AI-GR Pod 8 Ziad Obermeyer 07.24.23

[00:00:00] I think there's this tendency, everyone gets angry at the algorithms. It's like, oh, this algorithm is by, but the algorithm, the algorithm is doing exactly what we told it to do. It's not the algorithm's fault that costs are lower for black patients than for white patients at the exact same level of health.

[00:00:19] That's our fault. That's society's fault. But now the algorithm is mirroring that deep bias in society because we chose. to have the algorithm mirror that deep bias in society. So that's where the bias came from.

[00:00:36] That was Dr. Ziad Obermeyer of Berkeley, reflecting on his research that dissected the source of racial bias in a major healthcare algorithm that he published in a landmark paper in the journal Science back in 2019. I wanna welcome you, our listeners to another episode of N E J M AI Grand Rounds. I'm your co-host Raj Manrai, and I'm delighted to be here with my co-host Andy Beam.

[00:00:59] [00:01:00] Andy, I always really enjoy speaking with Ziad. He's truly creative and rigorous in his research, and I also find him to be an amazing communicator of complex ideas. In addition to enjoying hearing about his pioneering work on algorithmic bias and medical decision making, I thought his comments about the recipe for great interdisciplinary work were really thought-provoking.

[00:01:18] So to do truly creative and impactful work in medical AI, can you pair a great clinician with a great AI researcher, or do the two skills really have to co-exist within the same mind? I think Ziad, both articulates and himself, is a strong argument for the skills having to coexist. And all in all, this was a really fascinating discussion.

[00:01:36] I totally agree, Raj, and that was really well said. And I, you know, when I think of Ziad, I think of him as kind of a renaissance man. His work cuts across so many different disciplines from medicine to economics, and as we saw in this episode, he really investigates how these methods impact the lives of patients.

[00:01:53] And as you mentioned, he has this really unique ability to access and distill very complicated topics in a way [00:02:00] that I find inspiring. So it was a great pleasure to talk with him about his research. The N E J M AI Grand Rounds Podcast is sponsored by Microsoft and Viz.Ai. We thank them for their support.

[00:02:13] And with that, we're excited to bring you Dr. Ziad Obermeyer on AI Grand Rounds.

[00:02:19] Welcome to AI Grand Rounds Ziad. We're super excited to have you on today. Oh, it's, uh, wonderful to be here. Ziad. Let me also welcome you to the podcast. It's great to have you on. This is a question that we like to get started with, and please forgive our AI puns. Could you tell us about the training procedure for Ziad's neural net?

[00:02:37] How did you get interested in AI? What data and experiences led you to where you are today? I think the, the proximal experiences are definitely my experience, training and practicing as an emergency physician. I think if you are working anywhere in the medical system and you're vaguely paying attention, you are making mistakes all the [00:03:00] time.

[00:03:00] You're seeing other people making mistakes and you're trying to learn from those errors and minimize your loss. And I think a lot of the things that we need to do as practitioners in the healthcare system are things that look very much like things that machine learning algorithms do well. We're often trying to make diagnoses.

[00:03:20] What is diagnosis? It's like I look at a patient, I see a bunch of stuff about that patient and I try to convert that into a probability that they have, you know, illness one, illness two, illness three. When I'm working in the emergency department, I'm often trying to think about, well, do I need to bring this patient into the hospital or can I send them home?

[00:03:40] Are they gonna do okay at home? Um, that's a great prediction problem. That depends on their likelihood of bad things happening to them at home. And so wherever you are looking in medicine, you're seeing these prediction problems like cancer metastasis, great prediction problem, diagnosing the cause of pain on an x-ray, great prediction problem.

[00:03:59] So I [00:04:00] think that's the thing that got me really excited was working as a doctor, seeing all of these really tough decisions that we're having to make doing not as well as I would've liked on those decisions, and trying to find tools that would make me and my colleagues better. If, if I could, hopping on the theme of prediction models, I feel like your career breaks my own prediction model for what physician scientists typically do.

[00:04:24] Uh, you've written very broadly as we'll talk about on economics, um, things that sort of border on ethics and philosophy, also technical contributions. So could you talk about your training or educational background that sort of helped you tackle such an impressively broad array of subjects? Oh, thank you for the kind words.

[00:04:42] I, I agree that a large language model would not be able to predict the next word, uh, equivalent in my career because I often really wasn't able to predict it either. So, to go back in time a little bit further before my medical training in college, I actually studied history and philosophy and I did [00:05:00] a master's in that.

[00:05:01] You know, it comes in handy. Sometimes. Um, but it has had this pretty big influence on my worldview. I'm very interested in how new scientific fields form and how that is often a combination of genuine scientific and technical advances, but also a lot of social things and how groups of people come together and build a scientific field.

[00:05:24] Uh, it's a great example of the human enterprise and all of its beauty and fallibility. So I studied those things. I thought I might wanna be a historian and or a philosopher. It turned out I wasn't that great at either history or philosophy. And so then I thought maybe I'd like to just make a lot of money.

[00:05:42] And so I went into management consulting and that was actually really fun. It wasn't for me. I only spent a year doing it. But, but it taught me a lot about how to think strategically and I guess economically, how to think about the trade-offs that people and businesses face and how things [00:06:00] are often a lot harder on the inside than they look on the outside.

[00:06:04] And then I went to medical school, so I think I, I had that, that set of interests and that worldview and the background going into medical school. And I think that might have been one of the things that made me more attuned than other people to the huge opportunity that I see in applying AI and associated tools to, to medicine.

[00:06:22] Was there something specific about medicine that made your interest in it more durable than some of your other previous interests? It's a great question and part of the reason it, I've thought about this a lot is because I actually really didn't like medical school and so I had I been predicting in year one or year two of medical school, I would not have predicted that I would finish medical school, let alone go on to really love being a doctor.

[00:06:46] I think there's something about medical school that is a little bit infantilizing you, you just have to memorize a lot of stuff and I think you have to do that for good reason sometimes. But I think there's also this [00:07:00] huge mismatch between what you learn in medical school and what you need to be a doctor or a scientist in this area.

[00:07:07] And so I really struggled with that in medical school and I thought about not doing a residency, but largely I think thanks to inertia and probably not very good decision making. I ended up doing residency anyway. And I had enjoyed the parts of my medical school time that happened in the, um, in the emergency department.

[00:07:27] So I was like, whatever, I'll just do this. It's a cross section of medicine. So that, you know, I was, I was actually hedging. I was like, well, you know, if I end up not doing medicine later, then at least this will be good training. It'll get me exposed to lots of different parts of medicine. But then it turned out that I really, really liked being a doctor.

[00:07:44] And I think there's something about that exposure to the real world and the problems of patients that I think it's shaped. The problems that I work on in my research as well. And so I think it's a really important, interesting job. There's [00:08:00] drama, you know, it sometimes it does feel like being on a TV show.

[00:08:03] You go into a room and you can ask anyone anything about their life and they'll answer the question because you're trying to help them. And you have that bond of trust that you've forged. It's an amazing job. And then there are all of these really interesting intellectual questions that come off of it that are at the intersection of statistics, economics, behavioral science, computer science.

[00:08:25] So I just think it's the best job in the world, the job that I'm doing at the intersection of these fields, and I feel incredibly lucky to be doing it. Suzi, maybe one more question about your background before we jump into your work. So I believe we actually met through our mutual friends, Sachin Jain, maybe a decade or so ago at a great Cambridge, Massachusetts bar, the Plow and Stars.

