AI Grand Round Podcast #5 04.26.23 No Doctor Needed? Dr. Michael Abramoff on the Potential of Autonomous AI

[00:00:00] Another thing, which also happened in 2010, which was that I was doing all this research and my colleagues started to notice my colleague ophthalmologist, very esteemed colleagues, and certainly there was an editorial in the most widely read ophthalmology journal in the world called Ophthalmology Times.

[00:00:19] On the front page by the chair of ophthalmology at Johns Hopkins, the Retinator Revenge of the Machines, about my research, calling him out by name. Essentially saying that I'm the terminator of retina. I'm the retina specialist, right? It's. That’s that, that hurt and, um, and it's going to cost, uh, jobs and it's going to lower the quality of care.

[00:00:41] It was partially tongue-in-cheek and partially very serious. That was Dr. Michael Abramoff recalling some of the challenges that he's faced in pioneering AI for ophthalmology, where at one point he was even called the Retinator, the Terminator of the eye. So welcome to another episode of AI Grand Rounds. [00:01:00] I'm Raj Manrai and I'm here with my co-host Andy Beam. [00:01:03] So Andy, I was really struck by Michael's amazing amount of persistence and passion as we heard in that clip. He's had challenges, but he just keeps moving forward. He's been a real pioneer in medical AI for a few decades. He's a professor at the University of Iowa, and also the founder and executive chairman of Digital Diagnostics, an autonomous medical AI company.

[00:01:23] We got to explore both the genesis and the implications of his groundbreaking work, which led to the first FDA authorized autonomous device, so no doctor required. This is the IDx-DR device, and it detects more than mild diabetic retinopathy from digital images of the eye. I also have to say that at moments, the entropy, or if you prefer, the temperature setting, was a little higher than I expected, and to our listeners, if you listen long enough, you'll hear about asteroid mining, Lana Delray and Ethical AI, which is a sentence that I never thought I would say. Overall, this was an illuminating and really fun conversation.

[00:01:59] I have to [00:02:00] agree, Raj, I did not see asteroid mining coming either, and I think you hit the nail on the head when you said that you were impressed by his persistence. In addition to his deep technical innovation in this space, I think what struck me the most was his willingness to go deep on the non-technical side of things to make sure that he was able to improve patient outcomes, and specifically his work on defining new medical reimbursement models for AI I found to be very fascinating and his willingness to take seriously the bioethical considerations for medical AI.

[00:02:30] So again, I totally agree. It was a fascinating conversation. I certainly learned a lot from it. And with that, we're very excited to bring you Dr. Michael Abramoff on NEJM AI Grand Rounds. The NEJM AI Grand Rounds podcast is sponsored by Microsoft and Viz.ai. We thank them for their support.

[00:02:54] Well welcome to AI Grand Rounds, Dr. Michael Abramoff. We're excited to have you on today. Thanks so much for having [00:03:00] me. Very excited. So Michael, this is a question we like to always start with. Could you please tell us about the training procedure for your own neural network? How did you get interested in AI?

[00:03:10] What data and experiences led you to where you are today? Yeah, absolutely. And uh, you know, my neural network is protected by, uh, a set of white hair. So you, listeners cannot see me, but I'm the only one here with white hair that shows my age and also therefore how long I've been involved in this. I started training as about a medical student as well as a computer engineer.

[00:03:33] Long time ago in the eighties. Ultimately became a neuroscientist in Japan working on neural network simulations of the brain. So what we were trying to do at the time was mimic single neurons in a computer cloud and then try to. See how that related to cerebellar neurons where we had brain slices and try to mimic them and ultimately try to build larger neural networks, maybe three neurons, right?

[00:03:58] Compared that to the millions of [00:04:00] neurons we use now. And there was this larger group there, we were looking on various ways of, well, how does a larger neural nets work processing information, you know, about information theory of the brain, so to say. So I was very excited about that. Ultimately decided to go into the software industry for, worked for a long time in, in, in software.

[00:04:18] Went back to medicine, finished my residency in ophthalmology. I've always been interested in, in section, like you say, of, of the audience here of AI, as we used to call it, machine learning then, or neural networks and medicine and how you can apply it, right, not only as a theory of the brain, but also can it actually be useful for what we do in clinic.

