean, they made that happen. And so, in many ways it's just the people that I work with who are supportive of the path.

But I also think to the earlier comments about information diet and refining our priors, the things that I work on. So, research and computation and working with y'all at the journal, the work with Bessemer. Leading our health care AI practice and practicing medicine are highly synergistic in nature, and each one feeds into the other in different ways.

And so, whichever kind of context I'm switching into, I know it's serving something else I'm working on in a, in a really interesting way. And so, I've kind of given myself the intellectual freedom to know that's true, even if it feels like it's a little bit of a longer journey to connect those dots. And so, I guess that's my short answer to what's probably a much, much more complicated question.

Yeah, it just, as an aside but also meta related to this, you know, entire conversation, have you used AI at all to help with medical school? Like, is AI useful for you as a tool to go through the preclinical curriculum or through your clinical training [00:28:00] now? Has it been useful for you? Absolutely. So, when I was in medical school, I actually, on my blog, put out a blog post and I said the title of it was, you know, "Put your healthcare AI app to the test."

And so, I think in my unique role as an investor, you know, I've been meeting lots of founders, a lot of physician founders, building tools for other physicians, and I just dogfooded this stuff for four years, right? When I think about some of the companies we're invested in, hopefully we'll talk about a few of them, Abridge in particular, one of the AI scribing companies.

Oh yeah, I remember this. You, you actually, I was, uh, okay. Wait. We should tell the story. Right? So we were, Morgan and I were at a conference together, right? And this was a conference organized by Zak. And Morgan was talking about Abridge, and I'd heard about it, but then he was like, let me just show you a demo.

And I was like, okay. And I think I acted as a patient, right? I acted as a patient. You did. You took my history. You talked, first of all, I was very impressed with your bedside manner. Thank you. You took my history. You're a very good doctor. And Abridge, the Abridge app, very, very effectively summarized the note [00:29:00] there, right?

And did it in I think basically real time. And you showed it to me, and it was pretty cool. So, what, why don't you just tell, I think a lot of people have heard about Abridge, but maybe you can tell folks about Abridge. Well, I hope the Abridge team hears this 'cause always be selling, you know? But so yeah, at a high level, Abridge is a tool that clinicians can utilize to capture the conversations they're having with patients, and from a system perspective, capture that information in a way that is complete and accurate.

And that enables that data to be utilized for downstream use cases in terms of providing that information to other clinicians, providing it to the patient in an accessible format. And then thinking about all of the downstream activities a health system has to utilize that information for, whether it be revenue cycle, clinical quality reporting, and the like.

We invested in Abridge back in 2020. Interestingly, at the time, the company was a direct-to-consumer app. So, the founder, Dr. Shiv Rao, he's a cardiologist by training who partnered up with Zach Lipton at Carnegie Mellon to develop this technology. And initially the distribution was, let's go direct to [00:30:00] patients who have this just acute pain point of getting shuffled between appointments from their PCP to their endocrinologist, to their cardiologist, none of whom are talking together.

Can we reverse engineer a solution to interoperability, starting with a patient? It's kind of heralded as this concept of the personal health record. Can you update the patient's information and phenotype as they're bumping into the system and experiencing health care? Interestingly, that product as a consumer product did gain traction with a lot of rare disease patients, unsurprisingly, who really feel this pain point acutely.

But ultimately, as the company matured and frankly as the market matured and as products like ChatGPT educated the health care executive landscape that this stuff was real. This stuff was not just pie in the sky, founders selling a fake vision. These were real products. The company repositioned towards the enterprise and kind of the rest is history.

I will say I do think that the ambient scribing landscape, whichever vendor you're talking about, is the greatest unanimous adoption of technology by health [00:31:00] systems and hospitals since the electronic health record. And this has all happened without the Hitech Act equivalent or meaningful use equivalent for the adoption of AI and grounding ourselves in that reality is important.

So, let's stay on that for just a second because like the thing with Abridge and the ambient scribing technology that has been surprising to me is the traction and durability. Like, one of the hardest things in health care is business models. Who the customer is often like if you have an AI health care application is not obvious.

So, is this like the alignment between things that doctors actually want to use and some type of sustainable business model? Like, do the hospitals pay for it because it actually does facilitate billing and reimbursement more efficiently? I see why the doctors like it because it lowers their administrative burden. Why do the hospitals like it? Yeah, well, let's ground ourselves in the reality of where things are post-pandemic. Right? In the, in the years after the pandemic, half of U.S. hospitals were margin negative, meaning they were losing money on provisioning of [00:32:00] care.