[00:08:47] I was a grad student at the time. I think you were a new faculty member and you were pitching a very interesting project on machine learning for emergency medicine. And so first I have to give you credit as, uh, being on the machine learning wave [00:09:00] way in advance of much of the field. And I think you've really chosen problems well, uh, and we're gonna discuss some of those projects in detail.

[00:09:07] But maybe as sort of a background question in general thoughts, I can just ask you to elaborate on some of the cultural differences that you've had to navigate across these different fields. So we've already talked about economics, computer science, medicine. These are very different fields in some ways, right?

[00:09:07] And they have, uh, very different approaches to solving problems. How do you think about that in terms of your training? What have you self learned? What have you formally learned? What have you had to learn on the fly to solve the problems that you do? I think one of the interesting things about medicine as a field is that there's no field.

[00:09:07] Like there's no, I mean, you go to medical school, but medical school is, it's, it's a little bit like, learning to become a plumber or do some trade. It's like that there's a certain amount of really concrete stuff you have to learn how to do and you learn how to do it. But then there's this scientific backend to medicine that is very much at the intersection of a lot of these fields.

[00:09:07] Um, and I think that's, it makes it hard to do this work because unfortunately no one field is gonna give you everything that you need. And so one of the things that I actually admire about both of you and the work that you've done is that you've really invested a lot in learning the content. As the computer scientists say the domain expertise, and I think that's really a great model for doing work in this field, or, or it's at least what I've tried to do.

[00:09:07] So I, I've spent a lot of time doing like problem sets in textbooks and, you know, we, and this is, this is on your own right out outside of a formal class. Yeah. It, it was, and I think it was because there were just some questions I was really interested in that I wanted to be able to solve myself. I think, you know, it's different.

[00:09:07] I think what the work that both of you do and the work that I try to do. I think it's different from most people's model of how you do cross-disciplinary work. I think most people's model is like, you get a computer scientist and you get a doctor and you put them in a room together and then they'll produce something good.

[00:09:07] What doesn't really work about that model is that you put those two people in a room together and they often don't have anything to talk about. And in fact, they don't even really speak the same language. Uh, and so the work that gets done is kind of, it has this superficial feeling because nobody can really get into the interesting part of either one field.

[00:09:07] So even though it seems ridiculously inefficient, I think that the right model is the doctor actually needs to learn the computer science and the statistics and the statistician or the computer sciences actually needs to make big investments in learning about health. And I think that's, even though again, it doesn't seem, it, it seems like you should be able to do it the other way and it's like, wow, this is really a lot of work for both people to learn both things.

[00:09:11] But I actually do think that that's how the best work that I've seen in this space and the work that I've done, that I'm most proud of, um, that's how that's happened. So that's a great transition to your work and your research. So you've worked on a lot of different topics, uh, at the intersection of machine learning, economics and medicine.

[00:09:29] And I think one of the very impactful, uh, threads of work that you've led is on the topic of algorithmic bias. I thought we could start off with your 2019 science paper which is titled Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. So this paper, to me, was really a very powerful and clinically important example of the importance of how you formulate a machine learning problem, how you choose what your task is, what your target is to predict, and what this encodes [00:10:00] both about society and about the generalizability of a model that is trained to predict some outcome.

[00:10:06] So maybe we could just start off by you summarizing the paper, telling us what you set out to do, and maybe also you could share, how this project got started, how you got interested in this particular problem. I think to set the stage for this work, one of the things that I think a lot of people listening to this will know is that health systems have made huge investments in population health.

[00:10:30] So I think the old model was, you know, you just treat one patient at a time. They come in, you treat them, you treat and release. And I think that as the healthcare system has changed, health systems have changed to a view where they're responsible for a population of patients and they're responsible for caring for those people.

[00:10:48] They're responsible for the costs of those people. They're just like, take charge for this big population of patients. And one of the things you need to do in that setting is you need to make sure that the people that you're taking care of, [00:11:00] Don't get sick. So some of those people are going to get sick over the next year, and if you knew which patients were going to get sick, you could actually help them today, but you often don't know.

[00:11:10] And so there's a range of things you can do. High risk care management is the, the big category of things, but it's basically just extra attention. Like a nurse practitioner will come to your house, they will refill your medications, they'll examine your foot for ulcers, they'll, they'll do all sorts of things to prevent.

[00:11:27] You from ending up in the emergency department and seeing someone like me and getting admitted to the hospital for antibiotics or diuresis or whatever. So that's the kind of the big picture and, and you can see where algorithms could be super useful for this task because what you're trying to do is, okay, I see a bunch of people, I have a bunch of data on them, some of them are gonna get sick.

[00:11:45] I need to find these needles in this haystack so that I can get on top of their care today. And so this has become a huge industry. So we studied one commercial product. Um, it's an algorithm made by a analytics subsidiary of [00:12:00] a large insurer. That product is used by their estimates to make these decisions about who gets access to extra care for about 70 million people every year.

[00:12:09] The family of algorithms that works just like the one we studied about 150 million people per year. So like the majority of the US population is being run through one of these algorithms. And it's not just if you go into the hospital, you can be sitting on your couch and these algorithms are churning through your data.

[00:12:25] Gating, whether you get access to extra help with your chronic conditions or not. So that's the the landscape we, because at the time I had a partnership with a academic health system in the Boston area. We had access to the algorithm risk scores from this one commercial product. And we also had access to this rich data set of what happened to people.

[00:12:49] And so the way this health system was using the algorithm was three times a year, they would just run the algorithm that would generate a list of high risk people and then a small fraction of those people [00:13:00] would get fast tracked into the high risk care management program that gives them extra help with their health.

[00:13:04] The next 50% approximately would get shown to their doctor. So the doctor gets to decide, so there's a human in the loop deciding is the algorithm right or not? And then the bottom 50% just gets screened out so they don't get offered access to this program. So the score was really dictating whether you got access to help or not.

[00:13:22] And we were very interested in whether that score was biased. So one of the things that I learned from this project is that if you wanna study bias, you have to start by defining bias. Because even though, you know, I think in a lot of settings is very clear what bias looks like, especially in human settings where you see discrimination with algorithms, it's a lot harder to pin down what bias is.

[00:13:47] And so having a crisp definition of bias was really important for us to study the extent of bias, uh, in this algorithm. And so, given how that algorithm was being used, what we [00:14:00] thought was okay, so the algorithm score is determining whether you get access to help or not. And so given two people with the same score, they should need the same amount of help.

[00:14:09] They should have the same health needs, so, And the color of their skin shouldn't matter. But unfortunately, that's not what we found. So what we found was that if we took those two hypothetical patients with the same score, and we followed them over the year after the algorithm made its prediction and generated that score, the black patients did far worse than the white patients.

[00:14:29] They had far more flareups of chronic conditions that landed them in the ER and in the hospital they had worse blood pressure, they had worse kidney function, they had worse hemoglobin a1c, like everything was worse, even though the algorithm was viewing them the exact same way and gating their, the, the same level of access to high risk care management and, and extra help with their health.

[00:14:51] So we had a precise statement of what bias would look like. The algorithm failed that test. And then the question is, well, why did it [00:15:00] fail that test? And it failed for a very, you know, a deceptively simple reason related to how the algorithm was built. So, What we wanted to do in this setting was find people who were going to get sick, but in your data set, there's no variable called sick.

[00:15:19] There's a bunch of data in there, but they're not one variable. So you need to figure out what exactly is my algorithm predicting what variable? And the variable that the people who built this algorithm, and a lot of other algorithms like it chose, is a variable called total medical expenditures costs.