[00:04:38] And so I kept back and forth and many people told me, wow, this combination of computer engineering and machine learning and medicine is so useful, so exciting. And there was literally nothing. No, no one was interested. There was no funding, nothing. So I kept moving between doing my residency and ultimately ending up doing a PhD in machine learning.

[00:04:58] Finished that, uh, did a [00:05:00] fellowship in vitreoretinal surgery, really focused more on the surgical aspects for a few years. Got an opportunity, this was in the Netherlands issue here from my accent in Amsterdam. Came to, uh, to Iowa in 2003, where I was able to get NIH, National Eye Institute excited about the types of research I was doing where I said, well, can we use machine learning to mimic what a brain of a clinician like me does?

[00:05:23] And so that went was a long path ultimately with me starting to be interested in either why, why, why, why do I think it's interesting? And I started to see promise in healthcare, like the high cost, low accessibility, lack of health, equity access, and in my view, The solution in other fields like agriculture.

[00:05:43] I'm here in Iowa and there's, you know, John Deere, uh, tractors all around me, right? There's corn fields and silos, and so I said, well, clearly automation, fertilizer, all these things, mechanization have helped tremendously in making food, maybe [00:06:00] 50% or 80% of people's expense a hundred years ago or longer to now.

[00:06:04] The problem right now for some people may be lack of food, but most people, it's probably an abundance of food and how to deal with that. So that really changed to automation and in healthcare, I think we have tried it with things like electronic health records, but my own experience and have done some research and collaborated with others, electronic health records have actually slowed me down.

[00:06:24] They have slowed productivity. When you start looking in at the data, for example, the Bureau of Labor Statistics at the Department of Labor, you can actually see that productivity has been declining for physicians, outpatient, outpatient clinics. And so how can we tackle that? And that's when I started talking about autonomous AI, meaning the medical decision, whatever it is, is made by the computer solely without assistance by a physician.

[00:06:49] Until then, most AI was more assistive, meaning you have a physician, another provider or nurse in the clinic with the patient, and they're guided. [00:07:00] And I use it daily in my clinic, but this is not the same because ultimately the responsibility for the medical decision is still typically with the physician, even from a medical liability perspective.

[00:07:10] Also, you don't really move the needle on scalability and and automation. You still need that physician there with the patient. You don't improve productivity. So I went all in for autonomous AI. Um, and I think now I've talked long enough. Uh, but yeah, that's the background of my neural network, so that was great.

[00:07:27] There's a couple nuggets there that I wanna put on the shelf and come back to, especially around the problem with abundance, the problems with the electronic health record present. But I'd like to just focus a little bit more on your training, cuz I think a question that we get a lot is from the clinician side, how do I get involved in AI?

[00:07:43] And then from the computer science side, how do I get involved in medicine? And I just wonder if, given that you have an MD PhD, your PhDs in a technical field, and obviously you're a practicing ophthalmologist, is the MD PhD the gateway to this intersection or given the training that you've been [00:08:00] through, are there sort of side routes into that?

[00:08:03] There's a few answers if you, I, I, I'm afraid to say it, but learn to code is, is one thing which really makes a difference still, even if it's just Python. If you can build your own neural networks, there's just. A lot of ground still to be covered. I would just like to point out the sort of subtle deep computer science jab there of saying if you want to code, even if it's just Python, the clinical listeners, that that, that's a, that's a nice little jab there.

[00:08:27] I like it. Oh, okay. So I didn't want to insult anyone. It's, it is more, I mean, yeah, sure, you want some, want to build really sophisticated recurring neural networks in C++ and, and that is fine, but I think we need more people who can just set up an existing inception or whatever it is and start training it.

[00:08:45] Because I think what is exciting about, especially machine learning, is that once you have the engine, the tool, it's, it's more about the training data and how you use it. And that's where the, you know, the interesting discussion is going to be, I hope. But if you, if you [00:09:00] want to understand, I, I actually wrote, or I was asked to edit the educational material on AI for the ophthalmology residents around the world, which is called a Basic [00:09:09] Clinical Science Course of the American Academy of Ophthalmology, which is widely used even when I was in the Netherlands. We used that for our training and so they asked me to help, you know, what, what does a practicing specialist or physician need to know about AI? And it's mostly about what the potential, the risks and what is going to mean for your practice.

[00:09:30] So that's important if you then want to build it yourself. Understanding is good and, and trying it out yourself is good. But to really become involved, you need to be able to at least code, understand statistics, understand pre-arrest clinical trials, understand back and forth about what type of reference standard to use.

