AI-GR Podcast 17 04.12.24 D Koller

[00:00:00] And we've seen that there's more and more genetic insights that can come out. Now, when you think about the broader spectrum of data that can be collected about even just a single human organism, you have single cell RNA seq from different cell types in different states at different ages under different environmental conditions.

[00:00:22] That's just RNA seq. You can also think about protein levels, and we haven't gotten to protein levels at the single cell. We're starting to scrape the boundaries on that, but it's certainly not at scale. And then going down into the level of individual proteins, going up to the level of, now everyone's excited as they should be around spatial biology and the interplay between different cells.

[00:00:43] The number of data modalities, the number of ways that we have to measure biology, and the number of distinct biological contexts with, that we as a human live in and that exists within even the body of a single human is ginormous. It is way larger than [00:01:00] what I think the complexity that we've trained the large language models that we're currently leveraging in these in the more traditional

[00:01:06] LLM sense, what we lack are data collection approaches that achieve that scale. And I'm just really excited to be living in this time because the number of ways that we have to measure biology quantitatively and at scale is increasing, maybe not quite as fast as the capabilities of AI work a few years back on that curve, but you can see that exponential curve.

[00:01:31] And I think the synergy between those two is going to just unlock just an incredible tidal wave of insights as we start to bring those two tidal waves together.

[00:01:47] Welcome to another episode of *NEJM AI Grand Rounds*. This is Raj Manrai, and I'm here with my co-host, Andy Beam. We're really excited today to bring you our conversation with Daphne Koller. Daphne is the CEO of [00:02:00] insitro, where she's working on artificial intelligence for drug discovery. Daphne has really done so many things, Andy.

[00:02:06] She's had this illustrious career as an academic, as a professor at Stanford. She started Coursera. This was an amazing and really, really wide-ranging conversation. Let me admit something to you, Raj. I was nervous going into this interview. One, I think, you know, as you know, I was on paternity leave, so maybe not at my most mentally sharp, but Daphne is such a force of nature that it's very intimidating to talk to her.

[00:02:30] She is so accomplished and so bright that it's hard not to be a little intimidated going into a conversation with her. Having said that, I think that we covered a lot of ground and I think that the listeners are really going to get a sense of all of the different areas of computer science and biomedicine that she's had a significant impact on.

[00:02:47] It was, nerves notwithstanding, a real treat to get to have her on the podcast. And again, I learned a lot from this conversation, both about what she has done, but also how she thinks about the world. So, for a lot of reasons, it was a real treat to [00:03:00] get to talk to her.

[00:03:04] The *NEJM AI Grand Rounds* podcast is brought to you by Microsoft, VisAI, Lyric, and Elevance Health. We thank them for their support.

[00:03:18] And with that, we bring you our conversation with Daphne Koller. All right, Daphne, thank you so much for joining us on *AI Grand Rounds*. We're super excited to have you here today. Thank you, Andy. It's a pleasure to be here. Daphne, welcome to *AI Grand Rounds* and thank you for joining us. So this is a question that we always like to get started with:

[00:03:36] Could you tell us about the training procedure for your own neural network? How did you get interested in AI? What data and experiences led you to where you are today? So, I will say that when I got into AI, what I got into wasn't actually AI. I got interested in the question of how we can get computers to make better decisions, [00:04:00] initially using decision theoretic principles.

[00:04:02] And AI didn't encompass that at the time. I got into the field, I'm old at this point, um, at a time when AI was all about logical reasoning and, you know, uh, symbolic computations. And we, there were people who went around saying that, oh, you're using numbers. People don't use numbers. What you're doing is not AI.

[00:04:22] And so what really happened is that what I did, which was initially decision making under uncertainty and then learning models for decision making under uncertainty. So, learning models that enable the computer to make sound decisions that got swallowed by AI and eventually basically took over the field over time.

[00:04:44] Could we get you to go back a little bit further? So, you were doing this presumably in a computer science department. A lot of that sounds like classical statistics to me. Could you take us back before grad school? Like, what sparked your interest in these areas and how'd you get there in the first place?

[00:04:58] So I have a somewhat unusual [00:05:00] academic history because I started going to college at a relatively young age. So, I got into computer science, honestly, as it was becoming a field in the 80s as a very young high school student. So, when I was 12 or 13, and I completed my undergrad degree at 17, and computer science was at the time, really almost a branch of mathematics.

[00:05:24] It wasn't really in most universities, a field in its own right. It was really very much a mathematical study. And I loved actually both the mathematical component of the field, I was a double major in math, but also the fact that you could actually take these very abstract conceptual frameworks and then use them to get the computer to do stuff.

[00:05:45] And, and the computer was able, you know, at the time, these are very rudimentary capabilities, but you could build a game that got the computer to play Pong. And that was like, so cool that you could actually tell something what to do, and it did it. No one else that I tried to tell what to do did what I said.

[00:05:59] So it was [00:06:00] nice to have the computer do what I said. And so, it was really empowering and fun to do that. Um, and then the question was, well, okay, if we can get the computer to do stuff, how do we get it to do the right stuff? And initially, my interests were actually in multi agent systems. So, if you look back, to the work that I did even, um, in my master's degree and then subsequently in my Ph.D.,

[00:06:23] a lot of that was about multi agent systems and game theoretic models for multi agent systems. And then I realized that before you get the, uh, a community of agents to do something intelligent, you first have to get individual agents to do something intelligent. We were very far away from that, which led me to the study of these decision theoretic systems and decision making under uncertainty. And then realizing that the biggest obstacle to that was that the computer just didn't have a very good model of the world, and that people just weren't particularly successful at building usable models that really captured the [00:07:00] complexity of the real world around us, and then recognizing that the only way to get to that was via machine learning.

[00:07:06] So in some sense, I got to machine learning via the back door at a time that it really wasn't part of AI. I was the first machine learning hire into Stanford's computer science department and frankly the strongest proponents for hiring me weren't the traditional AI people who were the department at the time, it was actually others in the department who saw the value of this more modern style of getting computers to act intelligently.

[00:07:32] They were the biggest proponents for hiring me. Yeah. I remember I took an AI class before machine learning really took over and now we call it good old fashioned AI, you know, Russell and Norvig, uh, that textbook. Uh, and there's like a tiny sliver of machine learning in there. So, it is interesting that you were kind of at the vanguard.