[00:15:36] So looked ahead to see, okay, starting from this point, how much is the healthcare system going to spend on a given person? And let's use that as a proxy for how much care they need. So how much care they got. Denominated in dollars was the proxy. How much care they needed was what we wanted. But as anyone who has tried to access the healthcare system in this country, but in [00:16:00] many other countries knows needing healthcare and getting healthcare are not the same thing.

[00:16:05] And there are certain groups of people for whom that is particularly true. And those people are disproportionately non-white and they're disproportionately socioeconomically disadvantaged. And that was kind of a known fact, like we know that not everybody who needs healthcare gets healthcare, but translating that known fact into algorithm design, unfortunately didn't happen in this case or many other cases.

[00:16:27] And so we. Huh everyone. I think there's this tendency, everyone gets angry at the algorithms. It's like, oh, this algorithm is bi. But the algorithm, the algorithm is doing exactly what we told it to do. It's not the algorithm's fault that costs are lower for black patients than for white patients at the exact same level of health.

[00:16:45] That's our fault. That's society's fault. But now the algorithm is mirroring that deep bias in society because we chose to have the algorithm mirror that deep bias in society. So that's where the bias came from. Do you know if the [00:17:00] manufacturer tried to predict health state or need as opposed to costs before deploying this model in practice?

[00:17:08] So I don't think they did. And we actually, um, when we first got these results before we published, we struggled a little bit for what do we do with this result. We had some suspicion that this paper was gonna get attention. It ended up getting a lot more attention than we thought. But what we actually didn't want is for that company to learn about this when it was covered in a newspaper or, or something.

[00:17:34] And so we, because we actually just wanted to make this better, we wanted that company and, and all the other companies to actually change the way that we're building these algorithms. So we actually reached out through our networks to people at this company, and we ended up connecting with the technical team that built and managed this algorithm.

[00:17:51] And what I came away with from that work is that nobody, not this team, but not anyone who was in this loop of people buying the algorithm, [00:18:00] applying the algorithm to their patients, nobody had thought about this problem. And we all use cost as a shorthand for need. Like we do it in policy circles, we do like it.

[00:18:10] This is very, this is a very widespread error in health policy circles of using costs as a proxy for need. So I don't even think this was on anyone's radar screen, but when we pointed out the problem, they were really motivated to work with us to fix it. So we actually worked with that technical team to build an improved version of the algorithm that predicted a bunch of outcomes that were much more closely grounded in health than in cost.

[00:18:35] And we dramatically reduced the amount of bias and the algorithm that resulted from that process. And so my first. Reaction when I saw these results was, yeah, this kind of tracks, we say we're trying to take care of patients, but actually we're trying to reduce costs. And so of course the algorithm was predicting cost.

[00:18:52] That's what everybody cares about. But I actually don't think that's right either, because even if you just cared about reducing [00:19:00] costs, You still wouldn't predict costs and target those people. There's a bunch of high cost people that you can't do anything about. And in fact, in a lot of these, like when you look at the lists of patients who score very high and need extra help, what kind of patients are they?

[00:19:15] Well, they're patients with metastatic cancer. High risk care management is not gonna make a dent in the cost of metastatic cancer. That is just a cost we bear as a health system and as a society. People with total knee replacements, they cost a lot of money in that year, but they're not gonna keep costing money.

[00:19:32] And if they do, again, there's not. So we're we're interested in preventable cost, not total cost. And so if you're interested in preventable costs, it's really not clear that you want to be predicting total cost. You wanna be predicting these avoidable costs that come from emergency department visits that could have been handled over the phone, hospitalizations that could have been prevented, had someone's diuretic dose been increased in time, like things like that.

[00:19:56] So I think that process of going through the [00:20:00] problem formulation, as you mentioned, it's super, super important, but it's often an afterthought when we build algorithms because, you know, often, frankly, because it's such a, a pain in the ass to get the data and clean it and do all this stuff, that by the time you have the data you're like, all right, you know, let me just pick a variable and predict it because this has already taken six months and I'm way behind schedule.

[00:20:19] I think you're totally hitting on something that is, uh, so central, right? Which is that getting the data, accessing it, cleaning it up is often the bottleneck, the challenge, and the really big hill that we have to climb in medical machine learning. And I think the other thing you're, you know, you're pointing out there's a lot of attention on algorithms, right?

[00:20:40] But all this work beforehand that goes into it, that can actually take up most of your time. Um, so you mentioned that you contacted the manufacturer and that you worked with the manufacturer to then change the way that this prediction was happening, right? To change the model itself. I'm really impressed.

[00:20:55] You know, there's sort of the technical side of this paper, but there's the social side of this paper as [00:21:00] well, which is navigating that process to even get the risk scores, get the data, use it in this way, and then to work with the manufacturer. I think there's a lot of well-intentioned medical machine learning researchers who are gonna be listing.

[00:21:11] We're thinking about similar challenges where they're working with proprietary algorithms or technology. What advice do you have to them about how to navigate the social side of this process in addition to the technical one? I think first it's important to realize that it's such a luxury to have access to the scores.

[00:21:28] I mean, in, in criminal justice and in lending in a lot of other areas where people are trying to study algorithms and algorithmic bias, researchers just can't touch the scores. It's an amazing luxury to be working in this area in health where the scores are sitting in a data warehouse alongside all of the outcome data that you can use to validate the prediction.

[00:21:49] So it's a great area to be working in if you're coming at it from the health angle and it's, it's a lot easier than people who are trying to do this in other areas. I think what [00:22:00] I learned from working, at least with the, the technical team at that company is that, and, and it's, it's hard to remember when you're studying something as awful and, and shocking as bias at, at a large scale, but.

[00:22:17] What I tried to remember and, and what was born out in all of my interactions was that most people actually really wanna do the right thing. Like we were on one call with this team, it was like early people hadn't joined, and this one person on the team said, I am so glad that you all found this problem.

[00:22:37] Um, I took this job instead of a job at a tech company because I actually wanted to make a difference and I wanted to work in health and I wanted to get people care that they needed. And this is the last thing I wanted to do. And it's so great that we have the chance to fix this thing. So, you know, it's possible these people deserve an Academy Award and they're actually evil geniuses and they got caught.

[00:22:56] But I don't think so. I think people mostly [00:23:00] wanna do the right thing, and we're all learning about this together. This is a very new field, that's a new area. We don't have great understandings about how bias gets into algorithms, so we're all trying to figure this out together. And I think having that positive attitude at least made the work a lot more satisfying for me.

[00:23:18] And I think it was reciprocated in working with a lot of other people who wanna affect positive change as well. And I just would wanna say the, like in addition to highlighting these, how. Algorithms can operationalize structural biases on a huge scale. I think for me it also highlighted the urgency for really thinking about these things, cuz we tend to talk about the impact of AI in the future tense.

[00:23:42] Will AI replace doctors? What will happen in the future? But as your work shows these simple forms of ai, if you wanna, you know, be generous with the definition of ai. Already touching hundreds of millions of lives. And so this isn't something that we can kick down the road and worry about later. We actually have this problem now and it's only gonna become more [00:24:00] pressing over time.

[00:24:01] Yeah. Thanks for bringing that up, that it was very surprising to me how widespread these algorithms were on this backend side of the healthcare system. You know, I think you go into your doctor's office, you're sending faxes, you're filling out paper forms, you're like, what the hell? This is like, there's no AI going on here.