[00:09:51] There's a lot of sophistication there that maybe is not for every physician. You really need to spend a lot of time there. On the engineering side, I had a lot of graduate [00:10:00] PhD students in computer engineering. You know, I've now sort of focused more on, on digital diagnostics, but there, for me, it was really important that they were able to interact with clinicians, that they spoke the same language, that they understood what clinicians are dealing with and how getting your hands dirty and the noisiness of much of the data that clinicians deal with all the time, as well as of course, the privacy risks.

[00:10:25] Just the emotional aspects of medicine is really important, I think, for aspiring AI engineering students to understand. And then ultimately, when it comes together, there's the more general risks and potential benefits of AI that we all need to know about. Yeah, I think that makes a lot of sense to me and resonates with my own experience is that when I'm talking to clinicians, I try and tell them that a lot of the base models we use are essentially commodities now.

[00:10:52] And so if you're trying to think about how to use them clinically, you need to understand their operating characteristics, their inputs, and their outputs, and that's probably sufficient [00:11:00] if you want to do research, then as you said, that's like another thing, and you probably need to code and understand statistics.

[00:11:06] I think that there's a wide range of opportunities for clinicians who don't even know Python, but are still interested in this interface. Yeah, and, and ideally you have a collaboration ongoing with maybe a grad student in engineering or maybe a group there that that is Of course, in some institutions that's well done.

[00:11:21] In addition, not so much, much, uh, but yeah, I I even for when you want to commercialize it, ultimately right now where, uh, digital health companies are AI algorithms are typically seen as a commodity. Great. Thanks. This Biomark approach that, you know, I happen to have a patent on is so there's a conflict of interest here, but I do think there's certain advantages and I think some regulators also see these advantages.

[00:11:46] So I think that's a perfect segue, right, Raj? Yeah, that was, that was great. Um, so Michael, we want to transition to your research and also to your efforts to move beyond academic publication to deploy medical AI. So, you've [00:12:00] been a trailblazer for several decades, both in academia and in industry where you're founder and chairman of Digital Diagnostics.

[00:12:06] Back in 2018, the FDA issued authorization for your company's device, which I understand was the first FDA authorization for an autonomous AI without physician input. And this was for the IDx-DR device for autonomously diagnosing patients with more than mild diabetic retinopathy. So this is 2018, but I know this journey started several decades before that.

[00:12:28] Could you take us back all the way to the beginning of the technology journey? How did you choose this particular problem and what were some of the key early challenges, uh, that you had to overcome? I already sketched a little bit, you know, where my oranges were in, in, I, I really saw myself as a neuroscientist.

[00:12:46] Uh, we're interested in, in, in applying what, what we learned. And we still don't know much, much about how the brain works, but at least we, we know a little bit and start to be useful. It was really when I was in clinic [00:13:00] as a resident that I thought, I have all these patients coming to me and I look at them and they spend half a day in my clinic and there's nothing wrong with them.

[00:13:09] And meanwhile, there's all these patients coming too late because they don't have access. Or there was, was was another problem or they didn't feel like going. And in many cases it's too late. So clearly we're not finding patients who need their care and we have too many patients that don't need their care.

[00:13:23] Can we or can I not do something about it? Is there not a machine that can mimic my brain and that can find these patients wherever they are? We're talking 1996 now. And so this was still the time that there were no digital cameras most anywhere in medicine. It was just starting in radiology. And so it was not like there was a bunch of training that you could just build an algorithm on, but they said, well, let's try to see whether this can work and when improve it works.

[00:13:56] Surely people will say, wow, this is brilliant and uh, you know, we need to do this. Okay. So that [00:14:00] was the naive thought I had and which I will be going back and forth between the commercialization aspect and the research aspect. I, I cannot help that because I think applying research to actually benefit patients is in many cases what physician scientists are about and hopefully engineering scientists as well.

[00:14:15] And so part of that was selling up literally a telemedicine network to have the, the images so we could scan them, so we could train algorithms and see whether we could have a performance that was acceptable. And at that time there was no definition of, well, this is the threshold you need to meet, which has been a challenge for a long time and we will get back to that.