[00:07:49] of the probabilistic or data driven or machine learning approach to AI. So, what were those early days like when you were kind of this insurgent in a very historically renowned AI department at [00:08:00] Stanford? So, if you look back, I mean, 1985 was one of the AI hype cycles, one of the earlier ones. And there was a big AAAI or IJCAI conference.

[00:08:12] These were the big AI conferences at the time. And there was a panel on AI. What is the future of AI? And most of the panel, with the exception of Judea Pearl, who I consider to be one of the foremost leaders in this new vanguard, all said that you cannot use numbers in AI. People do not use numbers.

[00:08:33] Numbers are anathema to AI. Probabilities for sure, and they had a whole bunch of arguments about why probability theory would never be the basis for intelligent reasoning. And Judea, who's a close friend and mentor, uh, basically was the one holdout for this is the future. This is what we need to do.

[00:08:51] And I would actually say that Russell and Norvig were among the ones that actually were helpful in migrating the field more towards this decision making [00:09:00] under uncertainty and machine learning. The earlier textbooks didn't have any mention of either probability theory or machine learning.

[00:09:08] So in some sense, they were the transition point. And I did my postdoc with Stuart Russell, and I know that he's a very big proponent of, he was even then a big proponent of that, of that transition, but it took the field quite a long time to adopt it. And there was this schism where those of us who did the more probabilistic stuff, mostly published in a completely different set of venues than the traditional AI conferences for quite a number of years.

[00:09:34] I don't want to go on too much of a tangent here, but your comment about Pearl is very interesting because back in those days he was saying, we need probabilities, we need data and now I think he stayed mostly the same, but the field has kind of polarized around him. His argument now is like, we're too reliant on data.

[00:09:50] We need actual causal models of the world and, you know, the probabilities are going to lead you astray. I don't know that I've heard him say that the [00:10:00] probabilities are going to lead you astray. I think what happened was that the field really did swing as a pendulum. It swung from we entirely logic based, you don't need probability, you don't need data, you don't need numbers, migrated temporarily, transiently through something that had elements of both symbolic and machine learning, and then swung with the advent of deep learning much more towards

[00:10:26] we don't need any kind of world model. It's all going to be just the numbers. We're just going to learn everything from data. And I think now if you listen to the people who were in some ways, the leaders of that revolution, people like Yoshua Bengio or Yann LeCun, they are, are now coming back to we need to have some level of causality, some level of tie in to sort of more symbolic concepts because that is going to be important for common sense reasoning and [00:11:00] it's going to be really important for taking action in the real world and understanding what that action is, what the consequences of that action is going to be.

[00:11:09] You cannot rely on plain old pattern recognition, which is where, um, you know, deep learning has really had its biggest impact. And so, I think there's now almost like starting to gravitate back towards the middle a little bit. So, Daphne, if I can take us in a little bit of a different direction. So I think you framed

[00:11:27] the very exciting early days, let's say, of AI at the CS department or machine learning at the CS department at Stanford. And you, I think, are of this very rare group of people who've had real contributions both in computer science, so general computer science, probabilistic graphical models, support vector machines, your work there, but also, of course, in applications in biology and medicine.

[00:11:51] I don't think I actually know the story of this, uh, but could you tell us about how you got interested in biology or medicine, um, and where that, [00:12:00] let's say not transition, but where the work that your group started doing at Stanford in applications of machine learning, computer science, where it started and when it took off.

[00:12:10] So I'm going to actually go back a little bit further, if that's okay. And talk a bit about my own personal journey as I think about what to work on. When I was a Ph.D. student at Stanford, my work was highly conceptual. Very, um, a lot of theorems, a lot of abstract concepts, some very beautiful frameworks.

[00:12:30] And then when I went to do my postdoc at Berkeley, I had what turned out to be, I think, a very pivotal conversation with my postdoc advisor, Stuart Russell. He took me out to lunch one of the first weeks that I was there and said, so you did this beautiful Ph.D. thesis. You got a Ph.D. thesis award for it. If I gave you a group of really talented undergraduate computer science majors to work with you and you could code together something from your thesis, what would that be?

[00:12:59] And [00:13:00] I literally sat there, I think, with my jaw hanging open because no one had ever asked me that question. And if I had to answer it, honestly, the answer is nothing. There was from the thesis that I would have wanted to implement and turn into a useful product, a useful artifact. And that bothered me.

[00:13:20] And it kind of pushed me on a path of an increasing commitment to building things that actually make a difference. And that moved me from the more conceptual work that I'd been doing to probabilistic graphical models, which I thought were useful from probabilistic graphical models, all that started doing that even prior to my postdoc from probabilistic graphical models to machine learning, to applied machine learning, to machine learning in the service of actual disciplines, which at that time were broader than just biology.

[00:13:52] I worked in machine learning applied to robotics and to computer vision. And at some point, [00:14:00] I got interested in biology, not because I had any particular affinity to biology myself, because the truth of the matter is, I went to a high school that was highly tracked, and I was tracked math physics, um, and, you know, the biochem people were a different breed, and we didn't talk to them, and they didn't talk to us.

[00:14:17] And so, I really didn't know anything about biology, but this was a time when the datasets that were available to machine learning researchers at the time were actually kind of boring. So, there's only so far that one can get excited about spam filtering or, even worse, classifying articles into 20 news groups, which were the datasets that we all had to work on.

[00:14:39] And I wanted to do things that were more technically interesting. And this was the time when the first large for the day today, they're tiny, of course, um, datasets were coming out in biology and medicine. And these are things like, for example, the first microarray data where you could actually start to think about gene-gene [00:15:00] interactions and relationships between genes and phenotypic consequences.

[00:15:04] So the first, um, project that I did was actually with a wonderful colleague who was a tuberculosis expert. And we worked on the project having to do with machine learning for tuberculosis epidemiology on a dataset that when you think about the recent COVID datasets is minuscule, but it was one of the very first, um, network, if you will, of transmission.

[00:15:27] And then moved from there to working on the, again, gene expression and the ability to infer regulatory networks from data. And then from there to some of the genotype phenotype correlations and, uh, some of the earlier relationships between genetics and gene expression, genetics and phenotypes. And so, the nice thing about biology was that every few months, there would be another really cool dataset, oftentimes in a different modality that I was unfamiliar to me.