You had the great resignation of health care workers, and not just clinicians, but really across the system, people facing just tremendous amounts of burnout in the current way of working and then you had this moment in AI with GPT and these other technologies that allowed us to build these delightful product experiences that had been kind of unseen in health care. Right? Design has never been a strength of our industry, and so I think in many ways you had the perfect confluence of these trends coming together to facilitate this adoption in the ambient scribing category.

Now, as it relates to the business model? Yes, I think most of the players in the space are selling this technology to hospitals as a software product, as a tool for clinicians. And as they, many of these companies, approaching this as a point solution. We are an AI scribe and that's what we do. I think as they move into other spaces, like you mentioned, revenue cycle management, maybe clinical research, I'm a big believer that ambient scribe should improve the quality of

the data we collect about patients and will be a win for research and what's happening in life sciences and [00:33:00] real world data.

I think other business models will reveal themselves. The big thing that keeps me up at night, if you ask me about AI, it's not whether we'll be able to do certain things. I think I'm techno optimist and believe in the arc of progress bends towards interesting things. It's whether or not in our current system and all of the bureaucratic apparatus around it, whether we will be able to afford this stuff and not just afford it for our commercial populations but afford it for everyone.

And I think people working in technology should pay attention to that problem because we all wanna build products that people use. And unfortunately, the reality has been in AI that the adoption has focused on stuff that makes money. But when I think about the greatest potential here, sure, revenue cycle's cool.

Can we phenotype disease in a more precise way to actually get patients' answers faster and connect them to therapies that are curative faster? That's really, I think, where the technology goes and where the greatest impact will manifest. Cool. Awesome. [00:34:00] Morgan, can I ask you to talk a little bit about your general sort of investment thesis?

I'm sure you are getting pitched all the time. How do you cut through the noise? One of the funnest experiences in medical school, just as an anecdote before I answer your question, was getting pitched by some of my professors and, this is sometimes on rounds, you know, like after rounds. Amazing. Hi, Morgan, I have something to sell to you right now.

Yeah, yeah, yeah, yeah, yeah. It was, it was amazing. I met a lot of wonderful people through that. But look, I, you know, at, at Bessemer high level in investing is both thesis driven and opportunistic, right? So, we come up with what we call roadmaps, which are thesis about where we see the puck headed in technology over the next five to 10 years, per se.

We can also dream about where it goes in 30 years, but in the venture capital business, being early, being too early. Is often being wrong, right? And so, you need to believe that these trends are manifesting over a near term time horizon in health care and life sciences. Just for folks who are less familiar with Bessemer, we've been investing in the category for 40 years.

So, we have a dedicated focus and [00:35:00] commitment to these industries. Even though building in them can be very hard, we think it's quite worth it and

are long-term oriented. And we do that across modality. So, I describe us as modality agnostic, problem specific investors. We invest in software companies, tech enabled services, services, therapeutic platforms, and diagnostics.

And when I think about my approach, it's let's find the biggest problems in health care and medicine and let's figure out what the right modality is to solve it. It could be a diagnostic therapeutic combination. It could be a software plus a diagnostic. It could, you know, you can see how all these combinations can become quite interesting.

And we're doing that from the pre-seed to growth stages and in AI. When we started investing in the category about six years ago, we were really focused on what I describe as AI first founders. And I'll distinguish what AI first means from AI enabled. So, AI first companies and teams are those that are actually innovating at the methods layer in AI, right?

They're publishing papers, they're presenting at top AI conferences, and they're really advancing what our capabilities are from an applied AI [00:36:00] perspective. And I think Abridge is a great example of a company. We have a company called Subtle Medical that spun out of a Stanford radiology lab that was also in this category of being AI first.

I distinguish this from AI enabled because the teams at AI first companies look different, right? They're hiring AI scientists. They're often collaborating with academia. And you know, the actual value creation at the company will of course come from building a business, but it first comes from developing a technology that is proprietary and unique.

I'll caveat that with saying, of course, all moats are fleeting in AI right now, and so you're in many ways on a kind of hamster wheel trying to constantly keep up. And the company will need to finance itself accordingly, often raising a lot more money in order to do that. Let's kind of juxtapose that with AI enabled companies, which are companies that are much more in the business of deploying AI.