[00:14:35] But it started to work, came to the us, got actual funding from NIH that these algorithms actually perform. Equally good as me as as then a retina specialist. So clearly there was promise there. So it said, well what should the performance actually be? Because I knew that when I look at myself and other of my colleagues, were highly experienced retina specialists, we different about 20, 30% of cases.

[00:14:58] So the fact that I agree or [00:15:00] disagree on AI doesn't necessarily mean that AI is wrong or right. I mean there's a noise there. Can we do better? So that was, you know, we'll set it aside, but that was an interesting question. But also the research started to be fruitful. There were publications, algorithm, you know, successful publication, track record, but that didn't move the needle.

[00:15:20] It was not, people were knocking on the door saying, oh yeah, yeah, we need to do this CMS, or, I was very naive and, and maybe some, you know, listeners are also thinking, how can this be done? And. It's, it's a little bit more, it's, it's not, it's not just publish and solve all problems with a single people. No, it's, it's, yeah.

[00:15:38] And, and also I've been saying in other parts of my research about, uh, neuro degeneration in as a cause, cause of a diabetes. There's published and perish. So you can literally publish and still perish. So, but yes, absolutely. It's definitely, publications are important. They absolutely help with the IP, with convincing investors.

[00:15:55] But at that stage, I started to realize, [00:16:00] well, rather than me saying it's safe, I want the highest authority in the world, which I considered the FDA to say this is safe and effective. How do I do that? So I went, ultimately ended up going to FDA in 2010 at a meeting, said, I want a computer to make a diagnosis, nears why?

[00:16:15] And they said, ho, ho, ho. And you know, we worked together very closely, continued to work together very closely, but they were, uh, skeptical to say the least. And so I realized that. There's a long path here. I need money. Uh, and funding is not really, NIH funding is not really for that. They don't want to fund FDA interactions and a thousand dollars per hour consultants, right?

[00:16:37] That is just not what, what we as tech payer should be paying for when we pay, you know, researchers with an r01. And so people told me, well what works in pharma? The pharma model is that you patent your idea, maybe new protein or maybe falling mechanism, new CRISPR thing. And then, uh, Pfizer or Regeneron or wherever will come in and take a path all the way [00:17:00] to, uh, phase one, phase two, phase three.

[00:17:02] It's a well trot path. It's well known. And they said, well, why don't you do the same with your AI algorithms? So did, and you know, was hopefully waiting for IBM or Google to show up and nothing moved. Right? And so, okay, then other people are doing philanthropy very successfully. You know, you can raise a lot of money, but that's, for example, here in Iowa, very successful institutes for fishing research and other institutes.

[00:17:26] But that's for people who got blind at a young age because of inherited eye disease. People love giving money for making blind kids. See that is, you know, very motivating if you say, well, I'm going to generate a 10,000 pages of paperwork for ISO certification and FDA not so exciting for a retired farmer who made a few million dollars from their farm and, you know, can either contribute to making blind kids see again or a batch of paperwork.

[00:17:54] So I realized I needed angel investors and that ultimately led to funding company here. So it was not [00:18:00] like I set out to find a company, it was, Hey, I want to solve this problem. I realized I need to, I'm looking at a long runway I now call it. Um, and how do I solve that? And that turned out to be with a business plan.

[00:18:13] And I was successful rating first Angel investment money and later VC and now very proud to say for the first time growth equity last year, uh, stepped in. So for the first time ever, AI is not seen as a debt, right? Which is what venture capital really does, but now growth equity is, this is going to happen.

[00:18:31] This is not a bet anymore. And, and so I'm, I'm really excited about that newest aspect of our fundraising. But yeah, long journey, this is the sort of more commercial aspect and why that worked, because part of it was wanting to meet the highest standard of safety. And part of it was I've shared with, with at least Raj, another thing which also happened in 2010, which was that I was doing all this research and my colleagues started to notice my colleague [00:19:00] ophthalmologists, very esteemed colleagues, and certainly there was an editorial in the most widely read ophthalmology journal in the world called Ophthalmology Times on the front page by the chair of ophthalmology at Johns Hopkins.

[00:19:13] The Retinator Revenge of the Machines about my research, calling him out by name, essentially saying that I'm the terminator of retina. I'm the retina specialist, right? That's that, that hurt. And um, and it's going to cost, uh, jobs and it's going to lower the quality of care. It was partially tongue-in-cheek and partially very serious.