[00:24:18] And that's right. Largely on the clinical side, but on this population health management side, on a lot of other operational sides like clinic bookings, like things that have a direct impact on health, these tools are already in very, very wide use. We a after, because we got attention for that study, we got a lot of inbounds from health systems, tech companies, a lot of others that were interested in evaluating their algorithms for bias.

[00:24:44] So we did a lot of follow up work and case studies and efforts to diagnose and fix bias and other algorithms. And these things are. Everywhere and everywhere we looked almost we found bias in the algorithms that were being used. So I think that's a great transition point because the [00:25:00] next paper that we wanna talk about kind of flips that on its head and talks about how actually machine learning may be an opportunity to reduce healthcare disparities.

[00:25:07] So I'd like to talk about your paper, an algorithmic approach to reducing unexplained pain disparities in underserved populations where you go after longstanding pain disparities between different populations and actually show how algorithms can find things that explain sort of self-reported versions of pain.

[00:25:24] So could you, could you set that up for us and help us walk through that and, and was this a. Actually a continuation of the paper that we just talked about or, or like sort of what was the genesis of this one? So this paper was led by my colleague Emma Pearson, who's a fantastic computer scientist at Cornell and, and she incidentally is a great example of a computer scientist who's made huge investments, not just in learning about knee x-rays and MRIs and everything she had to do for that paper, but also learning about a ton of other areas where she's applied algorithms from covid to criminal justice to [00:26:00] lots of other things.

[00:26:00] So I think she's a really shining example of the kind of cross disciplinary expertise that we were talking about earlier. So this paper actually came out of a bet. I was at a conference and my co-author and friend David Cutler, who's a health economist, was presenting some research that he was doing on pain and he was looking specifically at knee osteoarthritis, which is one of the most common causes of pain.

[00:26:26] And he was preventing data that. Showed these very striking disparities in pain that exist across race and socioeconomic gradients. And so one of the many injustices in ssi, it is not just socioeconomic injustice, but if you think about the injustice that comes from the fact that some people experience pain daily, much more than others, that's just like another crazy and visceral kind of injustice that is largely missing from like public statistics.

[00:26:57] So we just like don't capture data on pain, but it's [00:27:00] really shocking if you just think about what that means for humans on a day-to-day level. Okay. So the research that David was presenting is a great example of a pretty consistent finding in the literature, which is that not only do, for example, black patients have more pain than white patients, but that difference, that gap in pain persists even when you control for.

[00:27:22] Severity of disease in this case, in the knee. So what you do is you take two patients, one black, one white, and you get two patients whose x-rays, their knee x-rays were graded in exactly the same Kelgrin Lawrence score. Both of 'em have the same score. That's like the system that radiologists use to grade arthritis.

[00:27:41] You look at all the compartments of the knee, you assign each a score, you sum it up, it goes from zero to four or five. Anyway. I'm not a radiologist, uh, far from it. So when you do that, you still find that black patients report more pain and that's interesting. And so I think David's [00:28:00] interpretation, and I think the interpretation that most people have is that there are psychosomatic and psychiatric stressors.

[00:28:06] There are differences in medical care. There's lots of stuff that's going on outside the knee that could be contributing to this and causing this gap because we've held the knee constant. And so after the presentation, I went up to David and I was like, I bet it's in the knee. He's like, no, it's not in the knee because we've controlled for that.

[00:28:22] And I was like, no, I, I bet it's still in the knee. So that was the, the friendly bet that led to, uh, what were the stakes? I neither David nor I is, an actual betting man, but you know, when your academic a gentleman's bet, then when your academic credibility is on the line, Andrew, it's, you know, it's, it's worth, it's priceless.

[00:28:42] So we were able to get data from this NIH study that David had been using, but not only did we get the data on pain and race and, and Kelgrin Lawrence grade that data had gotten, we also got the x-rays. And so what we did is we, we trained an algorithm to look at the x-ray. [00:29:00] And predict whether a patient would say, this knee hurts or this knee doesn't hurt.

[00:29:05] And it's worth drawing a distinction between that approach and what I think most papers in the literature do, which is like, well, how do you train an algorithm to read a knee x-ray? You train it to replicate what a doctor would say. Mm-hmm. About that knee x-ray. So you'd input the x-ray and you'd ask it to output the Kegrin Lawrence grade, but.

[00:29:26] If you're me and your hypothesis is, well, I think radiologists are just missing stuff in the knee that's causing pain. That's exactly what you don't wanna do. So you don't want the algorithm to replicate the limitations of human knowledge and the errors built into it. You want it to actually find new signal.

[00:29:42] So we trained it to listen to the patient instead of replicating what the doctor would've said. And so that gave us an alternative measure, an alternative fact, if you will, about the knee. So radiologists have their score, they have a Kelgrin Lawrence grade for the knee. We have our algorithmic pain [00:30:00] score, which is capturing the predicted amount of pain from the pixels that we see in the knee.

[00:30:06] And what we found is that if we then go back to that pain gap between black and white patients, there's some gap. So black patients are in much more pain. If we control for the radiologist's interpretation of that pain, we account for I think 9% of that gap. But if we control for the algorithm's interpretation of that x-ray, We account for about half of that gap.

[00:30:27] So the algorithm is doing a much, much better job of explaining pain overall, but it's doing a particularly good job of explaining the particular pain that radiologists miss and that black patients report, but that can be traced back to some pixels in the knee. Right. Um, yeah, I, so I, you know, putting on my stats hat like models are meant to explain variation, that's fundamentally like what a model does.

[00:30:51] And I think what you're saying here is that the deep learning model actually does a better job at explaining subjective pain variation than the human model [00:31:00] does. Is that sort of a fair con condensation? That's exactly right. If you look at the R squared of pain, so like of all of the variation in pain, how much does the model explain versus the radiologist's interpretation?

[00:31:13] The model, I think was about 50% more overall. Wow. Wow. Now again, it was still low. Yeah. And there's still a lot going on in pain that cannot be traced to the pixels, but when you're looking at the pixels, you wanna get the best explanation of pain as possible. Yeah. And again, this I think, really highlights what I like so much about your work is that all technologies have this duality aspect to it.

[00:31:37] They're sort of, uh, net neutral. They can be harmful or good depending on the context they're used, the data they're trained on. And I think you sort of have shown the dynamic range of possibilities for good versus ill with how we think about medical machine learning. Thanks for pointing that out. I could not agree more with that statement.

[00:31:53] I think they're just tools and the tool is neither good nor bad. It's just about [00:32:00] how we build the tool and how we choose to apply it. And that can do terrible things or it can do really amazing things and it's really just all up to us. Mm-hmm. Yep. Okay. So Raj, are you think we're ready to transition to the lightning ground?

[00:32:15] Should we ask Ziad about the data sharing and AI innovation? Oh, right, yes. Let's do that. So we've done a good job at exploring a lot of these early topics, but we do wanna give you a chance to talk about Nightingale, some things like that. So I'll ask that one. So in addition to your, your primary research, um, you've also taken some entrepreneurial roles around building companies and platforms that can help others do, uh, sort of what you've done in the past by getting access to these rich data sets to answer important questions.

[00:32:43] And I think importantly, the question for these data sets is not, not necessarily pre-specified, it's kind of more like a, a playground or a sandbox or just sort of a place to do research. So could you walk us through Nightingale and sort of, um, how that started and sort of what you hope to achieve with that?

[00:32:59] Yeah, [00:33:00] I think to, to go back to an earlier part of our conversation, Getting the data that you need to do your research is a huge, huge preoccupation of any researcher in this area. And I think the problem is that the data are essentially locked up inside of the health systems that produce the data. And it can be really perverse.

