

AI-GR P21 08.09.24 Vijay Pande

[00:00:00] And so many faculty are involved with startups and students from Stanford are involved in startups, some very famous ones. And so, I wanted to be in that ecosystem and I was. I was involved with being on scientific boards, um, being board directors and companies came out of my lab and so on. And that was exciting.

The part that I was gravitating the most to over time was venture capital. And so very much finding the fit was a key part of my mindset. And it took some time for that, but that first meeting with Marc and Ben, I think their mindset of really thinking about the future, being very intellectual about that, you know, how to put it.

I think to, they, they reminded me when I joined the firm in 2015 the way, you know, Stanford was in 1999. That the real desire to really put an impact in that landscape, but still building and still growing. And so, it was something where tech was a key part of the mindset for driving A16Z.

And Marc and Ben, this was something I think that was a key part of, could see the value in tech and [00:01:00] life sciences and tech and health care. So, I came in 2015 and launched our first life science and health care fund, first bio fund. And so, we're now on our fourth fund and now it's about over 50 people on the team.

Welcome to another episode of *NEJM AI Grand Rounds*. I'm Raj Manrai and I'm here with my co-host, Andy Beam. And today we are thrilled to bring you our conversation with Dr. Vijay Pande. Andy, this was really a lot of fun. You know, Vijay has this amazing journey from being a Stanford professor of chemistry, structural biology, and computer science, to now, uh, a general partner at Andreessen Horowitz, where he leads investments in health care

and life sciences. It was really fun to sort of pick his brain about that, uh, transition from academia to industry. And I have to say, I thought he was really balanced and thoughtful in articulating why someone might want to stay in academia versus, uh, versus trying to impact health care from outside of academia.

All in all, this was a lot of fun and really, [00:02:00] really great to get a chance to talk to him. I really enjoyed this conversation. Vijay has these two sides of his personality. He's a general partner at Andreessen Horowitz. He's done

amazing work leading their bio fund, so he opened up this entire new investment area for A16Z.

It has done a lot of amazing work there, but it's still really easy for him to tap into this, like, scholarly academic side, too. So, he can give you very thoughtful, very grounded answers on scientific questions and how they think about investment in these really hard areas. Like investing in health care is a very hard thing to do.

And he's able to articulate this very first principles, very well motivated, investment thesis. So, I really loved talking to him because you get to learn a lot. You sort of get some inside baseball knowledge about how these deals get made. But he's also able to pull it back and give you these very thoughtful answers to very hard questions.

The *NEJM AI Grand Rounds* podcast is brought to you by Microsoft, Viz.ai, Lyric, and Elevance Health. [00:03:00] We thank them for their support. And with that, we bring you our conversation with Vijay Pande. Vijay, thanks for joining us on *AI Grand Rounds* today. We're super excited to have you here. Absolutely. Happy to be here.

Thank you. Welcome Vijay. So, this is a question that we ask all of our guests to get started. Could you tell us about the training procedure for your own neural network, how you got interested in AI and what data and experiences led you to where you are today? Yeah, so here I ended up going kind of far back in that I started programming at a pretty young age, like when I was 11, and I was just doing things that were fun.

I ended up doing things a little more seriously, later on when I was, 15. I was in Naughty Dog Software, a computer game company with just me, Andy, and Jason, the founders, and so on. And so that was real, but between then I poked around a lot of things. And I don't know if you guys remember SHRDLU, I forget how it's pronounced.

The Terry Winograd AI thing where you could have the computer pick up this and that, and it was very old school AI, but that got my [00:04:00] excitement. And I coded some probably very simple-minded version of that early on. And then later on in college, neural nets start to get very hot and neural computing was getting hot.

So, this is like 1990 or so. And so, I was coding up neural nets and doing the math for Hebbian neural nets. And the thing that most of us saw at that time,

there was a lot of excitement, but neural nets couldn't really do that much because, as we know in hindsight, they were just single layer neural nets that they hit a limit.

And so, I think most of us put things down to do real things. It seemed like a toy and then things changed, a little later. Yeah. I remember when I was in grad school. The saying was that neural nets were the second-best way to do almost anything. And sort of the implication there was that they weren't the best way to do anything.

So, what was the serious thing that you picked up after you decided that neural nets were still in the toy phase? Yeah. So, a lot of what I did from grad school, postdoc and early days at Stanford were physical simulations. So, you know, go down to the physics of atoms to understand, how molecules work and so on.

The real benefit of physics is that they [00:05:00] you're hopefully not overfitting and you have real generalizability. The fantasy from the days of Newton is that you can watch an apple falling down from the tree and then you can predict planetary motion, something that seems very far from overfitting, right?

That you really understand the fundamentals. I think a big concern with ML was always overfitting, especially in an age where we didn't have really much data. And so that was a natural thing to do, but I think, probably 2012, 2013, 2014, things start to change both in terms of data, but also many of us,

and this was a key thing on my mind at that time, we're pushing for more compute. And so that combination really changed things. And correct me if I'm wrong, you were faculty at Stanford for about 10 to 12 years in this lead up to the sort of deep learning renaissance that we're in now. Yes. I started in 99 and I started at A16Z in 2015.

So, 15, 16 years. And I was primarily in chemistry, but had appointments in computer science, structural biology, and I was chair of biophysics. So, a lot of the work I did was at the intersection of those different fields. So, I think we're [00:06:00] going to circle back to this at some point, but could you tell us a little bit about the decision making

for leaving your prestigious, presumably cushy academic job to transfer into the exciting world of investment and venture capital. Yeah. So, even from the beginning, very much the rationale for me to go to Stanford, one of the things I

really got excited about Stanford was to be in the middle of the startup ecosystem.

And for those of you that aren't familiar with just the geography of the area, Stanford and Sand Hill Road, which is the home of venture capital in the area, and to some extent the world, you could in principle walk, but it's like a five-minute drive between the two. And the ecosystems very, merge very much.