[00:15:58] So it kept [00:16:00] creating new challenges and opportunities for novel machine learning to be developed. And so initially, my personal interest was mostly this is a great place to find good challenge problems for machine learning. But then over time I became interested in the field in its own right and ended up having this really weird, bifurcated existence where half my lab continued to do core machine learning work published in computer science venues.

[00:16:26] That was going to be my next question. That was going to be my next question because I feel like we have, you know, we sort of, we split right in the sort of methodological focus and then the applications. So, yeah, so those, the main machine learning conferences and then half the lab is in the general scientific journals.

[00:16:43] Yeah. And they, you know, and there was some interplay between them in my lab, but when you went out in the sense that the methods people were sometimes inspired by the biological problems and certainly the people working in biology were very much informed by the methods development that was being done in [00:17:00] the group.

[00:17:00] But when you look at the outside world, I have, even today, computer science colleagues that ask me, so why did you get into, uh, biology so late after you did Coursera? And it's like, no, no, there was this whole thing that you didn't even know about. And then you had, on the other side, biology colleagues who I think didn't even realize that I was in a computer science department because since when do you have people publishing in nature and science and cell in a computer science department?

[00:17:27] And that bifurcation was kind of odd. Yeah, you're really straddling a lot of cultures, right? So, like, as you mentioned, it's what the main journals are that those communities read or what they look to even I think the publication culture is very different in computer science versus in medicine or in biology. And so, it's very typical for example and in statistics, right?

[00:17:50] It can be wildly different too and so if you're being evaluated or you're being uh, you know, working in close collaboration with [00:18:00] researchers in those communities, I think there can be a lot of challenges and opportunities to overcome as you're sort of navigating. But you know, Andy and I are both at a med school and at a school of public health, and we both work in AI and medicine, and there are very different venues and very different criteria that are applied by those communities and judging research output and, and what is a paper even, right?

[00:18:22] The sort of foundational question. Oh, no, completely. So, first of all, the computer science field, by and large, um, thrives on conference papers and relatively small units of work at a very rapid publication pace, whereas the field of scientific inquiry focuses on much longer form pieces that can sometimes take

[00:18:42] years to complete. And that's a very different sort of cadence for how the work gets done and how people get evaluated. And so that's one piece. And then I would say there's a very deep, I would say, sort of, um, mindset shift between [00:19:00] how one thinks about science and how one thinks about engineering. And I think about machine learning often very much in the engineering side of things.

[00:19:08] When you're an engineer, you're looking for patterns. You're looking for the model that will explain the maximum amount of the data that you're observing with reasonable amounts of accuracy. And when you've found that, that is the victory. That is the winning state, is a model that generalizes reasonably well for a pattern you've been able to discern.

[00:19:29] When you're a scientist, oftentimes what you're looking for is the outliers, the exceptions, because those outliers are oftentimes the beginning of a thread that will lead you to a completely different and novel scientific discovery. So you have one group that's looking for patterns, and the other group that's looking for exceptions, and that makes that interdisciplinary communication quite challenging sometimes, and it's definitely something that even today, as I'm building a cross functional company with individuals from [00:20:00] both of these groups, getting people to communicate is more than just about jargon and making sure that you're familiar with each other's terminology, but also about how you think about science.

[00:20:10] So I was going to say, I've never sort of thought about it that way, that engineers sort of care about the mean or the first moment of the distribution, whereas scientists might care more about the second moment or the variance or the tails. I know that we speak different languages, but I've actually never thought of it like we care about different parts of the distribution kind of fundamentally.

[00:20:26] That's a very interesting like perspective on that. I wouldn't call it just the first moment, but I would call it the sort of the, the pattern that explains enough of the data. It doesn't have to be the first moment, but the pattern that explains enough of the data so that you feel you have the ability to generalize to new data points, whereas the outliers, the exceptions, are where you kind of like, well, my model doesn't explain this point.

[00:20:51] Why? Assuming it's not an error. Why is this at this point is different? And what novel insights does that unlock from a scientific discovery perspective? And so, it's [00:21:00] a very different mindset. Yeah. So, I'm going to fast forward a couple of decades from where I think we were when you were starting your lab. And Andy is going to dive into your current work at insitro in just a couple moments.

[00:21:13] But before we do that, I wanted to ask you, your perspective and your thoughts about a very interesting paper that you recently wrote, which Andy and I were happy to co-author with you, which of course was published in *NEJM AI*. And so, this is a paper called "Why we support and encourage the use of large language models in *NEJM AI* submissions."

[00:21:34] I think you really led this. And so, can you give us your perspective on what we're trying to say with this editorial, with this article, and how you anticipate LLMs being used by scientists in improving analysis of data, communication of results, and their dissemination? Yeah, no, thank you for asking that question.

[00:21:55] I think I was struck by the fact that the advent of LLMs [00:22:00] caused so much distress among some of our scientific colleagues in terms of, wait, so now machines are going to take over our job and they're going to write scientific papers. And so, we should prevent that from happening. And we can discuss at length why trying to prevent scientific progress is a bad idea because it's never been successful before.

[00:22:22] So there is that. But I think maybe even more to the point is tools like that elevate all of us and they allow all of us to do better work. They allow us to do better research in terms of understanding what's out there by summarizing papers for us before we dig into the ones that we think are the most relevant.

[00:22:46] They can help us do better data analysis, make better figures. They allow us to write better prose, especially for those of us that might not have, say, English as a first language, or have some other kind of language disability. Doesn't mean they're worst scientists. [00:23:00] It just means that maybe writing isn't what comes most naturally to them.

[00:23:04] So you're elevating everybody. You're actually equalizing the playing field for people who come in from maybe less advantaged backgrounds. So, I think both from an egalitarian perspective and also from the perspective of our goal in *NEJM AI*, and I think in science in general, is to elevate the quality of the science that is done and the insights that that provides us.

[00:23:28] It is not to try and judge whether someone writes better than somebody else. For that, there is college and exams and so on and so forth. But when you get to the point where you're actually a practicing scientist, what we should care about is the quality of the science that you're able to produce. So, I think it's absolutely the right decision that we took.