[00:33:23] Like I, I'm sure for example, that if you are appointed at the medical school at Harvard, but because the medical school is a different administrative entity to the Brigham Women's Hospital or the Beth Israel Deaconess Hospital, literally across the street, you do not necessarily have a right to access that data, except if you go through some other channel.

[00:33:44] It's not, we don't don't like people to know this until they agree to be a postdoc with us. So you're giving away all the, all the bad secrets. I'm sure you can, we'll edit this out later. Yeah. So, you know, the, so these things are really perver and if you wanna hire an analyst to, to do anything with the data, you have to hire them through the [00:34:00] hospital, which means you have to have a grant through the ho.

[00:34:01] So, It, it's Byzantine and it's very frustrating. And, and I think it's really holding back this space because it is multidisciplinary work. You need computer scientists, um, hospitals can't hire all the computer scientists that they would need to do this work, and they're not gonna win the war for talent, um, against Google or Facebook or, or even just computer science departments of different universities.

[00:34:28] It's just, you know, it is just a fact. So doing this work is just, you know, there are so many frictions and it's why I have a few papers that have, like, one paper took me seven years to write. Um, there's another one where we're at, I think year eight now. And like, you know, a lot of it is because I'm inefficient and an academic and just, I procrastinate a lot and I'm not that great at getting things done.

[00:34:51] But it's not all the blame. Some of the blame is on data use agreements and IRB amendments and de-identification problems and [00:35:00] things like that. So a few years ago, uh, thanks to initially support from Eric Schmitz Foundation, uh, Schmidt Futures, but then with follow on funding from the Moore Foundation, Ken Griffin, and a bunch of other donors, we set up a nonprofit called Nightingale Open Science.

[00:35:15] And what Nightingale does is we take that philanthropic money and we deploy it to build interesting data sets with partners inside the health system. So, uh, working with, uh, researchers at Providence, um, St. Joseph's a big health system all up and down the West coast. We bought two slide scanners and hired some pathology technicians, and we went to all of their breast cancer biopsy slides.

[00:35:40] These things were literally collecting dust on the shelf in a basement, and we just started digitizing those slides, linking them to electronic health record data, to cancer registry data to social security data on mortality. And if you go to nightingale science.org, you'll be able to fill out a very lightweight [00:36:00] data use agreement and within a day or two, get access to those de-identified data on an AWS platform with a Jupiter notebook where you can just start doing your research.

[00:36:11] And we're exploring models and as you and I have talked about before, Andrew, of like how to get those data also into the ecosystems of researchers at universities who already have clusters who don't wanna do things on AWS. But we've actually had a lot of traction from researchers around the world who don't have their own compute clusters, who are delighted to, to use the AWS, platform that we have things on.

[00:36:31] So, so we build up those data sets around interesting medical problems, in this case, cancer metastasis. We de-identify the data and then we make those data available in a secure and we think ethical way on this platform where researchers can use it to do nonprofit research. Um, so that's Nightingale.

[00:36:49] This is something that I think we've talked about before, but other than a charm offensive, which I know coming from Zaid is a considerable force, how do you get the hospitals to actually, what is in it for [00:37:00] them? Like why would they want to contribute to data from this? I know, you know, especially Boston area hospitals are notoriously stingy with their data, even with Harvard researchers.

[00:37:09] So how, what's the key to unlock that data? I think the key is that there are a lot of hospitals and so even in our case, if 200 say no, five say Yes. And so we were very lucky to find fantastic partners at hospital systems from big non-academic systems to small county hospital systems. And there were people who just, they had a problem that we could help solve, whether it was for research or something operational, and they were able to push that forward.

[00:37:40] And so we, we were very, very lucky to find the partners that we found. But you know, your question is such a good one because getting a hit rate of five out of 200 or so that we talked to, it was a frustrating process and it's definitely. Puts a limit on the scale. And so it was actually in the course of getting rejected, [00:38:00] uh, hundreds of times that actually, like, there was another idea that occurred to us at the time and us as, uh, my co-author, sandal Mohan, who's an economist and I and a few others, is that the same problem?

[00:38:12] That was, I think a huge bottleneck for research was also a bottleneck for people who wanted to build products that would have a much more direct way of getting into the clinic than a paper that we as researchers publish. And so there's a for-profit that I also co-founded called Dandelion. And Dandelion has agreements with a handful of large non-academic health systems where we can curate and create data sets and build those data sets to serve the needs of startups.

[00:38:42] Life sciences company, medical device companies that wanna build AI products with a view to deploying them into the clinic for patient benefit. So do you act as like a, a broker in between the, the industry partner and the hospital system and, yeah, that's exactly right. So I think the current way, if you're a, an AI [00:39:00] startup and you wanna get access to data to build a product, you have to build a bilateral relationship with a hospital and go through a very onerous contracting process.

[00:39:09] And, and so I actually know of a handful of startups that ran out of their Series A funding while they were waiting for these kinds of agreements to get finalized through the hospital. So it's really definitely stifling a lot of innovation that I think could benefit patients. And so the idea of dandelion is to reduce those frictions to make it easy for people who wanna develop products to access data again, in a secure and ethical way.

[00:39:32] And I think the ethical. The ethical way I think is an important one. Cause I think you, you have to think hard about, well, we're selling data basically, and I think the way I've been thinking about it at least is there's a huge opportunity cost to leaving those data not used. There are so many insights and products that could directly benefit patients that are not getting developed today because it's so hard to access those data.

[00:39:59] So [00:40:00] even though there are clear threats to privacy that we need to take into account and security, there's also this other risk that I think we spend a lot less time thinking about. But that is at least as large as the privacy risk, which is what, what about all of the things that could diagnose cancer early that we don't have today?

[00:40:19] Because nobody can develop the product. What about all of the ways we could do better triage in the emergency department and get care to people who need it faster and reduce the cost of care? Like what about all of those products? Where are they? Well, I think they're not being developed because it's so hard to get the data to develop the products to begin with, and that that's a problem too.

[00:40:39] Yeah, I, I agree. And the healthcare system is just very bad at trade offs in the sense of weighing explicit costs versus implicit costs. And I think like, well, they're very good at minimizing the, the risk of explicit harm. Um, but there's no accounting for the implicit harm done by lack of progress and lack of innovation.

[00:40:57] And I think that, you know, hanging out with economists is a [00:41:00] great way to get into that mindset. And if there's a way that we can reduce that friction, I think that it'll hopefully tip that balance in a, in sort of a more favorable direction. Yeah, very well put. Yeah. On Dandelion and Nightingale, how do you think of what is sort of a for-profit activity versus a nonprofit activity?

[00:41:18] Because in a way they're both data platforms accelerating AI and medicine, so how do you partition work or projects for one versus the other? Yeah, F fantastic question. I think the way I think about it is just to work back from the outputs. And so I think there's a lot of research that needs to be done to produce ideas and papers.

[00:41:39] Like there's a lot we don't know about what cancer is, which cancer is spread, like can we predict which cancer like, so there's just some basic science questions that we need answers to before we can develop. Products. On the other hand, there's some very practical things that we can start doing today, even without fully [00:42:00] understanding the nature of cancer and metastasis to improve the care of patients with cancer.

[00:42:05] And so there's also a bunch of things we can do with data that have an outlet to things that benefit patients directly on non-academic timescales and with non-academic skill sets. And I think one of the things that I've really come to appreciate about the private sector and basically my new non-academic friends and acquaintances is, boy do they get shit done.