And so many faculty are involved with startups and students from Stanford are involved in startup, some very famous ones. And so, I wanted to be in that ecosystem. And I was involved with being on scientific boards, um, being board directors and companies came out of my lab and so on, and that was exciting.

And being a part of that ecosystem, the part that I was gravitating the most to over [00:07:00] time was venture capital. And it was something that I'd been thinking about, but you know, venture capital firms are about as different as people are different. And so very much finding the fit was a key part of my mindset.

And it took some time for that, but that it'd been on my mind ever since the beginning. And so, what was it about A16Z specifically that made that sort of VJ market fit, feel so good? Yeah. I think first meeting with Marc and Ben, I think their mindset of really thinking about the future, being very intellectual about that, how to put it, I think to, they reminded me when I joined the firm in 2015, the way Stanford was in 1999. That the real desire to really put an impact in that landscape, but still building and still growing.

And so, it was something where tech was a key part of the mindset for driving A16Z and Marc and Ben. This was something I think that was a key part of could see the value in tech and life sciences and tech and health care. And that combination was very, very exciting to me. And I think a lot of other venture capital firms, I think took a more traditional mindset [00:08:00] that tech and life sciences are in tech and health care.

We're not going to mix. And again, correct me if I'm wrong, but I think you came there to start their biotech fund that essentially you originated that sort of whole arm of A16Z. Yeah, yeah. So, I came in 2015 and launched our first life science and health care fund, the first bio fund. And so, we're now on our fourth fund and now it's about over 50 people on the team.

And that was part of the other excitement was to actually really build something. And that's something I very much enjoy, but especially to build it in

this new mindset that it would not be looking backwards the way traditional funds are built, but really looking forwards. Okay. So, I think I'm going to take this in chronological order here.

And can you just talk to us about what your biotech investment thesis is? So, I think that a16z does have this unique perspective, like how do you understand value in this area? How do you understand good companies? And I think also like, how do you identify good founders? Because I believe that's a huge component of investing in successful companies.

Yeah, so this is something where I think there's, this was probably more unusual 10 years ago when I started the fund and it's fun to see [00:09:00] that I think the rest of the mainstream investing, I think has come a long way, but, especially when we launched the first fund, the concept of technology having an impact on life sciences or health care was pretty radical and even debatably heretical.

I remember people telling me that machine learning or AI would never have an impact in drug design. And that, okay, we've seen this before and people have talked about this before. And that was a key hallmark, but I think I don't want to over rotate on the AI side. I think the true nucleus of what tech brings is AI is an example of, but not the only example of, I think the true nucleus is a concept of engineering.

That we're going from bespoke artisanal discovery to something that's designed, engineered. And when we talk about engineering, what we really like about engineering is the fact you can make it 20% better year over year. And that sounds actually relatively modest, right? But if we compound that

over decades, that's what Moore's law is. That's what the cost of genomics going down exponentially is once [00:10:00] you can get something improving with that regularity and that regularity comes from an engineering mindset. Now things change. And so, the idea of bringing engineering to life sciences and engineering to health care, uh, that was a really foundational aspect.

I wonder though, if there's an important difference here that you had to grapple with, especially in the early days, I think this is appreciated now. So, you mentioned like Moore's law for the cost of sequencing. But in drug development, we often talk about Eroom's law, which is just Moore's law spelled backwards.

And it points to the fact that actually drug development is getting longer and more expensive. And so, you're often working on radically different time horizons, especially for return on investment when you're thinking about a biotech than you are for a traditional tech company. So, it was like some type of expectation adjustment, something that you had to do in those early days.

Again, now I think this is appreciated. There are some similar theses that you could identify between tech and biotech, but the timescales are just so different. Yeah, the funny thing is like the time scales for IPO for biotech is faster. The time scale to revenue is slower, but the revenue ramp can be really [00:11:00] fast.

To put it in tech words, like the product market fit is perfect, right? If you have cancer and this is the only first in class, it's going to do well, right? And so, it's just, it's a very different mindset, I think. The key thing that, um, we very much have tried to stay away from, though, is something that has a lot of single asset risk, something that's really hard to predict.

And the benefit of a platform company is that if the first asset doesn't work and the platform is really productive, you can have many beyond that. In practice, I think one of the big challenges, even for modern platform companies, is that once you have that first asset, that's where most of the value is.

And so, there's a very strong temptation to pile behind that asset. And so, building a true platform company, modern day Genentech or Amgen, or even more recently, Alnylam or something like that, those are hard to do. They're truly hard to do because of the temptation to pile into the asset. Yeah.

There's always a tradeoff between exploration, exploitation, and all of these things. I wonder if you could tell us about a couple of companies that you've been involved with that you think sort of embody [00:12:00] a successful platform approach. A natural one, and this is a fun one investment and investment I'm on the board of is a company called insitro.

We've had Daphne on the podcast. So, there you go. So, Daphne is a true OG in AI. And so, um, insitro is her vision for bringing AI to drug design. I think she's primarily, especially tackling the challenge of using AI to unravel biology and especially human biology. That the fact that drugs work so well in mice and so poorly in humans is obviously the fact that we can run experiments on mice.

We have lots of data on mice and it would be naturally unethical to do the analogous experiments on people. So how can you do that? And it's a perfect

role for AI to gather all that data and to build models that are predictive. They're not perfectly predictive, but they're way more predictive of human beings than a mouse would be, is the bar.

And so, for better or worse, the bar is actually not crazy high. And she very much wants to build a true platform company where there's not going to be a single asset, multiple assets. And what's intriguing about what she's built is that the ability to [00:13:00] understand biology, has implications for the whole drug design process. For not just finding targets, but you can imagine what you could do with that for trials and beyond.

Got it. Got it. So, I'd like to know, so we've gotten you through biotech investor, now health tech investor, I think is your most recent arm something that I think is a little closer to our wheelhouse here on AIGR. And I think something that even though I'm deeply involved with this, I have less clarity on than on biotech.

As you pointed out, the markets are exceptionally clear and well defined. It's very easy to understand value pools in biotech. In health care, it's not even always clear who the consumer is. You know, you've been on record as saying that the biggest company in the world, in the future, will be a health care company.