[00:19:32] And actually immediately got a phone call from my chair. Hey, what's going on here mean? And I was trying to make tenure at the time so you can understand it was somewhat upsetting, but actually I have it now framed in my other office. Uh, it was a really, really important moment because I realized that even though you think you do the signs right and you, the goal is to help patients, it may not always go over smoothly.

[00:19:58] And you need, [00:20:00] actually, healthcare is complex as we all know, to our maybe detriment, maybe not, but it's very complex and we need stakeholder support. We need everyone in the healthcare system to be comfortable. With the introduction of a new technology, that is very scary. So a, let's figure out what is scary about it.

[00:20:17] And I'm talking about racial bias, which we are, you know, we started talking about before then data usage liability, right? Who's medically liable for an autonomous AI making a erroneous diagnosis? Does it actually help patients? Are we going to pay for all of this? So these are the issues that I was summing up and starting to realize we need to address these.

[00:20:38] And on the other side, we're the stakeholders. These are patients, first of all, patient organizations. Clearly my colleagues, professional medical societies, the American Medical Association, it is payers, CMS, Medicare, private payers like United and, uh, Aetna ethecists. Very important. And then of course, AI creators.

[00:20:59] These are all [00:21:00] stakeholders, and then ultimately even investors. And, and then the P P P is what I call physicians, patients, and payers, right? That is the, the, the nucleus of that. We need all of them to support this to be able and value-based care organization, sorry, forgot one. I, I don't have a list in front of me, but normally have a slide where I list these like SCQA USPCF, uh, PCORI, and they all look at each other and they all are, have a little bit of a letter and take on what should happen in healthcare and, and we can talk about it, but I, I don't think it's necessary now more importantly that to get these all on the same page about something as provocative as an autonomous AI, a computer making a diagnosis which had never been done was what led me to develop with real bioethicists.

[00:21:47] I wouldn't call myself a bioethicists, but I had to learn it to create an ethical framework for AI that was published in the American Journal of, uh, Bioethics. And that was very important in the interactions, uh, with FDA, how do we deal [00:22:00] with bias? How do we deal with the clinical trial? Is there something you can measure about these?

[00:22:05] And uh, we now call it metrics for ethics. Where you say, well, there's an ethical principle, like autonomy of the patient. How can you measure that when you have a healthcare system? But you can measure it for a physician, but you can also measure it for an AI. No one ever did that for physicians, not for healthcare systems.

[00:22:21] We talk about it. Many papers on the in, in medicine are about words, but not really about measuring. And if you're an engineer, you need to have something that you can measure, so you can meet it or not, right? You, you develop something you want to know, does it meet this ethical criteria? Does it meet this ethical criteria?

[00:22:38] Same for justice or equity, same for patient benefit or, uh, non maleficience. We can go into the, the ethical discussion, but I don't think that's what your audience wants to hear more. More importantly, a are something you can understand as an engineer or an AI creator. More importantly, you can measure it and if you can measure it, we can optimize towards it.

[00:22:58] And then what is interesting, you [00:23:00] can never meet all of these ethical criteria a hundred percent. You can optimize the patient benefit by not allowing them to smoke, telling them exactly what to eat, but you lose a lot of autonomy and patients do not always want that. So you need to find a balance between the different ethical principles.

[00:23:16] And that is, I think, the contribution we did, uh, in that, in that framework. I'll just say I am completely struck by how many different parties, groups of individuals you've had to navigate in just taking this from a paper, this sort of early technology into something that's now deployed in the clinic.

[00:23:34] And so we asked you about your training in terms of MD PhD being the gateway, but I'm wondering how you learned all of what you needed to learn, uh, to sort of navigate all of these parties from the FDA to payers to other folks. And just thinking about the physicians and machine learning scientists in the audience who are interested in those skills and, and commercializing, did you have a mentor or did you have, uh, sort [00:24:00] of the right mix of people or did you just figure it out on the fly?

[00:24:03] You know, what advice would you give folks who are, are in a similar position and want to commercialize their technology? Ultimately, the objective function was patient benefit, right? And, and health equity. So that was the, the north star, the guiding principle that really helped. Now I had to figure it out on my own and then decided, part of it is, as a physician, you, you are aware of payers, you're aware of HEDIS and MIPS and, and, and things like that, and what they want to see versus what you need to record in the chart.