[00:42:28] They don't have projects like I have that have gone on for eight years. If it goes on for eight days, it's like, what's going on? What's taking so long? So there's an impatience and a raw competence that I've been trying to learn from, from that world. So se separation of timescales is, uh, is maybe one, one principle to think about it.

[00:42:48] And then I, I have to ask what is behind the name Dandelion. I'm thinking about being on the playground and blowing the flowers and the little petals going everywhere. I think that's, that's definitely one of the images we were going, it, [00:43:00] it's, you know, there's a dissemination kind of thing.

[00:43:03] It's a positive thing. Dandelions are pretty. Elliot, our CEO and, and co-founder, he'll say all that stuff, but then he'll admit that it's because his dog likes to eat dandelions. Maybe I'm misremembering my horticulture, but that's actually the way that they reproduce, right? That those things flow and then they Exactly, they plant it and create new dandelions.

[00:43:22] So absolutely a very apt metaphor. Okay, so now I think we actually are ready for the lightning round. So, um, buckle up, uh, and, uh, we'll let you know, uh, if you win at the end of the round. So we're speaking today on March 14th, 2023, which is the day that GPT4 was announced and is now a thing. So we're, well, it's also pie day, I thought, I thought this was gonna be a pie day joke.

[00:43:54] Yeah. So, uh, that's actually true also. So maybe we'll have to, we'll have to work into some, [00:44:00] uh, some pie questions at the end of the round. So, I guess the first question is, given everything that's happened over the last six months and given your work, thinking about the duality of machine learning in medicine, I think it's fair to say that we're on the cusp of large language models starting to affect clinical practice.

[00:44:19] So I would like your wager on if LLMs large language models will be net positive for medicine over the next five years. Be net, and this is the lightning ground. So I'm gonna ask you to be brief in your responses. Definitely net positive. I would bet that it'll be more positive on the backend in terms of summarizing and abstracting information than it will ever be on the front end, given everything we know about the problems with the front end.

[00:44:46] I love that I get to ask this question. I think I know maybe one of the answers to this, and it might involve a beach, so I'm priming you a little bit here. If you weren't in medicine, what job would you be [00:45:00] doing? Um, you can dream, you can dream big. Yeah. I think that there's a couple of things that I love.

[00:45:09] The real answer is that I get to keep being a doctor. I also get to do research. I also get to do this private sector stuff. So I feel like I, I'm getting a lot of the best of many worlds. The other things that I would love to spend time doing are making espresso. Uh, I spend a lot of time doing that in my off time anyway.

[00:45:29] I recently learned how to surf, and so being a, a surf bum would also be really fun. And I think those are my two big hobbies for the moment. Awesome. Uh, the, the barista beachbum. I, I like that. Uh, counterfactual. Um, uh, so our, our listeners can't see this, but you're sitting in front of an impressive display of books.

[00:45:50] So this question is, what is your favorite book? And maybe keeping with the Beachbum theme, if you were, what's your Desert Island book? If you were stuck on a desert island, what book would you take with you? [00:46:00]

What a great question. I think, you know, how can one pick one's favorite of anything? I think the book that I've probably reread the most and so would be my Desert Island book.

[00:46:15] It's a book called The Rise and Fall of Modern Medicine, and it's an amazing history of medicine through the lens of the amazing discoveries that have been made by hobbyists, people doing frankly dangerous and irresponsible stuff. And it is just full of amazing stories. And who wrote that book? Uh, James Le Fanu.

[00:46:40] Okay. Is a, I think a general practitioner in London, if I'm not mistaken, and a fantastic historian. Will AI and medicine be driven more by computer scientists or by clinicians? Clinicians with. Deep investments and expertise in computer [00:47:00] science, statistics, economics, and behavioral science. So we're getting back to this multiple skill sets and domains of knowledge existing within the same mind as opposed to collaboration across separate silos.

[00:47:12] Yeah, and and I should caveat that unfortunately by saying that those people are not being churned out at a very rapid pace by our system of medical education, unfortunately. Given your status as gentleman scholar and well-read philosophy student, if you could live in one other historical period, what would it be?

[00:47:33] I think the smart money is always on the future, and so I'm gonna say maybe 150 years in the future. So it's like stuff is better, technology is better, but it's not so different that I would just be totally lost. Uh, is that conditioned on there still being a future in 100 years? So existential risk being roughly constant or diminishing over that period, you're taking that as for granted?

[00:47:58] Yeah. I mean, who, who [00:48:00] knows? Maybe, maybe we'll be on Mars, but yeah, I, I think I would, given my probability estimate of that, I'd take it unconditionally. I got it. All right. Last lightning round question. Do you think things created by AI can be considered art? Absolutely. Yeah. I, I think, I think all of these questions about is it sentient?

[00:48:23] Does it pass the Turing test? It's so beside the point to focus on the AI rather than to focus on. The reaction that we have to the AI and what we can learn from the AI. So I think it can absolutely be art. I think it can absolutely teach us things, and I think it can absolutely do all those things without being conscious or, or sentient or whatever.

[00:48:45] It's really, I mean, not to sound narcissistic, but it's all about us. It's all about the humans, how we're using ai, how we are getting better and making society better with ai. Beauty is in the eye of the beholder. Um, so congratulations. [00:49:00] I'm happy to report that you passed the lightning round with flying colors.

[00:49:03] I know that was a relief, a real, uh, a nailbiter there. What would've happened to me if I hadn't passed? Andrew? Uh, we are not allowed to disclose that. Oh my gosh. Okay. If you read carefully in the, if you read carefully in the material that we sent over that you had to sign, uh, it was included in that.

[00:49:19] Okay, so we'd like to just have some big picture, I think topics that we'd like to discuss with you before we let you go. And again, this goes back to all of the work that we've discussed so far, and I think it's less about AI and more about us, like you said. Do you think that AI over the long term is going to reinforce or reduce healthcare disparities?

[00:49:42] And that can again be a, a societal question, are we gonna get our stuff together and figure this out? Or if not, which way will AI tip the scales in that scenario? I think from the point of view of today, I think we often [00:50:00] overvalue or overestimate big historical, secular trends and undervalue I. Personal agency and the fact that what we do matters.

[00:50:12] And so I think that question depends entirely on we're, we're at a very critical moment now, and we don't have a lot of regulation right now. We don't have a lot of practices that inform how we adopt or purchase or scale up AI in lots of different sectors. So a lot of things are gonna get decided now, and a lot depends on the current generation of researchers, clinicians, health system executives, regulators at the FDA and other agencies.

[00:50:39] Like what we do now is really gonna matter for that long-term view. And I think it's critical that we get it right now. And so all that said, If you look at the track record of most technologies, they tend to reduce disparities over time. And so I think, you know, my base rate estimate for [00:51:00] whether technology increases or decreases disparities is that it decreases disparities as all of the innovations in our healthcare system have done over time.

[00:51:10] I think that's a very reasonable stance. I guess I can just be pessimistic that structures that have been in place for so long are gonna get solved before. Cause I, I do feel like there, there has to be sort of an order of operations there, there's a sequential dependency that if those structural problems don't get solved before AI becomes entrenched, then you know, as your work and others have shown, then it can just operationalize all of that.

[00:51:34] So, but no question. I agree that sort of the longer term view of that, um, as trending towards, towards balance, I think there's, there's another thing that I've started. Becoming conscious of and trying to work towards, which is using the technology to actually fight against those structural problems. And I think it's hard to see sometimes because if you look at something like wearables Yes.