Could you walk us through sort of the investment thesis for health care companies, especially as they interface or interact with AI platforms? Yeah, so the health care part of the fund is a key part of it. I think one of the unique aspects of the fund as first constructed is that we would do [00:14:00] both.

That typically life sciences and health care investors and tech investors, those are typically three different firms, if not three different funds. And I think one of the rationales for especially including health care is that I was expecting then, and I think we're very much seeing now, a blending of these areas.

That modern drug design companies, modern biotechs, have to really think about care delivery in their choice of therapeutic areas and so on. And then if you're a care delivery company, the things that are coming down the pipe in terms of new pharmaceuticals radically change things. I mean, GLP-1s are just one of many examples.

So, in terms of health care leading to a huge company, I think the opportunity here is that we've yet to see a health care company that is built the way a consumer tech company is built. And if you think about the impact of these

companies could be way bigger than any FAANG company, way bigger than Facebook. In terms of what you care about, it's hard to put something more important than your health, or your parents' health, or your spouse's health, or children's health, and so on. The challenges, I think that's what you've alluded to, is [00:15:00] that the current health care system is complex, and actually who pays for what and how things work is fairly complicated.

Health care companies with a consumer mindset, and that could be a couple different things. It could be, in principle, who pays and that's one thing we could talk about, but I think the key thing about a consumer mindset is from a tech perspective is changing consumer behavior. And one thing that tech companies are particularly good at, and tech is particularly good at, is changing behavior.

And for better or worse, you think about all the things that I do, my colleagues do, the world does to try to impact life sciences, in the end, there's many things on care delivery that you could do that have as big or bigger impacts. The quote that's usually brought up is that even curing cancer would lead to three years of added lifespan.

And cancer is something that, we were very excited to have a huge impact in, and it's a horrible disease, but I think this really points to the fact that there's many other things that we should also be paying attention to. And that's, I think, part of the opportunity for health care is to think about how can we motivate

people in consumer oriented [00:16:00] experiences that we have in tech to be able to actually really take control of their lives. And we've made several investments in this space as well. And it's still early, but I think even the early days have been fairly exciting. Could you talk about some of those investments in this space?

I think Hippocratic AI is one that comes to mind and some of the others. So maybe you could tell us about that and how you see this, the space evolving. There're a couple different categories we could talk about. One is AI plus health care going, providing to the health care system.

And Hippocratic is a great example of that. So, what Hippocratic does is that it provides essentially AI nurses and they've been very clever in terms of going after areas that could have a still a big impact on the health care system, but for now avoid challenges such as diagnosing or prescribing drugs that a doctor would do.

And so, and also actually one of the big surprises is that I think we'll see in hindsight, when we look back at COVID, that as much of a mess and horrible COVID was, there were actually some interesting health care tailwinds that were created during [00:17:00] COVID, especially in terms of virtual-based care. And so, anything that you do with a nurse on the phone or on a Zoom,

in principle, you could do with AI. And right now, Hippocratic has AI nurses that one can talk with. They can be used for prep and other procedures. And if you think about one of the crises that we have right now is that there's a huge nursing staffing crisis. This is a very natural role for AI to play where it can do something that could have a huge impact.

Maybe we'll get into clinical on time, but not maybe immediately, but still today drive a lot of value. And I think about primary care as well, too, right? Even situated, as Andy and I are here in Boston, where we're attached to a major medical school and university, many hospitals, I think it's still hard to have time with your physician, right?

Just to actually discuss your care, go through the shared decision making. What is your personal values? What's, what are your goals? And I think this is why there's a lot of excitement amongst a lot of us that AI might help actually bridge some of this gap and fill this gap [00:18:00] in counseling and talking to patients and thinking about their care.

And ideally in a preventative and forward-looking manner as well. With Hippocratic, I think LLMs, obviously now these are, you hear about them all the time. The concerns are hallucinations, confabulations, safety, bias, integration with the existing health care records, and context and all that.

We're very familiar with this sort of set of challenges. What do you think is the sort of primary thing that Hippocratic is solving? I mean, there's so many ways to kind of attack this problem, right? And you don't want to do everything at once, but what is their bet on where they will stand ahead, where they will stand apart in the next few years?

And what are they really trying to solve amongst these, these challenges around LLMs? Yeah. So, the challenges you're talking about with LLMs are not unique to Hippocratic. They're true for anyone using LLMs. And the thing that's different about health— They're true for humans. There are many of them are true for humans as well.

Well, so that's a very interesting point that's worth getting to. I'll get to that in a second. If you think about hallucinations, LLMs, if you're doing [00:19:00] this for poetry, that's not a problem. That might even be seen as creative. If you're doing this for art and a cat has six fingers or I have six fingers or three fingers.

That's just artistic license, not a big deal. If you go to, if I have a surgery and I come back with three fingers, I'm going to be really, really pissed, right? That's going to be a problem. And so, the health care, we have to get it right. And we don't have that room for creativity. And what we're seeing today broadly is the use of LLMs as user interface and as a means for answering questions.

But on top of that, you'll have a mixture of experts of many types of models. Some of which may be LLMs, some of which may be more traditional machine learning and so on. That would make sure that if the LLM is hallucinating, that's something that we can just be on top of. And in a sense, this may be a poetic analogy, but I think it's deeper than that.

It's essentially a care team where you have one member that's good at one thing, another member that's checking on the results going out. And that team approach is something that can be done now with low latency, which is kind of amazing. And that you're getting this team of experts speaking to you [00:20:00] when you're talking to an AI.

And that part's really unique to today. Raj, I don't know if you have any more questions in that line. I wanted to ask like one more sort of big picture health care question before we go to the next section. So again, the biggest company in the world of the future could be a health care company. However, we're currently spending about a quarter of GDP on health care costs in the U.S. and your friend and colleague, Marc Andreessen has this blog post about how the cost of the TV has gone way down.