[00:23:48] I think that frankly, in not a very long amount of time, I would say no more than a couple years and probably less, the idea of banning the use of these tools will be as [00:24:00] laughable as the idea that we shouldn't let people use computers to do data analysis or, or even calculators. It's just going to be a, or that we shouldn't let

[00:24:09] people use spell checkers, which of course is laughable today. But if you go back not that long ago, there were people who were advocating against the use of spell checkers and calculators and computers to do data analysis. And I think the banning of LLMs from scientific study is going to look equally laughable in a couple years.

[00:24:27] Yeah, I completely agree. And just as a disclaimer, I think we are all aware that now there's a startling number of papers in Google Scholar that if you search for, as a large language model, I can't do blah, blah, blah, will actually show up in published papers. So, we are certainly not advocating for people to turn off their brains.

[00:24:43] Or, or anything like that. This is responsible use of LLMs. And like you said, writing is a very specific skill that strangely science is highly selective for. And it's kind of an orthogonal skill, like the ability to think, the ability to reason, the ability to be rigorous, are somewhat orthogonal to your ability to communicate those [00:25:00] ideas in written prose.

[00:25:00] And I love the idea that this is a great leveling tool, uh, for that. And the ability to, um, summarize the ever-growing literature so that we are able to, um, potentially, better contextualize our work in terms of what's already been done. I think it's going to make for better science overall, but you're absolutely right, Andy.

[00:25:20] This does not mean that it absolves the scientists from the ultimate responsibility for what their paper says. Ultimately it is your responsibility as a scientist to make sure that you have conviction behind the correctness and the novelty of what you produced. And that is absolutely, and we stated that very clearly in the paper as well.

[00:25:38] Agreed. Daphne, what is your sort of favorite use of language models yourself in either, I'm going to guess just from some of your comments, maybe consuming or summarizing some of the scientific literature or some other aspect of analyzing data or preparing, editing, how have you found them useful in writing papers?

[00:25:57] I think that right now, [00:26:00] for me, the killer app really is the summarization of papers and the scientific literature because the amount that is out there is just overwhelming and growing so fast. The ability to sort of very quickly get a read on whether a paper is likely to be relevant to the question that I'm studying, which oftentimes you're not going to get from the abstract because the thing that you're looking for is somewhere on page five.

[00:26:26] That to me is, I think, a real killer app, personally. I will say that when I think more broadly about the impact of this technology, I think one of the biggest impacts is going to be the democratization of programming. Right now, programming is one of the more challenging disciplines for people to manage, to learn, and it's yet a huge empowerer of people in terms of making sense of the world. Getting computers to do things that they personally find interesting, even if it's organizing their, you [00:27:00] know, their, their photos or their recipes or, or something that is more important, um, from a scientific perspective, like analyzing data.

[00:27:08] Right now, there's a gap. Even, you know, when I think about scientifically very talented colleagues that never learned to program and they're really dependent on having a data scientist kind of glued to them at the hip, helping them do their data analysis, which really limits the number of hypotheses that they can interrogate.

[00:27:26] And so if we create something where you can program by natural language and ask questions of the world in, uh, without needing to learn how to program, I think that will be hugely democratizing. I do think that it will create an obvious gap in the next skill set up the hierarchy, which is structured thinking, which is something that, unfortunately, we as a community do not take the effort to explicitly teach to our students.

[00:27:55] We kind of figure that they'll learn it on their own or they're born with it [00:28:00] intrinsically. I don't think either of those is true. Um, and I think teaching structured thinking is going to be an imperative for educators going forward because that's going to be the thing that unblocks your ability to leverage LLMs for problem solving.

[00:28:15] I agree. I think the thing that I always take away from my degrees in stats and computer science was not my ability to write Python, but my ability to think algorithmically and probabilistically and reason under uncertainty. It's more of a way of thinking than it is a skill set. A hundred percent. I have to say, unfortunately, I haven't programmed in a large number of years at this point, but I, but those skills of really taking a very mushy, abstract problem, one that's ill formed and breaking it down into manageable pieces that together create a solution to the thing you were originally looking to solve.

[00:28:49] I think that is a skill set that's not going to go away anytime soon. Agreed. I'd love to hop forward now to insitro. Um, so, for those keeping score at home, you've [00:29:00] been, a child prodigy, a Stanford professor, Coursera co-founder, and now we're going to transition to founder and CEO of insitro. So, could you tell us a little bit about the founding story around insitro?

[00:29:12] Like what made you want to take this on? Cause being a CEO, again, is a very different skillset than being a researcher and being an academic. Uh, so yeah, could you, could you walk us through that? Um, so I think I'm going to go back a little bit earlier than your question and to the time that I departed Stanford to go to Coursera.

[00:29:29] And that really emerged from what had been an increasing sense of urgency to make a difference in the world and trying to think about what could I do that will really have much more of a direct impact than just writing papers and hoping someone reads them or training students and hopefully go on to do something meaningful.

[00:29:47] And so at that point, work that I'd been doing at Stanford, um, for technology assisted education basically led to the launch of the first three Stanford, so-called MOOCs, Massive Open Online Courses. And when I looked at that [00:30:00] impact where we had 100,000 learners in each of those courses in a matter of weeks, I had a choice of, well, I could just go back and write some more papers, or I could actually leave Stanford and do something to really bring that vision to life.

[00:30:12] And I decided to do the latter, left Stanford on what was very much supposed to be a two-year leave of absence rather than a permanent departure to really try and bring that vision out. And it was an absolutely terrifying experience because not only had I never built a company, I'd never been at a company.

[00:30:28] I'd been an academic my entire life. I had no idea what an org structure looked like. I had no idea what a one-on-one was like. It was just like completely jumping off a cliff and hoping for the best. And, um, and so I ended up doing that and it was definitely a huge and terrifying learning experience to build a company from the ground up, especially a company that was a rocket ship like Coursera, where we were on an exponential curve for quite a large amount of time during those first couple years.

[00:30:57] And at the end of those two years, Stanford [00:31:00] basically came and said, well, we have a cap on two year of two years on your leave of absence. So, are you coming back now? And I said, I can't come back right now. We're still building. And they said, well, you have to pick. And so, I picked, and I ended up resigning my endowed chair at Stanford.

[00:31:15] And my mother thought I was nuts because who on earth leaves an endowed chair at the world's top computer science department. But anyway, there we are. And I stayed at Coursera, I think for a total of about five years. And at the end of those five years, it was a good moment to sort of step back and reflect.