[00:51:59] Who, [00:52:00] who has wearables right now? A bunch of rich people. And so right now, if you look at the impact of wearables, it's like, yeah, it's not, it doesn't look great. But if you think, if you step back and you think about wearables as a tool to fight inequality, well it would be great if we could cut out all of the problems with accessing the medical system.

[00:52:22] Through the use of something that is in everyone's pocket, you know, like on your cell phone or on your wrist, and cost keeps going down. So I think that's not a crazy thing to think. So it would be fantastic if you could get access to the quality of medical advice that rich people have access to on your phone tailored to you and your heart rate and your blood pressure and things like that.

[00:52:43] So I think these technologies, to pick up a, a recurring theme of this conversation, they can be used, their dual purpose. They can be used to reinforce disparities or they can be used to fight against them. And I think that given how entrenched the current structures are, using [00:53:00] technologies to disrupt some of those things and get access to people who don't have access is, is a very reasonable goal.

[00:53:07] I think another one of the big recurring themes of this conversation to me is that where really high impact, high quality research tends to happen, or at least one of areas, uh, one of the sort of set of circumstances where it tends to happen is when multiple skills exist within the same person, right?

[00:53:24] So someone who's a computer scientist, who has a deep appreciation of medicine or economics and can speak both languages and can be creative about how they see connections across these fields. So we have a lot of early career clinicians who are listening to this podcast, med students, residents, and they are engaged in medical school, right?

[00:53:43] They're in medical school or residency or training right now. What do you think they should know about AI to help them prepare for a career in medicine? Uh, what could they maybe self-learn? Who can they talk to? What could they do at that point in their training journey to really be able to, uh, to [00:54:00] speak to, to both skills?

[00:54:02] I think, you know, that. There are a few things that worked well for me that I'll just restate. I don't know if this is optimal. Uh, it's almost certainly not, but it worked for me. So I think the first thing I would probably suggest is that AI is a subset of statistics. It's an applied version of statistics with real data sets and, but, There is no substitute for learning the basic statistical stuff.

[00:54:32] And I think as a starting point, that is an amazing place to start to get a handle on thinking about how AI works, where to apply it, where it can go wrong, all of all of those really important things. We're blessed to live in an age where you can just go online and take a class by some of the best statistics teachers in the world.

[00:54:54] You just like do it instead of watching a bad TV show. It's about the same amount of time. [00:55:00] So I think that's an amazing resource to take advantage of. Now, Khan Academy is great, like there are so many great resources for learning things online, but starting with statistics or in my case, for me, the toolkit that always made the most sense or resonated with me the most was the economics, the microeconomic.

[00:55:19] Toolkit for dealing with data that's produced by humans and is messy and error prone and and driven by incentives and things like that. That seems a lot like medicine to me. And that way of thinking about data and analysis was always very compelling for me. So there's a book called Mostly Harmless Econometrics that I really liked.

[00:55:39] Uh, there's another version of that same kind of curriculum called Mastering Metrics, which is in some ways even better. But I, I think those are two fantastic books. Um, so would you be in favor of the U S M L E adding a block on econometrics to, uh, to licensing? You know, I think the question with these things is like, you can't just add, you have to cut.

[00:55:58] So what do you Cut? Cut, right. [00:56:00] I think there's a lot to cut, but yeah, I think absolutely. I think being a good doctor over the next 20, 30 years, it is just hard to imagine you can do that without knowing how to evaluate claims about vaccines and masks and a lot of other things that have come up, you know, a lot over the past couple days, like medical knowledge is constantly evolving.

[00:56:19] It'll evolve even faster as AI tools come into play. Being able to critically evaluate those claims is just incredibly important. So evidence, appraisal and statistics as sort of the foundation for the house? Uh, yeah, absolutely. Into it. Yeah. Yeah. And just thinking with data, I. But it's just like, it, it's, it's related to your earlier problem formulation point where it's just like thinking about how to take an abstract question and then think about what is the data frame that would answer this question.

[00:56:49] That is just like an incredibly valuable skill. I'll also say that one, to go back to ChatGPT, I don't know if you all have been doing this too, but I've been using it to write code [00:57:00] a lot and it's actually like dramatically increased the amount of code that I write. So I, you know, I have some research assistants, some PhD students, and you know, because I'm, I don't code that much anymore, it's like often easier for me to just ask someone to do something for me, but it's like a little unsatisfying, like I don't, like I know what I need done.

[00:57:22] I just like remembering which libraries to load and how to specify the stupid reshape command and all these, it is just like so painful. But now I can just type it in English. And then it generates a code snippet and then I can paste that in and like 80% of the time it works. It's incredible. So I actually think that for a clinician who has access to some data and knows some basic, like knows what are the statistical operations that she needs to do, this is an amazing tool.

[00:57:51] It's an incredible tool that I think for me, it's gonna honestly reshape a lot of how I do work over the next few years. Yeah. Andre Carpathy [00:58:00] noted AI researcher and now at Open AI had a quote that was like the number one fastest growing programming language is English. In the wake of ChatGPT love. So I think a lot of people are having a similar experience to you.

[00:58:12] Yeah. Speaking of things that people waste time on, uh, I spent too much time on Twitter arguing whether or not AI is a proper subset of statistics. Oh God, I shouldn't have wandered until this. I will not take the bait this time. I've learned my lesson. Change topic. Change topic immediately. Yeah. Well, maybe I like, here's, uh, something that I, I heard recently that I like, which is that it's as engineering is to physics, data science is too statistics.

[00:58:41] How do you feel about that? Um, analogy? Yeah. I, I, I find that unobjectionable and, uh, I think that I would just, there are parts of AI that have nothing to do with statistics, is all I would say. I, I totally agree with that too. Have a penance argument at best. So one of the last questions for me, which I think that [00:59:00] you're gonna have probably a few things to choose from for this one, but what is your most heterodox opinion?

[00:59:06] I. Um, with respect to accepted wisdom in, in the field. I mean, you know, uh, how, how much time do we have? Like, I can, we can talk about essential oils or, you know, uh, I think, you know, maybe, I think it's, it's probably related to medicine as a field, and I think my view today is that medicine is not its own field of study.

[00:59:40] And I think that that's hard to argue with that on, on some literal level. Like you can't get a PhD in medicine. And I think that it's because the data that make up medicine, like the things that get generated about patients in the electronic health record, in radiology PAC systems, [01:00:00] they're very hard to handle with traditional statistics.

[01:00:04] And so I think the field is stuck between a bunch of rules of thumb and kind of like an arts and crafts sort of approach to, to doing things where it's like maybe more charitably. Artisanal. Yes. It's it. Thank you. Yes. A small batch. Uh, locally. Locally crafted knowledge. Yeah, exactly. So, so I think that my view is that not only is there a field that is like ML plus medicine that is medicine, I think that is the thing that medicine will be as a science now.

[01:00:44] It'll incorporate a bunch of other stuff. Like you need a bunch of causal inference stuff that's not. Really part of ml. You need behavioral science, you need economic, like there, there's a, like a, but at its core, I think just given the data that we have, it's [01:01:00] high dimensional. It's like it's everywhere.

[01:01:01] It's like super complicated and given the types of problems that medicine solves, which are, you know, really like often around prediction and description and accounting for variants. I think this is medicine. I think our common friend, Zach Johanne, who's the chair of the department at Harvard, has this phrase that I think we've all heard a thousand times.

[01:01:21] That just becomes like a chant at some point. Or a mantra, which is medicine is an information processing discipline. And I think that at its core, I think that aligns perfectly with what you're saying. That one, what medicine should be doing is processing information. And it's kind of like this discipline of doing that and how well are we doing that is, is actually the the key question.