The cost of health care has gone up faster, multiples of the price of inflation. So, I guess a question that we often ask guests on this podcast is, like, will technology make health care more affordable and better, or is it just going to increase the spend? Is this like a new sort of spending mechanism on health care?

How do you see that playing out? And if you see it driving down costs, how do you reconcile that with sort of Marc's take on this? Yeah, I think the fundamental resolution of this is that health care as sick care, basically dealing with you once you're sick, that is inelastic from an economic point of view.

What would you pay to save your spouse? Everything, [00:21:00] right? There's, I don't want another spouse. So, I would give everything for that. And if that's the case, that's something that's inelastic like that is going to be hard and you'll just pour whatever money into you can into that and technology will create more options, but it doesn't change the fundamental elasticity of that.

Where I think this gets interesting is something that is easier said than done, but where I think the future of health care really lies is not in sick care, but in keeping us from getting sick. And that's something that's very natural for tech. That's something that obviously will, will bend the cost curve on the sick care side and will change things.

That's where I think tech has to go. And we're going to use tech to develop new therapeutics for rare diseases, and cancer, and so on. And that's going to be part of that sick care cycle. But I think where it gets really interesting and where the curve gets bent is ideally, you never go there. And so, from a health care delivery point of view, that's getting on top of things, that's running diagnostics.

Companies like Function is a great example of that. If you're a care delivery company, that's avoiding admissions or readmissions, and [00:22:00] using that as your North Star. Trying to avoid the sick care system as much as you can. That's the intellectual goal. How to get that done is the real challenge, and something that I think it's easy to poke at, and I think there's a lot of work to do, but that's, I think, the real goal.

Is there an opportunity, from what you've seen just to deliver care more efficiently? So even if you look at GDP per capita spent on health care versus mortality, we're awful by that metric. So even outside of like health care versus sick care, can we just like be, deliver health care more efficiently?

And you mentioned some tailwinds of COVID for telehealth and things like that. Is there just a way we can like, even if we're not keeping people out of the hospital, once they're in the hospital, not spending \$500 on an aspirin or something like that? Yeah. So, this is where I looked at companies like Devoted, Devoted Health which, is, has a Medicare Advantage plan and dual plan and is, but more broadly a health care company with both as an insurance company and as a provider. What they do is really handle all the logistics and thinking about the right care at the right place at the right time.

If you think about like maybe [00:23:00] a non-medical analogy would be something like Amazon, and the cost of things before Amazon. Like, Amazon is

an amazing logistics company and is able to drive the cost down through a marketplace as well. So, in that sense, there's no silver bullet algorithm that makes Amazon, Amazon.

It's about that mindset everywhere. And Devoted has that very much as well. I think their ability to bend the cost curve, I think is a great example of that. And naturally avoiding readmissions are on their mind and handling things ahead of time is something they're incented to do as being a payer and a provider.

Great. So, Vijay, I want to transition us a little bit and we want to spend a few minutes just digging into academia versus industry. So, Andy mentioned this, you know, you were a tenured professor at Stanford. You'd made it and you decided to leave that position to join A16Z. And I think you've clearly done very well there, but maybe you can take us into your mindset.

I think this was back in 2015, your mindset behind that decision. And maybe I can even ask more provocatively, thinking about [00:24:00] today, where should a smart person who's interested in impacting AI for biomedicine be? Academia or a startup and why? Yeah, so this is a question very dear to my heart and I think there's a couple different aspects to this.

So, first off, I think just purely intellectually, forgetting about any sort of macro arguments or anything like that, I think the first question is, what is the timescale for the impact that you want to do? So, what you want to do is going to require a decade or two decades. That's not something you do in a startup.

That timescale is really only available in academia. And so, when I started 99, I wanted to really make computational drug design a reality, uh, that was something that wasn't going to get done in three years and it wasn't ready to roll out to design drugs. That's something that companies are doing today.

And, even Genesis, spun out of my lab is a really beautiful example of that, that needs some time. And so that was the reason to be in academia. And the other reason to be in academia is, let's say you don't even know what you want to do. You want to explore. And I think there's [00:25:00] very few places where you can just go and explore.

And you think about some of the most fortuitous discoveries like CRISPR or something like that. That's something that I think could probably only be found through a love of basic research and exploring and finding things and seeing what happens. And it's impossible to even put a timetable on those types of things.

And then finally, it's your mindset, right? Some people, I think, have much more of an academic mindset. The academic mindset is, what do you get excited about doing? Do you prefer to read *Nature* or *The Wall Street Journal*? That type of thing. And actually, the funny thing is— We generate PDFs.

Yes. We generate PDFs, yeah. Yeah, there's also that. It's like, what do you, also, what do you hate less? Yeah. Proofreading papers and, and so on, I think you'll be gravitated, you'll see yourself being gravitated to one or the other, but now macro wise, I think the one thing about academia today is that it's a lot more complicated than when I started in 1999.

And I think, for someone to raise \$5 million in a pre-[00:26:00]seed round is relatively straightforward, especially someone who's well known at the caliber of a strong academic professor. And to get a \$5 million grant, has a lot more work and a lot more strings, a lot more complexity. And so, if you're in the space that we're talking about, let's say machine learning AI for drug design and health care, the opportunities from a macro perspective are really quite juicy right now.

Very exciting on the startup side. Now obviously you have to build a company and so if you don't want to build a company, that's not the place for you. But if this is something that now can be taken from the world of ideas to the world of implementation, something where I think this is useful to like actually, I can really positively impact patients' lives. That I think is the reason for the transition and finally one last point is that I don't know if this is the best way to characterize it, but I think it's not uncommon for people to start in academia and then move over.

And you learn a ton being a junior academic. The one challenge is like if you're there too long sometimes you get comfortable with one type of approach to another. And I [00:27:00] think if you're there for long enough, you may just be content enough and stay. So, there's a balance there too.

I see the physicists and the kind of the chemists just come out there. And a little Michaelis-Menten separation of timescales as a way to think about this. I was trying to guess what you would say as the sort of the first axis, first dimension to think about this. And that wasn't what I was expecting, but it makes total sense.