[00:24:31] Medical liability. So that's a little bit more. Uh, each to understand for a physician, then for an engineer. But that's, you can learn that. Don't forget, all of this has been done now and we've shown it can be done. So it's not like everything needs to be revamped again and again. There's a lot, the stakeholders are known.

[00:24:49] There's not, like every year there's a new, new group of stakeholders. It's, it's sort of a fixed system. I literally drew maps for my board showing almost like D-day, you know, those maps with all those [00:25:00] arrows to, uh, to the different Normandy beaches. Mm-hmm. Well, so I drew maps of, well, this group influences death group because it's all interrelated, right?

[00:25:08] Uh, National Committee of Quality Assurance looks at the standards of care developed their patient organizations and professional societies. So if you go to one group, it may help another group get more comfortable. So part of it is just navigating that. Part of it is learning. Who are the stakeholders being very open and not thinking that because you have this cool technology, what I call a glamor ai, people should just swallow it.

[00:25:30] It's glamor, AI's AI technology that is really cool and I love technology doesn't benefit the patient, or at least that has not been shown. And so it is. It is more than that. Ultimately, your North Star I think should be the patient benefit. How do you can benefit healthy populations and then it's just necessary to do this.

[00:25:48] But no, there was no mentor. I wish, yeah, I wish. But no, we had to reinvent it and I think it couldn't have been because autonomous AI is so confrontational, [00:26:00] right? We never think about how many rules and regulations are written, but it all assumes a human because that's been the way, that way for thousands of years.

[00:26:07] And suddenly with a computer, you need to change these almost implicit regulations. And now people started think, well, do I actually want that? But they never thought about a doctor doing it. They just, it was all implicit and now it had to be made explicit, and that's why, you know, some of the swirl was created.

[00:26:25] Got it. So you, you also mentioned that you faced, you know, some resistance here and there along this journey. And I really like that you've now framed the copy of the Ophthalmology Times that had the Retinator as the, uh, the cover page. Uh, which honestly is a, I know it, it sounds like it was a, it was a difficult time, but it is a pretty impressive nickname, uh, to also have at the, uh, the same time.

[00:26:48] Um, but you know, I, I think what this signals to me and what you've clearly overcome is that there's sort of resistance from folks at key moments along this journey. And so I'm also [00:27:00] curious, you know, you've been able to overcome that, but what advice you would have for young machine learning scientists or clinicians who are also facing resistance as they try to develop and, and deploy new technology?

[00:27:12] If it doesn't destroy you, it makes you stronger. So that's, that's, there you go. Totally true. If, yeah, when in doubt go to Nietzsche. I think more importantly, do you want to benefit patients or not? And then everything else is, sure is a hurdle on the way, but can be overcome if you, the more evidence you collect that A, this works, maybe, you know, hopefully you, you can show patient benefit.

[00:27:35] And I, I want to say that about this. I'd like to show this timeline where we had an algorithm lit, literally a multilayer known network beginning, and you have patient outcomes at the end, and there was a long journey to bring it to patients. A lot of evidence, a lot of stakeholder support, ethics, the reimbursement, don't forget, we haven't even talked about that, because if it's not reimbursed, if it's not part of the HEDIS MIPS, people won't use it if [00:28:00] it's not usable.

[00:28:00] And there's an entire team of digital diagnostics that is focused on what we call customer experience success. Meaning you take a clinic, you introduce AI, you don't drop it. And walk away. You know, you literally continue to support 'em to make sure it works for their patients. So all of that is necessary.

[00:28:18] And now finally, and I'm really excited about the papers coming out very soon. We showed in randomized clinical trials that this autonomous AI improves, that it not only improves health equity, it literally allowed clinics to go from very low percentage of people getting the exams they need to almost eliminating health disparities.

[00:28:38] So it is really doing what we envision to do. So part of why I'm mentioning is this, it was worth a journey. I mean, it was rough, but it was absolutely worded. We, we did it and with an entire team of more than hundred people of Digital Diagnostics. And thanks to stakeholders, support from everyone, because what you say is true.

[00:28:55] There was initial resistance and the Retinator, just one example. But you can [00:29:00] imagine me going into CMS and saying, Hey, you know, that said, reimbursement at $55 because here's why it should be $55, and there's a whole reimbursement framework based on an ethical framework that went in. Why that specific number that you can read up on, right?