[01:01:41] Yeah, I love that. Yeah. All right, so we have one last question and it's a pair of questions inviting you to imagine two different scenarios here. The first is, what is your best case scenario for AI and healthcare? And then of course opposing that. What is your worst case or what are you most worried about for AI and [01:02:00] healthcare?

[01:02:01] Well, I'll start maybe with the worst case cuz it's, it's easier to imagine because it's closer, unfortunately, to today. I think the worst case is it's not just that we get biased tools and, and deeply flawed tools of the kinds that we've talked about, that I've studied, that I'm really worried about. It's that all those tools are used for is this local optimization of a system.

[01:02:32] That sucks and that isn't proactive. That's very oriented towards billing and coding. There's a very unappealing path where we just get a hyper optimized version of our current shitty system. Can we count the beans faster? Yeah. Yeah. And, and I think that, There's a certain lack of ambition in how people are applying AI [01:03:00] today.

[01:03:02] That's really, I think, constrained by the imagination that people have and the kinds of problems that are front and center for decision makers today, which are unfortunately, you know, and, and you can't blame health system executives for being really worried about their profit margins. You know, given that inflation is really hard on hospitals that have fixed contracts, um, and, and rising costs.

[01:03:26] It's just like, I, I get it, but it's not good. So I think that the optimistic version of that is that, man, there are a lot of great problems to solve with ai, and I think there are two types of problems, and I think I've worked on the former, but I'm starting to work on the latter. The former is just, there are just these first order misallocations of resources that are tragic.

[01:03:49] Like I am working on a project on the theme of getting data. Uh, this is with a, in collaboration with a regional health system in Sweden, because in Sweden you can get all of [01:04:00] the electrocardiograms that were ever done in that region, uh, and you can link them to death certificates. And this is with a, a cardiologist in Sweden called Marco Lingon and my other co-author, Sendal Mohan.

[01:04:10] And we're studying sudden cardiac death. So sudden cardiac death is exactly what it sounds like. People just drop dead. It's really hard to know why even expost after they've dropped dead, because it's really hard to study people after they've dropped dead. And what makes it particularly tragic is that for at least some fraction of those deaths, we have the cure.

[01:04:29] We have defibrillators that we could put in people to terminate the fatal arrhythmias that cause these sudden cardiac deaths, but we just don't know who to put those defibrillators in. So one part of that project is just asking the very practical question of can we identify people who would benefit from defibrillators before they die rather than after they die?

[01:04:50] Super important, like very practical question around, we have this resource, we're currently putting it into the wrong people. 75% of defibrillators [01:05:00] never terminate a fatal arrhythmia. It's crazy. Um, and on the other hand, lots of people are dying of arrhythmias that could be prevented by these defibrillators.

[01:05:09] So crazy Misallocation first order problem, hundreds of thousands of deaths in the US every year. So that's one kind of problem that I hope more people will work on. And to go back to Nightingale, the data sets that we put on Nightingale are. Constructed around exactly those kinds of problems that are interesting and important and, and critical.

[01:05:27] So I hope more people will go look at those data sets and start solving those problems. I think there's a second type of problem where AI can really help, which is not just the practical parts of like allocating resources in the healthcare system, but actually making scientific discoveries in health.

[01:05:44] And so we're applying our machine learning algorithms to ECG wave forms. So you know, maybe the algorithm like sees a little squiggle, but. Unlike, you know, when you're doing this stuff in ImageNet, it's like, okay, yes, this is a Frisbee. [01:06:00] That's why it's say okay. But in ECGs we actually know a lot about how the heart produces the ecg.

[01:06:07] We know what part of the heart leads to what part of the wave form. We have simulation models of the heart that we can get to produce waveform. So I think tying that pipeline together of biological understandings of the heart and how the heart generates data to things that we're learning about the data via machine learning from patient outcomes, I think that's super promising.

[01:06:30] And I think that's the kind of stuff that I'm really excited about working on, um, over the next few years. So discovery of novel pathophysiology, um, in the heart from ECG data. Yeah, exactly. Yeah.

Perfectly put. Yeah. With, with a view. And then once you have that, Pipeline connected up. You can use it to discover drugs, you can tie it to animal models.

[01:06:50] You can, like, there's a lot of things you can do once you get the data. Talking to the biology. I have one wild card question to ask you at the end, just cuz it ties in nicely with [01:07:00] what you said. So, uh, you wrote, uh, perspective in any JM a few years back called Lost in Thought. And one of the points that you make in this perspective is that plausibly medicine already exceeds the capacity of the human mind and that we need AI as kind of like an integration tool to synthesize all this data.

[01:07:17] So can you sort of square that with what you just said? Is new pathophysiology understanding the goal, a goal relative to sort of the limitations of the human mind? I think, you know, you first have to give the human mind a lot of credit. Like we've figured out a lot of stuff about a very complicated system, but we really have no business figuring out.

[01:07:42] It's very striking, like when you take your car in. To the mechanic and you're like, I don't know. It's making the sound and the mechanic's like, yeah, I, I don't know. Like we built the car. It's like we, we can't figure out the car. I mean, the body is pretty complicated. [01:08:00] And I think that was the other thing that bothered me about med school is like, we have these drawings of like the cell and the nucleus and it's like, oh yeah, here comes the dna, here's the ribosome.

[01:08:09] This is happening on the, like millisecond scale. Things are bumping into each other all the time. We have no idea what's going on at, at a, at a macro level, and yet we still manage to figure out a bunch of stuff. It's amazing that we figured anything out that said we could use some help because, you know, If you think about something like a decision rule, you know, we use these decision rules in the hospital.

[01:08:32] It's like, okay, if your EKG does this and you're over this many years old and you're on this medication, then your score is this. But I don't think anyone actually believes that heart attack is a phenomenon that can be adequately described with like five integer weighted variables. It's just, you know, there are useful constructions that help, but things are complicated.

[01:08:55] We, we have this paper that looks at actually physicians in the emergency department [01:09:00] diagnosing heart attack, and that finds that. If you constrain a machine learning algorithm to a very small number of variables, it actually looks a lot like the doctor. So if we put ourselves on even ground and are restricted to a small number of variables, we actually are doing pretty well.

[01:09:18] It's not optimal, but it's like pretty close. The problem with us is that we can't go up to like 10,000 variables. We're low dimensional creatures. Yeah. And medicine is a high dimensional field, and so I think being able to. Capture that high dimensionality while also getting it in dialogue with our ability to think about models and counterfactuals and reason through things.

[01:09:43] Like that's, I think, the thing that's really gonna drive science forward. I think over the next few decades, there's, there's a wonderful paper that you probably know already, but it's, um, by my friend and colleague, Gregina Barzilay at m MIT, where she and colleagues [01:10:00] actually train an algorithm to discover new antibiotics in classes that are Yeah.

[01:10:04] Hallison. That, that were completely unknown to people whose living is to discover antibiotics. And it all just came out of this amazing. Exercise that tied data discovery to new drug libraries, to an experimental model where you can see whether things are killing e coli. So I think that paper for me was really like opening the door to this new way of doing science that integrates human reasoning and data collection and hypothesis generation with data-driven approaches from ml.

[01:10:36] Yeah, I think that's a great note, an optimistic one to end on as well. So I want to just thank you so much for being on AI Grand Rounds. Uh, this was a really, really fascinating conversation. This was so much fun. It's always great to hang out with both of you.