And your other point about, exploration, kind of ambling, wandering through idea space without something very, tangible, immediate in mind, but really

trying to have the room both in time and in, what you even ask, right? Like what you even sort of evolved to ask over the next few years.

I think academia is still very wonderful that way, even with the grants and with everything else and all the pressures that we face as junior faculty. That is very, it's honestly very encouraging as a junior faculty member to hear that and to hear a vote of confidence for certain ideas as being uniquely, in certain pads being uniquely academic.

I was wondering as a follow-on, [00:28:00] maybe I can get you to talk more about medicine versus biotech. And so, medicine is set largely by clinical practice guidelines, clinical societies, uh, major clinical bodies that I think still largely exert influence via kind of academic channels or conventional academic channels.

And I think, this is probably, there are analogs for sure in biotechnology broadly, but I think biotech seems to me at first brush and at, you know, probably oversimplifying here, moving more as driven by industry and driven by currents that are happening in, you know, outside of academia.

And so, if you are someone who's interested in influencing medicine, via what is considered best clinical practice. Do you see that as something that is still uniquely academic, or is that even something that is changing and that you should contemplate industry a little bit more? I'm asking clearly for myself right now, but I think there are many people who are also in my bucket.

No, it is a good question and an important one. And I think you make a useful distinction [00:29:00] between life sciences and care delivery, because there's not a life science biotech equivalent of academic medical centers. There is not like academic biotech centers where— That's the much more eloquent, that's a much more eloquent way of what I was just trying to say.

We're very much in agreement there. And the role of academic medical centers is really critical, right? And, and that's where a lot of innovation still come from and so on. And that, um, uh, that may change, but that's not, that's very much the way today. I think a lot would have to change for that to be there for the way providers work.

So, I think from that perspective, there is a huge opportunity on the health care side for driving innovation in AMCs. Uh, the real question is like, to what degree will the reality of running a medical system also complicate things?

Because it's, it's a rough time to be running a hospital right now. And, uh, I was at Stanford during the UCSF–Stanford merger and unmerger.

I was chair of biophysics then, which is in the med school. So, I got to have [00:30:00] the delight of watching that, you know, from the chair's perspective. You know, that's a huge business proposition to have to deal with. And that's different than like what a lot of academics want to be dealing with. And so, the future of AMCs from a business perspective is also, I think, going to be important part of this and how can we keep the best and keep it sustainable is the question I always have.

Awesome. Thanks Vijay. So, I think we're going to run you through the lightning round next, if you're up for it. Sounds good. Awesome.

The first lightning round question requires a little bit of setup for our listeners. Um, but you sort of mentioned in passing that you're one of the first employees at Naughty Dog. I think one of the sort of like blue chip game developers now they develop things like, uh, The Last of Us and Uncharted. And I think what Naughty Dog does uniquely well that other game designers try to emulate is storytelling.

So, they have these amazingly cinematic games that have these well fleshed out characters. And so, I guess my first lightning round question is, [00:31:00] is there something in the DNA there that you learned about storytelling that has served you well as both a professor and venture capitalist? Yeah, I think so. I mean, in the early days, uh, it was me, Andy and Jason, and I think a lot of what you're describing came after, even after somewhat, uh, significantly after, but I think the storytelling that is perfected later is true of all video games and video games as a media, a medium for, uh, and so I think that was very much on my mind.

I think though, part of it too, is that, uh, I think part of being a physicist is, uh, physicist culture and being trained as a physicist is that there is storytelling there too, you know, uh, and even one thing that doesn't happen as much in biology is like even family trees and academic family trees and the stories associated with your academic family tree.

So, my advisor's, advisor, advisor was Lev Landau, the storied Russian theoretical physicist and so on. And so, like hearing about Landau's culture and all these things, it's, it's a different type of storytelling and narrative to live up to. So, I think there's [00:32:00] many different traditions there. Um, in, in terms of venture capital, the fun thing about venture capital storytelling is that

it's both a bit about predicting the future and a bit about making the future. And that if we can lay out a future that's plausible, but who knows what the future is really going to bring, right? So, a future that's plausible and we can, rally brilliant people to come help us and help our founders, and we can put billions of dollars behind it, that future could become a reality.

And that part's particularly intriguing, especially in health care where the stakes are so high, and the potential is so great for change. Got it. Thanks. Vijay, if you weren't in science or engineering, or science or engineering or investing, what job would you be doing? Yes, I've thought about this and like, um, one of the delights about my family is that I have a lot of cousins.

Uh, and you know, so obviously my cousins have the same, uh, grandparents. Uh, and so like one of my cousins is a chef. I could totally imagine being a chef. That would be like fantastic. And I [00:33:00] like cooking. Another one of my cousins is a psychologist and is quite famous in that. And so, I could imagine doing that.

And so, I look at my cousins and imagine all the different lives that I could have had, and these are sort of things that I enjoy. But probably the deep, deep fantasy, which I've not made a ton of progress on, but maybe made a little bit of progress on would be something like some sort of musician, like a jazz musician or something like that.

And so, I have music is a key part of my life. It's too late to be a jazz musician. But, uh, actually briefly, when I was at Berkeley, Vijay Iyer, the famous jazz pianist and piano's, uh, my instrument too, had the email Vijay at Berkeley, uh, physical Berkeley to do right before I did. So, I got all this, uh, Vijay Iyer fan mail, which was, uh, obviously not intended for me, but was, uh, was, uh, inspirational, uh, and it was fun to get.

Yeah. Our sound engineer, Mike, who is a professional jazz musician, uh, is nodding along in approval, I believe. Excellent. Excellent. Absolutely. Yeah. I'm a saxophonist, so, you know. Well, and so my daughter, my middle daughter is a [00:34:00] saxophonist, and so we try to do stuff. How's she doing? Like, she's doing, she's doing well.

I think, uh, she has a lot of things going. We have, um, Doug Ellington has a band here. He's the grandnephew of Duke Ellington. Wow. And he, we often have him at our house and I'll, we'll, we'll step in and, and they are very kind to play around us and make us sound way better than we do. That sounds great.