[00:31:32] And notably, if you think about the timeline, I left Stanford at the end of 2011, early 2012, which was just when the machine learning revolution was starting in 2012 with ImageNet and the deep neural networks and so on. I'd missed all of that and I'd been far too busy at Coursera to even pay much attention to what was going on.

[00:31:52] I said, yeah, there's a lot going on in machine learning, but I didn't have time to even track. And then in 2016, when I started to look around, it was [00:32:00] like, wow, machine learning is changing the world across pretty much every sector, but where it's not having much of an impact is in the life sciences. And one of the reasons for that, I felt as I still do today, is that there's just not very many people who speak both languages, who both truly understand the problems that really make a difference in biology and medicine.

[00:32:22] And at the same time, also understand what the tools can actually deliver and are able to bring the two together. There are certainly more now than they were when I started, but it's still a diminishingly small fraction of, say, machine learning researchers who really are able to take those insights and apply them to life science or the other way around.

[00:32:40] And so I decided that this was an incredible opportunity for me to make an even bigger impact and Coursera was in good hands and there's not really a lot of AI in Coursera, certainly not at the time, and I felt like if I was going to make an impact, this was the place where I could bring the biggest value.

[00:32:58] And so I ended up [00:33:00] at that time going to Calico, which is a drug discovery company within the Alphabet umbrella. I didn't really know a lot about Calico, but it was an incredible opportunity to work with unbelievably talented leaders like Art Levinson and Hal Barron and others. And I figured it's certainly at least a place where I could learn and work with wonderful people.

[00:33:21] And so I did that, and that was my first exposure to drug discovery. And when I looked at how drug discovery was done, even at a cutting-edge place like Calico, it was like, wait, this is how we make medicines? No wonder so few of them are successful. And so, I realized relatively early in my journey there that what I really wanted was to build, I mean, I'm, I'm an engineer, so I build products, I build things, and I wanted to build a, a system that would help us make better medicines faster.

[00:33:56] And it didn't make sense to build a platform like that within the [00:34:00] environment of a company that focuses on the biology of aging, which is what Calico's mission is. And so rather than trying to create a xenograft of, you know, these two companies that don't really make sense together, um, I ended up leaving Calico in February of 2018 and launching a company, insitro.

[00:34:18] If you think about the name, it's the synthesis of in silico, which means in the computer and in vitro, which means inside the lab. And really bring those two groups of individuals, these two ways of thinking, these two types of technology together into a single integrated whole that is going to allow us to discover and develop better medicines.

[00:34:38] And that's the vision behind insitro. Uh, it was founded as an end of one with a very substantial amount of capital from a group of investors who had actually been looking to make an investment in the machine learning enabled drug discovery space. They had done diligence on a number of companies, found all of them lacking in credibility or [00:35:00] something else.

[00:35:00] And so when they heard that I was kind of looking to build something and said, yeah, here's a hundred million dollars, do something. And um, so here I was with a hundred million dollars as an end of one, uh, without any network of peers or connectivity in the biotech ecosystem to build a team around me. So it was,

[00:35:18] let's just say challenging and terrifying in a very different way from the Coursera journey, which was my first industry foray. This wasn't my first, but it was a deep dive into a completely different ecosystem that I knew very little about. It took a while to get there, but um, I'm privileged to now have an incredible team of people with very complimentary types of expertise because coming back to the name and the vision behind the company, building this new kind of drug discovery company that requires truly an equal partnership between life scientists and computational scientists and drug [00:36:00] discovery experts, you really need to have a group of people who come in with a genuine intent to understand each other and work together and they all have a seat at the table, which is a very rare thing to find in this industry.

[00:36:13] And so is it fair to say, I always try and distill things down. Do you think of insitro as a new kind of drug company? Kind of putting aside like pharma baggage and things like that. Is that fundamentally at its core? Like you're making new medicines and you hope to carry them from inception all the way through phase three clinical trials.

[00:36:32] Is that kind of the vision for insitro and AI can sort of help at all, all stages? So, I think the answer is absolutely AI can and will help at all stages as we go from the concept of this is a disease we're trying to deal with a group of patients. We're trying to help all the way through the creation of a novel therapeutic hypothesis, turning that hypothesis into chemical matter, taking that chemical matter and going through the clinical [00:37:00] development and even ultimately beyond that.

[00:37:02] I mean, we, I think over time we're going to have to understand what our drugs do to patients in the wild, in the real world, and use that to inform the next stage of our drug discovery effort. So, all of that are places where I can help. As a small company, I think it's unrealistic for us, at least in the early stages, to imagine that we would take every single one of our insights and turn every single one of them into a phase three clinical trial simply because it's a very expensive process.

[00:37:30] And so I expect that we will partner some of those projects to other companies, whether they're big ones like pharma or other biotechs that have other capabilities that we don't have. That is certainly in our future, but our ultimate goal is to make at least some, and if we're successful over time, more and more of those projects all the way through, because I think that AI truly can enable us to become more effective and efficient throughout the life cycle of a drug.

[00:37:59] Can I get [00:38:00] your thoughts on how AI at insitro and kind of like biotech for more generally, like where the right sweet spot for what we can currently do is? And I'm going to give you two axes here. So, one axis is kind of like biological risk or uncertainty. So, like low down on this axis are biological things we understand.

[00:38:17] We have a target, we have a pathway. Things high on this axis are we don't even know sort of what the mechanism is. The other axis is kind of technological risk or uncertainty. So, like if you're low on the biological risk, but high on the technical risk, we know what the target is, but we don't know how to hit it.

[00:38:31] So is AI good for helping us sort of reduce biological uncertainty, or is it really good at prosecuting targets that we currently know exist, but we don't know how to hit them? I think AI is good for both and different companies have taken different trajectories. So, if you look at the field of AI-enabled drug discovery companies, broadly construed, you will find that the preponderance of those companies are actually, I don't know if I would call it technical risk, but taking targets that have [00:39:00] been reasonably well validated and turning them into chemical matter, and they each have their modality of expertise.

[00:39:06] They're a lot of the earlier ones were in the small molecule space. Now there's a growing number in the large molecule protein antibody space, I think, driven by the successes of AlphaFold and its follow-on successors that allow us to design proteins very effectively. There're even more companies now looking at novel modalities or more cutting-edge modalities like, for example, gene therapy where you can design the capsid, or um, people now with all of the excitement around RNA looking at RNA therapeutics, and I think there's a lot of companies in that bucket.