Whether it's something straightforward, like a large language model, or kind of patching together a system, or thinking about each individual model as an instrument in a broader orchestra. [00:37:00] The pejorative for this, although I'm not endorsing it, is GPT wrapper. Right. Which is, yes, that's the pejorative.

Yeah. And so, these are, you call them AI enabled companies. So, I'm curious if the dynamics recently between like DeepSeek and OpenAI, has this changed

your view of where the, I think nothing is durable as you're saying, you know, nothing's infinite, but where the semi-durable moat is like, has it switched at all from the sort of AI first to the AI enabled companies?

Is syndication having distribution, having people enjoying and using your products and, like, knowing your brand, is that the bigger, like, is, does that have a relative advantage after just the last year of open source and DeepSeek, or still AI first has its place? The greatest moat in health care and life sciences is distribution, because the way that enterprises adopt technology is far slower, oftentimes in hospital environments.

You know, once technology's adopted, we see that technology [00:38:00] flourishing in that environment for the next decade. So, if you can weather the storm, the 12-to-24 months, it often takes to sell into these large enterprises and get on the inside. That in itself is a tremendous moat, and that's where AI enabled companies make their mark.

I think on the AI first side, there's still tremendous opportunity for those companies, but they need to be quite specific in where in the stack they decide to invest in innovating at the methods layer. And I think we can all agree on this call post-DeepSeek, but even we could have seen this coming a couple years ago.

That's not going to happen at the foundation model layer in LLMs in biology. Will it happen at the foundation model in PLMs? I'm not too confident. Right. But are there very specific areas. Whether it's in clinical entity recognition or if it's in a particular kind of variant interpretation protocol or, you know, you think about getting a level deeper and becoming more nuanced about the problem you're trying to solve.

Designing epitopes to mitigate immunogenicity as one of our companies, Seismic Therapeutic has focused on. And so, I think the moats will come from these more niche applications [00:39:00] for AI first companies, and then they also have the joint burden and task of also nailing the distribution.

Can ask a follow up to that, Morgan? This is something I feel acutely both in my academic and entrepreneurial life. One thing that has been a truism for the last like decade is that general approaches beat problem-specific approaches. And so, like when I was working on the USMLE, I had this little LSTM that I was trying to get to answer medical questions.

And I was trying to be cute and clever about the data that it had and like how to represent the knowledge. I mean, it turns out that like next token prediction completely just bulldozed that problem. The problem that I needed to be solving was, like, next token prediction very well. So what lessons are there for folks starting companies?

How to pick problems that won't be bulldozed by the general solution. 'Cause that to me feels like the great existential risk for all of these companies is that you won't get GPT-5 or GPT-6 or 04, 05, 06 out of existence. Like what distribution is a great moat for there, but that if [00:40:00] you're starting something new that may not be an option.

Like how do you think about that core existential problem? Well, if I knew where the generalist solutions would be the next 12 months, I would be the smartest person on the planet. I think it's very hard. A quote that I love that Daphne Koller often says is, it's very hard to predict what's going to happen in a short window of time before you, when you're riding an exponential curve, right?

That's like a very hard problem to solve. What I would say is I would focus on the areas that are more data sparse. So, when I think about what's happening on the clinical side, I fully agree with you. I think that generalist systems will continue to perform extremely well, and I would not be betting on a clinical-only first LLM unless it's in a very, very deep, you know, domain category.

Like a particular segment of rare disease where you fine tune something to be quite performant there. I think in biology there's much more opportunity because we just don't have the equivalent of the Internet for biological data available to us to train these models. And so, I think in the near term, I can't say what's gonna happen in five years.

There will still be a plethora of AI first companies that are able to [00:41:00] maintain moat in biology. I do not think it will be at the PLM protein language model layer, but there will be other segments where this is possible and I'm investing a lot in those areas. Cool. Awesome. Are we ready to move on to the lightning ground, Raj?

Yeah. Alright, Morgan, are you ready for the lightning ground?

Hit me. So, I think we asked the inverse of this question earlier, but the first lightning round question is, what has being a venture capitalist taught you about medicine? Interestingly, being a venture capitalist has taught me most of what I

know about the business of medicine. So, when I'm on a mitochondrial medicine service at a children's hospital and I try to order a particular genetic test, and I hear from the billing department that that ain't happening. I would say that my work in venture and deeply understanding the reimbursement structure of the industry has afforded me, I wouldn't say an empathy, but an understanding of why certain things [00:42:00] happen the way they do.