Yeah, yeah, absolutely. All right, so next question. Maybe I have some sense, but maybe I'm also way off base. So obviously, again, your colleague Marc Andreessen is a prolific generator of opinions on certain things. And so, I'm wondering if there's something that, what's the thing that you most disagree with Marc about?

That's a good question. What do I most disagree with Marc about? Uh, part of the draw to the firm was actually, we actually agree about most things. And, uh, he and I are fairly similar in terms of all sort of pre-a16z experience, in that I was never a CEO of a company, but I've been a chair of companies, and so has he, and so, and I like that [00:35:00] role, and he likes that role.

That's a great question. There must be something, but maybe diet. Like there's, there's a bunch of things on the diet. He's given up, he's given up booze, I think, if I remember correctly. Yeah. So, so, okay. So, so if you asked me this question six months ago, it would totally be booze. But recently, and he was actually somewhat of a catalyst of this because I was at his house and he showed me this hop water stuff, which was surprisingly not as bad as I was expecting.

So, my wife and I have greatly diminished alcohol, but I think the difference is that he's gone zero. And I've gone like a little bit and so the fun thing there and so maybe that's where you disagree. And I know he likes scotch and so on but now like the scotch is so much better when you have it much less frequently. And I was making that pitch to him.

I think I don't think he'd want to hear about any more of that. Yeah, so, so, he's dry, but you're still a little damp. It sounds. A little damp. A little bit, a little bit. Yeah. Got it. Amazing. Vijay. Will AI in medicine be driven more by computer scientists or by clinicians? Yeah, so this is a great question [00:36:00] and you're gonna feel like this is a cop-out answer, but this is really the right answer.

The right answer is that there are clinicians that are gonna be AI specialists, and that's what I was gonna have to be, and I think this probably for you two, this is not a shock, right? That it can't be somebody with just one mindset or the other. And actually the other thing, it would be hard to be.

It's not impossible. It's hard to have co-founders where one's the AI guy and one's the doctor. Even that's hard because they don't have telepathy. It's really different when it's all in one mind. And so, one of the unique things about

doctors today is that, freshly minted M.D.s have been around computers their whole lives, which is different than people like 20 years ago, even 10 years ago.

And so, I think we're going to have that. And there's plenty of exemplars of this. The Med-PaLM team is a great example. There's plenty of exemplars of people who are brilliant in both domains. And I think that really has to be the future. I totally agree. And this has been a big theme of the conversations we've had on the podcast too, where we're asking people about their background, physicians who are having a lot of impact and it really has been, I think a common story [00:37:00] that the latency is too high when the skills exist in two separate minds. And you just sort of get together and you don't really have anything to talk about as opposed to rapidly eliminating all these ideas and then coming up with the right path and it's kind of eliminating those thousand ideas

along the way to quickly get to where you need to go. So that's, that's great. So. Well, and so maybe actually the question I'd ask you guys, if I may, how's that going to work? Is that going to be someone with an M.D. learning AI or is that a computer scientist learning medicine? Andy, do you want to go first, or you want me to go?

Yeah. Yeah, I don't I, I think it's not either or so I think like I came into medicine through the side door. Yeah. But I think I, I think you have to have a deep appreciation of both. I think that it's very easy to be a computer science person and tell me like, what AUC should my model be calculating?

And that's going to be superficial and only get you so far. So, I think that like, as long as you're willing to be a nerd about both, um, you're going to have a lot of success, but you can't have a superficial [00:38:00] interest in either side of the equation. Yeah, I think that's right. And you can imagine someone who's like a CS pre-med undergrad.

Yeah. Yeah. It's easier said than done, but not impossible. And then they go to med school and they're like the master of two worlds. Or some M.D. Ph.D.'s who just, take time to study both subjects. I have to, this is my obligatory, Andy's heard this like seven times. So, he's gonna, he can simulate me fully at this point.

But I have to give my obligatory nod to this Ph.D. program I went through, here in Boston. So, this is the HST program, the Harvard, MIT Health Sciences Technology Program. And basically the Ph.D.'s take. The first, you know, a good chunk of the first two years of medical school, and then spend a couple of

summers in the clinic where you're taking histories and physicals, rounding with the teams, presenting, really understanding how doctors think.

And then the M.D.'s are very technical and they do an extra year or two of research. One of my friends of the program actually is at a16z now, Vineeta. I'm sure you, you know, you know where Vineeta Agrawala. And so, uh, it's just a, it's a, it's a great program. And it's like a, the [00:39:00] core thesis there is that you can't really have one skill set or the other, but you really have to invest in, in kind of education in both.

And so, I think we need more of those training programs of real domain expertise alongside the technical skills. I agree. And I think there's real potential for that. I think at Stanford, the MSTP group was a little bit like that, but HST is really special. Awesome. All right. Thanks. So, this is a question that now has almost become cliché because we ask it to everyone, but it's elicited such interesting responses that we're going to keep it going.

So, if you could have dinner with one person dead or alive, who would it be? I should have like a stock answer for this, but, uh, I believe we're not, I've probably done like a hundred podcasts and no one's asked me this yet. Nice. Uh, so, uh, yeah, no, no, but like, so I feel like I should do better, uh, with having an answer.

It's such a hard question, but like, uh, who, the people who come to mind, are the real old school people like Claude Shannon, John von Neumann, uh, Alan Turing, that generation [00:40:00] where I wish, just for them like, I could tell them like, oh, actually all this cool stuff came from like the seeds that you planted, the ideas that you had. And it would be so much fun to just sort of hear their view of what building the foundations was like, and just to get their take on where we are now.

There're few people that were such the polymaths, like von Neumann's maybe the canonical example. And I think I'd be probably if I had to pick one, maybe he would be the one that would be super interesting. That was a fun time. I mean, I think today is an amazing time too, for sure. But like, um, someone like that, I think would be, would be just, uh, fun to see the direction, the discussion is going in both directions.