[00:39:41] The number of companies that actually focus, as we do, on the discovery of novel therapeutic hypotheses is actually quite limited, and I think there's a number of reasons for that. One is, biology risk scares people, it's also something that takes you a lot longer to know if you've been [00:40:00] successful.

[00:40:01] When you're in the context of a designing of a drug, there's usually a fairly well-established set of assays that you can perform on a drug to know at the end of whatever your two-year design period, if you know, if you're lucky, that you've succeeded. It has certain binding affinity, selectivity, solubility, whatever, and you know that you've succeeded.

[00:40:20] And everything downstream is, you know, you don't need to worry about that. For biology risk, ultimately success is when you've put the drug in a patient, and it actually helps the patient be better. And so, the timeline is much longer. The risk is, feels to a lot of people to be much larger. So, I think that is one element.

[00:40:39] And the other is that there's not a lot of training data in therapeutic hypotheses, and the very naïve approaches of we're going to rely on successful drugs as training instances for machine learning or AI models doesn't work. There's just not enough successful drugs, and we don't understand what it is that makes them successful.[00:41:00]

[00:41:00] So what we've elected to do is really design the problem in a very different way where we have different ways of generating training data. We also have a much greater reliance on unsupervised and self-supervised machine learning algorithms, where the need for supervised examples is much, much lower.

[00:41:18] But it, let's just say it's a much harder AI problem and a much harder scientific problem. And so that's why I think we're relatively, I don't want to say unique. There are a couple of other companies that try and do that, but it's certainly not the majority. Yeah. I want to ask kind of like a follow-up question about that, especially given your experience straddling both of these worlds. I know in medicine, when we kind of look over at what's happening in mainland AI and try and make analogies to what's happening in our world.

[00:41:49] So we'll say, like, we're going to train a big language model of the electronic health care record. And we reasoned by analogy quite a bit, or we'll say, you know, a patient's clinical record is actually just an image if you [00:42:00] squint and look at it the right way. And I think that there's a lot of this in biology too, where reasoning by analogy, where we're going to build an LLM for the language of life or for biology.

[00:42:09] In what ways do you think those analogies are helpful? And in which ways do you think that they can be hindering? So that's a great question. I do think that those analogies, and I would say it's beyond analogies, it's actual reliance on technical artifacts of actual products that is helpful because I think if we are going to train a, whatever, language model, or in a machine learning model, if you will, on, on medical images, the number of medical images that are available to us is usually quite limited, and if we don't see the connection to the models that were trained on cats and dogs and airplanes, and leverage those in our work, we're going to end up with performance that is very much suboptimal.

[00:42:55] So I think those connections are very helpful. At the same time, I [00:43:00] think that there are definitely places where that over simplistic view can lead to unintended consequences, um, where people, for example, don't really appreciate the challenges that you have with, say, batch effects that are very, very subtle sometimes.

[00:43:20] And you get misleadingly high performance on your supervised task because the machine learning of the model is latching on to something that is, you know, some weird, uh, artifact of the x-ray machine that took the image and this x-ray machine was used in this hospital and a different x-ray machine was in a different hospital and the patient population is just different and one of them has more of a certain kind of disease than others and so the machine is making wonderful predictions based on something that has absolutely no biological relevance.

[00:43:54] So I think that is something that certainly happens in other contexts as well but is much more prevalent in, um, in [00:44:00] the biology and medicine space. And then even more so, is the sort of recognition that the ability to extract insight from these models, because ultimately when you're building a predictor and all you care about is can it find images of cats and dogs for me on the web, you don't really care why, you don't care how, there's not a need to sort of ask what the model is latching onto as long as it's doing a good job.

[00:44:24] In the context of the scientific discovery world, the insight is often the thing that you care about rather than the quality of the predictions. That's one of the things, by the way, that makes, for example, the difference between discovery systems and diagnostic systems. In diagnostic systems, you can make the argument that all I really fundamentally care about is, is it making good predictions, albeit out of distribution?

[00:44:48] If you're doing discovery, the ability to trace back and understand something that is going to be the, whatever the therapeutic hypothesis is much, much more important. [00:45:00] So one last question before we get to the lightning round. Um, so again, a big analogy that we often make is an appeal to the scale hypothesis.

[00:45:08] And so for large language models, the scale hypothesis has continued to prove to be true. And just as a reminder, that's the idea that if you have an algorithm that's scaled with data and compute, you can just keep throwing more of both of those at it and keep getting better results. How do we think about that in the context of biology specifically?

[00:45:24] Because a lot of biological data is highly redundant. And I'm thinking of like genome sequencing data is highly redundant. And even if you have, if even if you've sequenced every person on the face of the planet, the sort of information per bit there is actually pretty sparse just because we're all very, very similar to each other.

[00:45:41] So how do we think about the scale hypothesis for biology and to sort of what extent is that a useful analogy? I actually believe in the scale hypothesis, even in the context of biology. I think if you look at images of cats and dogs and airplanes, you've [00:46:00] seen 50 images of airplanes. I don't want to say you’ve seen them all, but you see them in slightly different perspectives, from slightly different angles, with slightly different tail markings and so on.

[00:46:09] And you continue to learn. And yes, the incremental benefit of each new sample diminishes, but it’s still valuable, and that’s how we’ve gotten to the performance that we’ve gotten. I think we’re nowhere close to saturating the ability that we have to learn new insights from biological data. Now, you gave DNA as an example, so I’m going to argue about DNA, and then I’m going to talk about the broader phenomenon.

[00:46:36] Even in DNA, I think the more people we sequence, the more rare variants we discover that potentially are highly deleterious or, or highly, um, or sometimes highly protective, as well as interactions between different variants. Things that are protective only in a certain, um, environmental context, only in a certain genetic context.

[00:46:57] We’re nowhere close to saturating [00:47:00] the insights that we have, especially given the fact that the vast majority of the individuals that we’ve sequenced so far have come from European backgrounds. So, yes, we are very similar to each other. But there’s a lot of different other genetic backgrounds where we’re nowhere close to having found the relevant genetic variants.

[00:47:16] And we’ve seen that even in the, you know, recent publication from all of us, as well as many of the others, that there’s more and more genetic insights that can come out. Now, when you think about the broader spectrum of data that can be collected about even just a single human organism. You have single cell RNA seq from different cell types in different states at different ages under different environmental conditions.