[00:29:15] And so very grateful to all these others who were willing to do what is right for patients. There's literally 30 pages in the federal register from CMS discussing their worries and concerns, what they called guardrails and whether they should do this, which they ultimately ended up doing. You can see it live in their proposed rules, how they were dealing with this, with, with this new thing in a healthcare.

[00:29:42] So very grateful to everyone involved, at least. And finally then you need to stop me again. Don't forget that healthcare was the first to use autonomous AI. So we always complain that healthcare is so complicated and how can you change it? There's no self-driving car. John Deere had the first autonomous [00:30:00] tractor at CES last year, but it has been in wide use in US and now elsewhere.

[00:30:05] It was an announcement for, you know, the Middle East yesterday, but in widely used on patients in healthcare and nowhere else until last year with autonomous structure. So we were actually the first, isn't that exciting? That doesn't get you excited that we can actually change this and make it better.

[00:30:24] Yeah, I think that's amazing and I think it's a perfect segue for something that I think we wanna focus on for our next topic. So just to pin down a couple things, one of your first release products was autonomous AI for the diagnosis of diabetic retinopathy. And just to be clear, that means that you have a device.

[00:30:40] That can be operated by a technician that can render a diagnosis of that disease without the oversight of a physician. Correct? Correct. And so that's the autonomous in the autonomous AI. And so I think the thing that is the real differentiator and the real thing that you figured out first that no one else has is how do you get someone [00:31:00] to pay for that?

[00:31:00] So I would like to like focus on the reimbursement model because. Just to be clear, sort of what I do is kind of cheap, like technology development is difficult sometimes, but it's kind of cheap and easy if you have the right data going that last mile, getting an actual device that can go into a clinic with a patient and then actually getting a payer to pay for that is millions of times more difficult than the original tech dev is.

[00:31:24] So I'd like to understand a little bit one, how you got it from whole cloth, got a reimbursement model for that, and then how that reimbursement model operates today. So if someone goes into their clinician's office and an uh, and a digital diagnostics device comes in, sort of what does the back end of that transaction look like?

[00:31:41] Absolutely. So reimbursement is extremely complex. Set it out there. It is fraught with legal risks if you do it wrong, to what seems like in order feels normal, it's a felony. So I will be careful in how I [00:32:00] respond. I hope you understand that. We published our reimbursement framework for AI half a year ago, and it's probably best if I go a little bit through that.

[00:32:10] I think there's two challenges. As an AI creator, what do you charge for that? Do you And I literally went through that and in interactions with Congress, so I went in front of the Senate's Finance Committee and explained early on, just after FDA, uh, de no authorization, actually, it's not approval. I'm not allowed to say that it should be de no authorization.

[00:32:31] FDA doesn't allow me to say approval. And I went in front of and say, well, how are we going to pay for this? We can pay it based on the marginal cost, meaning you have a million diagnosis with AI. So the, the R&D and all the science that you did is, is amortized. And now let's do one more diagnosis.

[00:32:48] Well, that's maybe a little bit of electricity and maybe a lot of wearing down of some atoms in the, you know, GPU circuits. Right? That's very little and definitely it's not a sustainable business [00:33:00] model that no investor will go for that. And AI creators will probably say, I'm not going to do that because it'll be a money drain.

[00:33:06] There's no money to be met. And to go back to your previous example, is that kind of the pharma model also where you're recouping R&D and then you have some sort of marginal costs that you are getting on every patient that, um, would be drug, is that Yeah, it's a little bit more complex because then the part of CMS that we're talking about here is only allowed to reimburse what a physician charged.

[00:33:27] That's the way the Social Security Act was written. So it, it might seem that you want to pay for something, but there is actually very strict laws and regulations and. Jail time if, if you don't do it in the right way. So that's why it needs to work. But, but that's the other aspect. The fir, but the first problem is literally what do you charge an AI creator?

[00:33:48] Another way what people then recommend is what is called cost effective analysis. Meaning you look at, there's this service, we know it is benefit to patients. And for example, the diabetic [00:34:00] eye exam is a good example, but there are many others where we know that if you do diabetic eye exams and I as a retina specialist get about $170 to $300 for those per patient, even if you do $500 to $600, even if you pay that, it's still worth the cost savings on the back end because of avoided blindness, visual loss, uh, very expensive surgeries.