And I think recognizing that has been a superpower in terms of, you know, you're making clinical decisions and wanting to do so in a grounded reality of the way that we fund health care in this country. So that's probably been one of the more specific takeaways. I also think that adventure has taught me a lot about what scales and what doesn't.

If you think about the asset class, there's many, many ways to fund a business in 2025, including not taking venture capital. But what venture capital is particularly designed for is technology that scales nonlinearly. I actually think we've gotten this kind of muddled in our current environment where we use venture capital a lot for companies that are only taking operational risk, which is fine.

It's actually quite an expensive asset to utilize for operational risk. It's far better positioned, given how expensive it is for a company to utilize venture capital for taking some sort of scientific, clinical, or technological risk. And so, I think in terms thinking in that regard, bringing that that mindset of what actually is capable of nonlinear scale and thinking about that in medicine [00:43:00] shines a light on where the opportunities will be.

Awesome. Morgan, which is a harder job, medical student at Brown University or vice president at Bessemer Venture Partners. Can you define what you mean by harder? This is my lightning round. I'm flipping it back to you. Uh, more. Yeah. Let, let's say, let's say harder is, uh, one that leads to more stress in your mind.

Definitely medical school and medical training. I mean, I would describe the highs in medicine are way higher than the highs in in business for me personally. And the lows are far lower. The stakes are significant, right? In business, if you lose money or something doesn't go as planned, like the business shuts down and it's sad, but those people likely go on to live their lives.

And in medicine, we're not always afforded that second chance. And so, I would say that's one of the harder parts about being a medicine. The other hard

part, and maybe you get this from my background, is there's a lot of structure and a lot of bureaucracy. And as someone who loves autonomy and freedom, I [00:44:00] often find myself feeling a bit defeated by the appending pressure of needing to conform.

But again, as I mentioned, I've been fortunate to find people who will kind of help me break down some of those barriers and silos to stay sane. Awesome. Next question. If you weren't in medicine and venture capital, what job would you be doing? So, I originally went to school to become a classics professor.

I love Latin and ancient Greek. And I know Emily, who was just on the podcast, shares that. And actually Zak shares that. I don't know if it's like a weird pediatrician, informaticist, classicist thing we have going on here. I would've become a classics professor. I love ancient Roman history, and I love to study the literature from those time periods.

The reason actually why I didn't end up becoming a classics major is because my dad told me he wouldn't pay for college if I didn't get a science degree. So that's why I made up my computational decision sciences. And I did end up integrating a lot of classics in that, including discovering the field of computational classics where I was using NLP methods on Livy.

So that's a [00:45:00] fun area that I dig into sometimes in my spare time. Nice. What's your favorite piece from the classics era? I'm a big fan of a lot of Catullus's work because it's funny and, I think lighthearted and kind of punchy. So, I think generally his body of work is interesting to me. And I'd also have to say, I know it's kind of a canned answer, but the Aeneid is something that I return to often. Forsan haec olim meminisse iuvabit.

I know my Latin pronunciation is terrible, but it is actually a quote that I carry with me often when the going gets rough. Which is, you know, perhaps sometimes it will be meaningful to have remembered even these things. Awesome. Excellent. Morgan, if you could have dinner with one person, dead or alive, who would it be?

I knew you guys were gonna ask me this question. Can I say the both of you? There's a few, there's a few people. We have had dinner before, Morgan, you have to say something you haven't had dinner with before. Well, I'm subtly mentioning this because it's been a while, so I'm hoping you'll take me to dinner soon, maybe for, for graduation, but I think.

There's a couple people that come to mind. Bud Rose, who's the founder of UpToDate, unfortunately, who [00:46:00] is, who is no longer with us, I think was just a remarkable human being in terms of thinking about where informatics and medical knowledge intersect, and how, frankly, one of the first people to grapple with how human physicians and computational systems can and should interact.

And I really wish I could ask Bud, Dr. Rose, what, what he thinks about our current environment with language models and where we should take this stuff. I think the other person, I deeply admire who's still with us, is Ted Nelson, who is a computer scientist and great computational thinker who designed Xanadu, which was kind of an alternative way that we could have organized information on the Internet.

And I think he had a lot of really great interesting ideas for his time. Everything from, you know, micropayments for creators to thinking about bi-directional linking and hypertext formats, which I think still in intrigued me as a way that we could have organized information at scale.