[00:47:41] That’s just RNA seq. You can also think about protein levels, and we haven’t gotten to protein levels at the single cell. We’re starting to scrape to, you know, the boundaries on that, but it’s certainly not at scale. And then going down into the level of individual proteins, going up to the level of, you know, we’ve now everyone’s excited as [00:48:00] they should be around spatial biology and the interplay between different cells, the number of data modalities, the number of ways that we have to measure biology and the number of distinct biological contexts with,

[00:48:12] that we as a human live in and that exists within even the body of a single human is ginormous. It is way larger than what I think the complexity that we've trained the large language models that we're currently leveraging in these in the more traditional LLM sense. What we lack are data collection approaches that achieve that scale.

[00:48:33] I'm just really excited to be living in this time because the number of ways that we have to measure biology quantitatively and at scale is increasing, maybe not quite as fast as the capabilities of AI, but you know, we're kind of maybe a few years back on that curve, but you can see that exponential curve.

[00:48:53] And I think the synergy between those two is going to just unlock just an incredible [00:49:00] tidal wave of insights as we start to bring those tidal waves together. Well, that I think is one of the most forceful endorsements of the scale hypothesis for biology that I've ever heard.

[00:49:11] So I, you know, consider me convinced. I'm glad. So, let's go to the lightning round.

[00:49:23] So the lightning round is, we'll ask you a bunch of different questions. The goal is to keep the answers relatively brief. Some of them are silly. Some of them are serious. And we'll let you decide which one is a silly one and which one's a serious one. And can I not answer ones that I don't want? So, abstention is not an option, unfortunately.

[00:49:39] Well, I don't know. I'm an independent human being. I can decide later if I believe that. That is, that is true. That is true. That's fair. Um, so the first question is, what's an example of a frivolous thing or something that you do just for fun? I do the New York Times crossword puzzle and I love to hike.

[00:49:58] Nice. Nice. [00:50:00] What is your all time favorite book or movie? I'm often asked that and I usually refuse to answer because I have too many favorites in different contexts and different moods. So I'm going to abstain. How about a softer question? What is a recent good book or movie that you read or watched? I really liked *Oppenheimer*.

[00:50:27] It's a recent movie that I watched. It is a wonderful synthesis of the challenges and opportunities and risks in science, as well as a very human story about a scientist and challenges that he faced. And so, I thought that was a really good movie. I need to see that. It's on my list and it's just won best picture, right?

[00:50:50] Best picture and best actor. It's one of a bunch of them. Yeah. It's cleaned up. Yeah. Yeah. Cleaned up. All right. Excellent. Um, what is one of the core guiding [00:51:00] principles of your life? I'm gonna name two that are intertwined. One is that I believe that it is the responsibility of each of us to try and leave the world a better place by virtue of us having been here.

[00:51:16] And the more you were born to privilege, the greater your responsibility is. And I also believe strongly in leverage, which is the fact that the benefit that you bring should focus on places where you can be disproportionately impactful. There are a lot of things that I could do, even I as a human being, I as an AI researcher.

[00:51:41] The reason I'm doing what I'm doing today versus the many other things that an AI researcher could do is because there's a lot of very talented AI researchers who can work on computer vision and robotics and natural language and other things, and probably do it as well or better than I can. AI for biology is a rarefied [00:52:00] group, and I think I can bring disproportionate impact by doing that.

[00:52:05] I don't want to interrupt the lightning round, uh, too much, but, uh, I think this has been a real recurrent theme amongst our guests, which is speaking both languages and not being an AI researcher who sort of dabbles in biology or medicine or a biologist who dabbles in programming, computer science, but really committing to both skill sets and really having both skills in the same mind, um, as, as leading to success.

[00:52:29] Uh, so it's great. You know, Ziad Obermeier spoke very, very cogently about this and I think several other guests have as well. So great to, great to hear that. All right. The next lightning round question, what is harder for AI, biology or medicine? Biology. Um, I think because there's, the complexity there is very, very large and it's incredibly intricate.

[00:52:57] I think medicine obviously is [00:53:00] complex as well, but I would say the, the bar in terms of really impacting clinical care by things that are relatively simple for AI to do is, is lower, so I think over time we'll start to hit the point where those relatively simple things are not enough and we need to go beyond and then the equation might switch, but right now I would say biology.

[00:53:32] Do you think some of that is tied to explainability or having to understand the mechanistic or the causal diagram that is required in biology versus in medicine, we have many examples of things that work well, where we don't exactly understand why they work well. So, is part of that wrapped up in explainability?

[00:53:54] I think part of that is wrapped up in explainability, in complexity, in [00:54:00] multimodality, and again, coming back to there's a lot of things that we currently do with patients that are so incredibly suboptimal. The fact that we do not tailor our treatments to the specifics of an individual patient, even though we know that one size fits all doesn't work.

[00:54:19] And I think there's just a tremendous opportunity to take these rich data that we're able to collect, although oftentimes we don't and don't record them, but we could around patients and really impact the care that we provide to those patients is going to be, um, there's a bunch of, I would say, relatively simple things that we could do, even if it's only helping clinicians take better notes so that they have more time to think about their patients.

[00:54:47] There's just a lot of, I don't want to call it low hanging fruit. Nothing in this space is low hanging, but, um, a lot of, you know, things that we could build that would be very, very helpful to how we care for our patients, um, [00:55:00] and I think biology is just a longer journey. All right, um, so I think this is going to be our last, last lightning round question, Daphne.

[00:55:09] So it's one that we like to ask lots of folks. So, if you could have dinner with one person alive or dead, who would it be? It's going to be really lame. Um, but Albert Einstein. Oh, not, that's not lame. How could Albert Einstein be lame? Well, because it's so cliche. It's so cliche, right? It's all relative.

[00:55:27] Sorry. Yeah. I'll just say it. I'll show myself. I'll show my way out. Yeah. All right. Very good. It didn't land well. It didn't land well. There's some smiles for the, for the recording. There's some smiles here and some smirks. Some golf claps happening now. Um, okay. So, I think we're going to wrap up with one big picture question before we let you go.