Yeah. I, it's your choice. I, I always worry about meeting someone like that in real life though, because whenever you humanize them, the legend gets instantiated as an actual person and you no longer have the legend in your head.

So yeah, I've seen this where like, what was it was a comedian or a sports figure,

oh, it was when Willie Mays dies. [00:41:00] What, died. I was watching PTI, and I don't know if you've seen Pardon the Interruption. Pardon the Interruption. Yeah, yeah, yeah. And they're talking about, Kornheiser was talking about meeting Willie Mays and he's like, just for your sake, never meet your legends. It's not going to go the way you think.

So, it's kind of the way it is, but like in my mind, it could be amazing. And like probably with AI, we can make a fake version that would keep me quite happy. All right. This is our last lightning round question. Given recent progress by Anthropic and Google, do you think OpenAI will have the most capable foundation model in one year?

Yeah, so it's a good, good question. And a year timescale is well chosen. Because if you asked me in one day, I think I would say no. If you asked me in one decade, I would say it would be OpenAI for a day. A decade is much more open. And a year. Probably in a year, I think there's going to be, there's all this talk of clout doing well.

And so, it's possible things have switched. I think the one big challenge is going to be is like, what is the right metaphor for AI? And I [00:42:00] think my favorite metaphor for AI is probably the microprocessor in terms of business. And if you remember the early days of microprocessors, there were a few, like the 4004,

the 808, and then 8080. And then there's a little bit of explosion with others and then, uh, basically Intel again, uh, and so on. And so, I think short term, there's only a few people that have the technology. Midterm, a lot of people have the technology and then long, long term, I wouldn't be surprised if maybe there's reasons that it would converge back again.

I think we're going to see a bit of an explosion and we already see that to some degree with all the open-source stuff with LLAMA and Mistral and so on. I think the other thing is that we're probably going to get less hung up on one LLM to rule it all. And that there's going to be lots of different things, much like you don't just work with one person or talk to one person or so on.

So, I think we're going to have lots of different uses, and that will be part of the explosion as well. This is back to your mixture of experts and agents all operating together. Yeah, exactly. [00:43:00] So well done Vijay, you've passed

the lightning round with flying colors. Thank you. It was so much fun. Thank you guys.

So, we want to just wrap up with some big picture questions. And I think the last lightning round question was a natural segue. The purpose of this line of questions is to just get your sense of how the next five-to-seven years are going to go. But we're going to anchor it on this essay by Leopold Aschenbrenner called "Situational Awareness."

I don't know if you've had a chance to read this, but— No. So, he's been making lots of rounds on the Twitterverse for writing this 183-page essay on what the next five-to-seven years are going to be like. And what's helpful is he makes some specific claims. So, I'd like to throw some of these specific claims at you and sort of get your reaction to them.

So, first he says that we're going to have AGI by 2029. And by AGI, he essentially means superhuman intelligence. I think we already have AGI now. I think that by any reasonable definition, LLMs are a general kind of intelligence. They're not superhuman in lots of different categories. So, he thinks 2029 essentially is the singularity, the event [00:44:00] horizon.

We're going to cross over it. They're not ever going to be able to come back. Getting there though, he says that the rate limiting factor is going to be energy. The GPUs that we need to train these models are going to be so power hungry that we're going to need a gigawatt power center. And that essentially, we're going to have a trillion-dollar data center.

That's a combination of these huge compute clusters, along with a nuclear power plant that's hooked up to it, to power it. So, when I guess, maybe I'll just throw all that at you and sort of get your reaction to it, and then we can like sort of pull on some individual threads. Yeah. Yeah. So, let's take the power thing first, because I think that's, uh, so and I think you can think like what fraction of the world's energy should go to AI.

And I think the answer is a relatively large fraction concerning the implications of this. And so back of the envelope calculations, 10%, 20%, 30%, that's a lot of energy that to add on top of things. And it actually has obviously real geopolitical implications because you just can't make energy just anywhere.

And so, places with large scale [00:45:00] energy production, whether that be oil and natural gas, or solar or nuclear. Become appealing with obviously nuclear becoming appealing because of low CO2 generation and so on. But

also, you want to do this someplace where there's cooling. And so, places with oil maybe not be the coolest place to have this.

So, I think all of that becomes logistical problems, infrastructure problems, but I'm excited about spending more energy for AI. I think that could be one of the best uses of energy that we could use as a society or as a, as a world. Terms of AGI, I think this is where it gets very complicated because well, first off, like even just intelligence, like, I don't know how deep you want to go, but you could talk about IQ, you could talk about G factors, but even G factors is very human centric.

Like AI has intelligence that's just really different. And that even like LLM right now, I can ask it physics questions, which I know the answer to. I can ask it music questions. I can have it do things like, I think it can do things that probably no individual could do, but like a group of people could do.

And to some extent it's trained on a group of people, so it's [00:46:00] reproducing a group of people. And so that's a kind of super intelligence. And then, we've already seen with things like radiology, another type of super intelligence is where AI is better than any individual radiologist. But it's, comparable to the, the group of experts and that's considered to be super intelligence.

And then I think what we really are looking for is something where AI comes up with some new physics or some new medical breakthrough or something that is just as different as Einstein was, or as, something that is that level of creativity. And even Einstein is such a hackneyed example, but it's just the obvious one because like his, uh, general relativity is something, or even special relativity is something that was such a major leap of thinking, albeit built on work of others, as well.

So, that part, I don't know if we're going to see that or not. And I think what could easily happen is that LLMs are very much a next token predictor and next token prediction does really well at these current games, but is that the thing to really take us to the next level? I do think that learning [00:47:00] is building latent spaces and this could be jazz.

The reason why you do scales in jazz and all different types of scales is that's a latent space. Um, I do martial arts, all the martial arts stuff is learning latent spaces and a language for, for combat and so on. Uh, learning latent spaces and then doing direct products of these latent spaces and all these things.