Cool. Awesome. Final lightning round question. What is your P of doom? Can you define that? So, P of doom is this [00:47:00] probability of doom that where the doom is specifically a consequence of AI run amuck. So essentially, what's the probability that AI becomes sentient and kills us all? Fifty percent. Over what time horizon? A century.

Okay. It is a safe answer. Okay. Okay. And this is why I'm investing in AI.

Wow. Alright. Yeah. So, we have a couple like big picture questions to wrap up with Morgan. So, you've been great so far and just wanted to zoom out a little bit to wrap up the conversation. I've listened to a lot of investing podcasts and read a lot of investing books, and it seems like identifying successful founders is like one of the things that VCs think they do very well that differentiates them from their peer group.

So, when you are looking for companies to invest in, what are the things in the founding team that you look for [00:48:00] that are strong predictors of success? One of the most important things I look for is what I describe as the learning rate of the team. So, if we have a call on Monday, and then we catch up again on Friday, what have you learned about your business, your customer base, your technology, and really tracking that curve over time?

And when I've backed founders who have that kind of just off the charts learning rate, I think that they, that compounds it. It enables them to frankly get to the truth, the objective truth, and make decisions in a highly informed way. In health care and life sciences, I think it's easy to often look for the domain experts in a space, so the smartest person in this field of biology or in this specialty of cardiology.

And I think those people are really important to companies either founding wise or from an advisor perspective. But I also look for people who are students of what's happening outside of health care and life sciences because so many of the great ideas come from these other industries that are adopting [00:49:00] technology much more rapidly than we are. And the transfer learning across industries is quite impressive and can really move the needle forward.

So, I look for that as well. I think the reality is that making an investment in a company is like getting married and, in fact, it's often harder to unravel than a marriage in the U.S. And so, I'd underscore the point that team is everything. But I'd also underscore that a few other things are top of mind.

I think one is timing. The why now, right? We talked on this podcast about a lot of ideas in AI and medicine that have been percolating for decades, and yet we're grappling with them in a real way in 2025 because there was some unlock and some change. That's forcing us to revisit this concept with a new lens.

And so, when I think about technology evolving and I'm making an investment in an evolving technology, I wanna believe that that inflection point is happening during a period of time that that company can be funded to realize it. And you wanna be sufficiently early to where, as people joke, if there's a market map of this category, you're way [00:50:00] too late. But you also don't wanna be too early.

That the commercial opportunity for that technology and that company doesn't materialize in what I describe as a venture timeline. Yeah, being early is often the same as being wrong. Spot on. In many cases, so. Period. Yeah. Morgan, what's your vision for the future of medicine? How do you see medicine itself?

The practice of medicine by physicians and maybe medical school, too, evolving over the next, let's say five or 10 years. I'll hit the medical school point first. I hope we move beyond this approach of information inundation and memorization, and more towards a system of thought and learning that forces critical thinking.

We're not there. And I think the hard part is that learning medicine is learning a new language. And so, we need to balance learning that vernacular and vocabulary with forcing people to actually think. And I would argue, think with AI and teaching that in an [00:51:00] intentional way. In terms of the field of medicine, we didn't touch much on this and maybe we'll revisit it on a future pod, but I've been completely pilled by the field of genomics and multiomics and where we are headed in that realm.

I had the opportunity as part of my rotations to spend time at NIH for three months, where I joke that I saw how medicine could be practiced if we actually funded things that were right for the patient. And I acknowledge that the NIH people joke, it stands for "not in a hurry" because things don't happen there in a super-fast or always efficient format. But the people who work there, the physician scientists who are leading trials there are truly empowered to think objectively about what is best for this patient.

And the institution is funded to support those decisions. And in a dreamy world, I'm excited about the genomic learning health system. And I'm excited about this notion that a larger percentage of people are going to undergo molecular testing in the next few decades, whether that's [00:52:00] actually a exome or, or a genome, or even more specific testing that helps us understand how protein expression is happening in a particular disease state. And I'm excited about our ability to marry that information with what we now have in terms of clinical data in electronic health records.

And so, if you think about what the Genomic Learning Health System stands for, it's this notion that people undergo molecular testing and then they bump into the health system and they receive care. And as they receive that care, we're constantly updating their phenotype and describing the new things that have gone on with them.

And as we have that information available to us, we can then go back and update how we've interpreted their particular genomic and multiomic context to have a greater precision around how we think about their health and wellbeing. I think this will be kinda first manifesting in areas of rare disease as it already is.