[00:55:49] We've heard you made comments recently about how it's hard to know where you are on the exponential curve of AI progress and how understanding where you're on this curve or the [00:56:00] difficulty with understanding where you are leads to bad intuitions about what the next five to 10 years is going to look like. So, could you sort of expand upon that and help us like, um, maybe give us a more sensible intuition about what it might look like in the next five to 10 years?

[00:56:16] I'm going to answer that question by explaining why it's so hard to know where on the exponential curve you lie and the fact that you are in fact lying on an exponential curve. If you go back a decade, a lot of people, certainly outside of the field, but I would say even within the field, would look at how much progress had been made in the last, whatever, two, three years, and they would extrapolate effectively a linear, um, a linear interpolation of those points and say, this is what we're going to be in 10 years.

[00:56:43] And inevitably that was wrong. But even then, it didn't sink in to people that this wrongness, that you keep doing this year after year and year after year, you're wrong, um, didn't really sink in because the wrongness was relatively limited in the early days. As you start to get to the [00:57:00] point where you're even your linear extrapolation, uh, is, gets you to these ridiculous places.

[00:57:07] And that's when I think people realize that, oh my God, AI is suddenly here. But it's not suddenly. It was here, you know, um, 10, 15, 20 years ago. It's just that, it's just that we didn't realize we were on that exponential curve. I think we're in the same place. I mentioned that earlier in terms of the ability to interrogate biology at scale in a quantitative way.

[00:57:29] And people often don't realize that we are on a similar exponential curve because we are at the earlier stage. But when you are on an exponential curve, the base of the exponent, just how quickly that exponential curve goes up is very hard to sort of extract and small differences make a very significant delta in where you will be in five to 10 years to the point that, you know, even, even I, for example, I, I was, I realized we were on an exponential curve, [00:58:00] um, quite a while ago, um, I did not predict the large language models.

[00:58:04] I did not predict where we would be in 2023 in terms of the ability to perform tasks at this level that involve language. So given that I wasn't able to predict it back in, say, 2020, that this is where we would be in 2023, even recognizing that we were on an exponential curve, I'm not going to shame myself or embarrass myself by trying to make predictions about 2026.

[00:58:32] So maybe I'll just sneak one last question in here, a big picture question. And I would ordinarily like to end on a positive note, so you could spin this into something positive if you so choose. Is superintelligence or existential risk, existential risk from advanced artificial general intelligence. Is this something that you take seriously?

[00:58:54] You know, we've had some very prominent scientists and folks in the industry and developers of [00:59:00] some of these AI models too, including folks like Jeff Hinton, who've really been very involved from the early days, voiced strong concerns about the path that we're on and about what the capabilities are of the next wave or the next wave after that of these models.

[00:59:17] Is that something that you take seriously? Um, or something that, uh, you're not, you're not so concerned about. So, I will say that this is definitely not ending on a positive note. I will answer the question, but then maybe we can find a positive note to end on. I'm not worried most about the form of existential risk whereby computers develop a level of autonomy and whatever sentience that makes them want to wipe out humanity and the *Terminator* scenarios that emerge from that.

[00:59:48] That is not my biggest concern. My biggest concern is that we are unleashing and democratizing a very large collection of very powerful tools that humans who do [01:00:00] actually do evil things to other humans are able to leverage in ways that are we're already seeing today. And what I think is enabled by these tools is evil at scale.

[01:00:11] And so one of the things that I actually think the existential risk conversations around the *Terminator* scenarios is doing is it's diminishing people's ability to focus on the much more immediate, in fact, I would say present today risks of humans using those tools to facilitate human trafficking, child pornography, the erosion of truth.

[01:00:36] I mean, I can go on and on and on to talk about all those risks that are, because we're not focusing as much on them, I think, are being allowed to thrive. And it's going to be in the same way that cybersecurity is. It's an arms race. The bad guys develop a technology the good guys have to develop probably using the same set of tools. In this case, AI is going to be the defense as well as the offense. [01:01:00]

[01:01:00] How do we use AI to detect? Deep fakes to prevent the erosion of truth to create watermarks that can't be forged. I mean, there's a whole bunch of things that we could and should be focused on as a community, but we're not. So, I think that to me is one of the biggest challenges of this narrative. Yeah, I agree.

[01:01:21] And I really like your focus on the present term and AI is obviously dual use. And I think you did a good job at enumerating what some bad uses of it may be. In an effort to fulfill your request to have this end on an optimistic note, what gives you hope? What gives you optimism about AI? How will it make our lives better?

[01:01:39] And what keeps you working on, on these problems? So, I'm going to divide the answer into two. I think that AI is going to make our lives better and easier in ways that I think are pretty much visible to most people today, you know, agents that do our bidding so that I don't have to deal with a lot of the [01:02:00] minutiae of day-to-day because an AI will be able to take my verbal commands and go and deal with those minutiae for me, and I will be able to have more time to do other things.

[01:02:10] It will empower people who might have really great ideas, but aren't able to write, or aren't able to draw, or aren't able to make movies, to create in that way. I think there is a lot of, or to program, we've talked about that earlier, there is a huge empowerment that is going to come from the availability of these AI tools. That's one element. The other element, which I think is often not as visible to the broader communities, the AI for what you might think of as deep tech or deep science or solving really, really hard problems that humanity has been grappling with for decades, centuries, and that we're not able to solve on our own.

[01:02:50] And that can be things like how do we bring better therapies to patients, or maybe how do we address climate change, um, and carbon capture and [01:03:00] things like that, where, um, we need all the help we can get and those tools are going to be super powerful. So, I think, yes, we do need people who work on the more, if you will, consumer facing aspects of this.

[01:03:11] The AI agents, the creative tools, the programming bots and all that. I think that's really great. But we also need people who are willing to kind of take the challenge of grappling with things that are going to take a longer time to solve. It's not, you don't get the immediate gratification of seeing people use your agent bot, but you are doing something that potentially can be transformative to the benefit of the world.

[01:03:35] And I think both of those are important opportunities of this technology. All right. I think that's a great note to end on and, uh, Daphne, thank you so much. This was a wonderful conversation and we really appreciate you coming on *AI Grand Rounds*. Thank you so much, Daphne. We know how busy you are, and this was super great.

[01:03:53] Thank you, Andy and Raj, you asked some really insightful questions, including ones that I'd never been asked before, which is [01:04:00] unusual, and it was a wonderful, far-ranging conversation, so thank you for taking the time to speak with me.