That's very natural thing to do. LLMs can do that to some extent, but I wouldn't be surprised if there is an algorithmic shortcoming. And if we were to study the history of AI, it's a series of S-curves where eventually it's like, it's going to solve everything, like single layer neural nets are going to solve everything.

Uh, except for XOR, except for this, except for that. And then, we get to a plateau. It would not surprise me if we get to a plateau before HEI. And this 2029 number would be a very natural thing if the S-curve were an exponential, but what if it were a sigmoid? And if it's a sigmoid or logistic, it's going to bend over.

Maybe we just don't get there. And that extrapolation doesn't work. I would not be surprised if [00:48:00] there's something really missing. And that's the cool thing, because I think we'll get it, but maybe just not this soon. So, I think one common rebuttal that I hear about AGI by 2029, I think you were kind of alluding to there is that like, they are just trained on Internet data and Internet data is only get you so far.

And we've already extracted all the signal from Internet data. Do you have a sense of how true that is? So, like the models that we're using, have they actually extracted all of the information that humanity has produced so far? Or are we still at some like small fraction of that? Cause I think that would suggest, how much room we're going to have

for improvement on the current paradigm. The one thing though, is that like the way computers learn is not the way we learn new things. It's the way we learn existing things. So, like, um, I don't know if you ever had friends that like grew up in foreign countries that moved here. I've had a few that they learned English from watching TV.

It was shocking for me to imagine that's even possible, but it's a common thing. And so that's kind of very LLM-ish, right? You're just watching and absorbing and [00:49:00] maybe there's a little RLHF in there when your parents or your teacher says that's wrong and that's wrong, but it actually reminds me a lot of learning English as a foreign language.

But that's different than learning, like in a Ph.D., we're into a new discipline where nobody knows the answer. And it's a whole different process. And I think a lot of things that we're excited about are things that nobody knows the answer to. And that's a different paradigm. That's maybe even learning is not the right word that's exploring, or creating, or discovering.

And that's not what LLMs have really been geared up for. And even, building tasks and Q-star and so on, that's not going to, I think, teach it how to do that. And there may be something to do on top. In a sense, LLMs are like a really good undergrad that's read all the books but is not ready to be a grad student yet, maybe, I don't know.

Yeah, I agree with that. I think that there's something to being embodied and interacting with the world that's missing with the current paradigm. So, one thing that I always think about is LLMs are great at selecting from a set of [00:50:00] hypotheses that are compatible with what we currently know, but often the only way to know which one of these hypotheses is correct is either to assume your way there.

So, I assume that I know how the world works and therefore I can rule some of these out. Or to do an experiment or interact with the world or push the cup off the table and falsify or confirm, your hypothesis. So that does seem like, I don't have a sense of how big the fraction of that missing piece is, but it definitely seems to be a missing component.

Yeah. So, towards that end, it could be that the human feedback that LLMs get through massive chatbots. Could be, but like, that's different than still like the type of advanced learning that we're talking about. So, um, I don't know. I mean, I'm, I'm very curious to see, and it could be, you can imagine creating different types of bots that are the professors and they create students and then you cycle over.

I don't know, reproducing academia and AI may not be the ultimate end goal, but that is an interesting paradigm to consider. I guess Raj doesn't have any questions. Just one sort of like forward-looking thing that I'd like to get your take on. My experience has [00:51:00] been that when you're talking to people about AI, especially people who are hearing about it for the first time, there's either elation or fear.

And that like it seems to be pretty polarizing along those two axes that, oh my God, this is amazing. Or, oh my God, this is the most terrifying thing I've ever heard of in my entire life. How could you possibly be working on this? Um, at least externally, a16z seems to be formally, very firmly in the optimist camp.

So maybe you could give us the bull case for optimism and about why AI is going to give us the sort of future that we all want. Yeah. And I can also give you a, what my understanding the psychology of the fear case is. I think, this is something that has happened over and over again. Part of predicting the future

isn't, like, you're a genius. Like I can predict what time the sun's going to rise tomorrow. It has a daily cycle. Some things have monthly cycles. Some things have yearly cycles. Technology and industrial revolutions people have studied, historians have studied, and they have like 25-year cycles. And we're in the middle of industrial revolution.

This is not a first. This is our fifth or sixth, depending on how you count them. And in these [00:52:00] cycles, there are all these people that are scared of the new technology and there's a lot of reasons to be scared. It changes things like when, before cars, if you were a master horse breeder, like you work really hard to be good at this really important thing.

And then that goes away. That is scary. And there's a lot of reasons, the fear of that. And if cars, the first-time cars are on the market, like cars kill people, they run them over. They're this new thing that you're not used to thinking about, and not in a way that horses run people over, it's just different. And sure, cars are ambulances, too.

But if you're just looking at the negative side, being run over a car is more scary than having cars to help you and all the things that you can do. And so, every fear that we could attribute to AI, we could also attribute to electricity. You know, you put this thing in everyone's house and that could kill you if you touch it, you know, or the positives of electricity or the negatives of cars, the positives of cars.

And so if you harp on the negatives, it can, anything can look really bad. Cars and electricity can look really bad. If you consider the net positives, the net positives are dramatic. Nobody wants to [00:53:00] go back to a world without cars or without electricity. If you go that direction, you're in a hunter-gatherer society where you know, the bully beats up everybody.

And that's the person who runs the little village. I mean, technology is the great equalizer that makes, uh, democratizes everything that the technology we have today gives us. Things that like a royal king a hundred years ago would never have. We all have the same Spotify and the same iPhone. I think on the medical side, I think the exciting thing is we will all have the same specialist.

We will have the best doctor, just like we have the best Spotify or the best iPhone we can all have that. And the democratization is something that just doesn't exist in something like health care today with such disparities of quality access and care and cost. I think that's the real hope and I think that's what gets me excited about AI in health care.

But I think this is AI broadly and it's technology broadly. And as a firm, we're very excited about technology because we've seen all the positive things that it's done. Thanks. I think that's a great note to end on. [00:54:00] Oh, fantastic. This was so much fun. Thank you. Thanks for coming on, Vijay. Thanks so much, Vijay, this was great.