But I'm hopeful that we get to a point, as we saw even just recently with the Mayo Clinic Tapestry study, that we think about this in terms of a broad-based population approach to multiomics and genomics. That is both diagnostic, but as we know, increasingly [00:53:00] preventative and interventional. In so many ways, oncology has paved the way in demonstrating how a deep molecular understanding of disease unlocks both novel diagnostics and therapeutics.

And just think today, naturally, if one were to be unfortunately diagnosed with cancer, you often not only know the cancer type, but also the driver mutation underpinning it. Yet, when we look across other specialties and disease states, that same level of molecular precision is often missing. Think heart failure, you might be told you have garden variety, heart failure.

And so, one thread I'm hoping to pursue over the next decade is, I guess what I would call, the molecular of other specialties, cardiology, neurology, nephrology, metabolism. Because I see tremendous opportunity to bring multiomic insights to each of these fields. What do you see as getting in the way of that, that future? Payment models.

What gets in the way of every great thing in health care? I feel like we don't have the payment models to support it. If you're a commercial insurer today and your average patient is going to be with [00:54:00] you or your member, let me use the right terminology, is going to be with you for two years. Why do you care about paying for their molecular testing to prevent some disease that might manifest for that person in a 40-year time horizon?

Why do you wanna be on the hook for the gene therapy or the siRNA therapy that would, in theory, you know, cure that patient if you're not going to be responsible for the cost on the line? These are the real questions we have to grapple with. Again, if people take away anything from this conversation that we have, I hope they hear me when I say:

the science, the computation, the technological innovation is no longer the barrier and maybe it never was. Maybe it was a fallacy that, that it was, but just given compute data access, where we're headed in terms of the arc of progress I mentioned, the tech is not the problem in health care and life sciences.

The problem is the business model, the economic model, the way care is paid for, and the incentive structure underlying that. And more and more technologists, I hope will spend a little bit of their time solving those problems, because I think [00:55:00] there will be an outsized benefit for us all. Well said.

Yeah, well said. One last question, Morgan. So, I think again, like you, you've been able to strike this like super interesting balance between like being a leading VC, investing in the space, uh, being a future physician for aspiring medical students who hear this and be like, man, that sounds like an awesome life.

What can they do to become the next Morgan Cheatham? No one should strive to be the next Morgan Cheatham. I'll start there. I'm a big believer that everyone has a unique contribution to make, and in fact, often when I meet people for the first time, I make a point of trying to figure out what that might be for them.

What I can say though is there are a few principles that have guided me in pursuing my journey thus far. The first and most important is I would say, clarifying your vision. Everything else will stem from that. What strongly held beliefs do you have about the future state of the world in technology, biomedicine, or whatever domains you happen to traverse? And how do you position [00:56:00] yourself to contribute to that over time as it unfolds?

For me, I talked about having a vision of what medicine could be like where I described the genomic learning health system or the multiomic learning health system, and I'm personally taking action by spending time in AI research, clinical genomics and biotechnology investing. All interdisciplinary fields to contribute to that future state.

The best way to clarify your vision though, I would say, is to lean into your innate curiosity, and I would say most of my major regrets in life so far stem from not going all in on something that I found interesting or felt curiosity about. Whether it was playing around with Bitcoin in 2015, everyone has their story, or dabbling in genomics and undergrad, but putting it on hold for about a decade only to revisit it now.

So, I would say following the threads of your curiosity is key. Even if your interests seem niche, esoteric, or intersectional and complex. And in doing so, finding your people along the way is also critically important. And then I would say the third and most tactical piece of advice would be parallelizing your career when [00:57:00] possible.

For me, this happened on accident, but it's obviously manifested in pursuing venture capital and medicine simultaneously. One informs the other, and on the other side of it, I almost feel like a toddler who was learning to languages. It may take me longer to start speaking, but once I do, I am inherently bilingual.

And so, it can be difficult to parallelize your career though in fields like medicine, which are bureaucratic and hierarchical, and I think we have to acknowledge that. But if you have that strong vision, you owe it to yourself to at least ask what might be possible. Another quote I love from the classics, uh, is faber est suae quisque fortunae, or a person is the maker of your own fortune.

That’s from Appius Claudius Caecus. You just have to ask. And the worst thing that can happen is someone says no. And in other cases, if you’re in a field where you don’t need to ask, then just do. Awesome. I think that’s a great note to end on. Thanks again for joining us today, Morgan. Thanks, guys, for having me.

This was a blast. Yeah. Thanks so much, Morgan. This was great. [00:58:00]