[00:34:23] So one way for an a creator to say, I will just go to just below the cost effective threshold, but then you would make it more expensive than it currently is. When my, my whole goal, my North Star was. Improving excess lowering costs. So we don't pay a hundred percent of our healthcare, of our income to healthcare in 20 years.

[00:34:42] And so not acceptable for me as an AI creator. And so there are various strategies that you can do for what you charge as an AI creator that is irrespective of whether you get it reimbursed, what should it be based on? So ultimately what I said, and that is not, may not be [00:35:00] possible for all, but definitely for autonomous AI, assume fields, is there's a willingness to pay currently for about 20 to 30% of patients who are getting these exams.

[00:35:08] So as a society where we're paying for other healthcare, there's a willingness to pay. We're already paying a certain amount of money for 20 to 30% of patients. Let's take the same amount of money. We wouldn't be paying more and just get it to a hundred percent of patients who deserve and needed. And that's what we based it on.

[00:35:25] I called it the equity enhancing payment. And so that ended up being $55. So it had to be $55 for clearly, right? I mean, there's a formula here. And so that was the basis for what you charge. And then ultimately many, many meetings with payers. Why it should be this. And then, you know, we, we could talk about the mechanics, but I think that's less interesting that the fact that you need to start with what is an appropriate payment.

[00:35:55] If you start with $600 for this case, where currently CMS is paying when it's $50, [00:36:00] very unlikely that someone will use it because the reimbursement will probably not go up. So you as an AI created, can charge $600, but then the physician only getting way less than that, no one uses it. It's a tricky balance to be found.

[00:36:14] But this is for what is called a physician fee schedule. Now I'm getting into the really wonky details. There's other payment schemes which are very different. Other AI companies do things with NTAP, which is a temporary, uh, reimbursement. It's a challenge right now. As you can see from these internal discussions in CMS that they exposed and in Congress as well, are we going to pay for all of this because we want to make sure it doesn't increase healthcare costs?

[00:36:39] Hopefully. And my goal is definitely to decrease healthcare costs, make it more affordable, more accessible for, especially for people who currently are not getting the care they need. Stepping outside of like digital diagnostics, if there is a wave of similar devices that are created that are essentially sort of screening tools, I believe you in your [00:37:00] sincerity that your goal is to drive down healthcare costs.

[00:37:03] But are you at all worried that this sort of WRI large will increase cost for the healthcare system as a whole? Meaning you can look to other areas of medicine where screening hasn't improved patient outcomes and we're sort of spending money for nothing. Are you at all sort of worried of that at sort of outside of the digital diagnostics ecosystem?

[00:37:20] Absolutely. So I care greatly about patient benefit, population benefit and all the things. These are the north stars, right? And then you work your way back from what is the evidence for what is best for this patient and what role should can AI play? That's how I would do it and how we see it. So it's not glamor AI.

[00:37:36] If the school technology, let's find the use for it now it's the what does the patient need, what does the population need? And so if you start with that, there's a lot of things to be done and there's a lot of evidence already. So let's start with that. And then what you see is all stakeholders aligning around that.

[00:37:51] And it becomes much harder for glamor AI where, you know, people get excited about technology and it's, you know, a thousand dollars per patient and it [00:38:00] sounds great, but it doesn't, like you say, actually move the needle. Less willingness to pay clearly right now, especially. And so sure you can try that and maybe patients self pay for it.

[00:38:11] And that's the choice. But I think especially based on the failure based care movement, which your listeners probably are aware of, there's more and more resistance to just doing things because of coolness of glamor. And so I, I think the guardrails are in place already. There's a regulatory process, there's an ethical framework, and there's a way to get payments.

[00:38:32] So why not? That's probably the easiest part. You can fight for a long time to get a thousand dollars for your AI. That is probably, maybe have a negative benefit for a patient. It's a long journey. Why do that and not do the easier journey that we just sketched. Mm-hmm. Yep. And so I think I, one more reimbursement question before we conclude this, uh, policy walk corner and move on to the Lightning Round.

[00:38:54] So if I could summarize why digital diagnostics has been successful in this area of [00:39:00] ophthalmology. There's a clearly defined task that a physician already does. There's a billable interface to that task. So reading a retinal scan, there's already sort of an existing way to bill for that service. And then you can sort of productize this into a medical device that essentially does that service.

[00:39:15] So there's a lot of pieces in place. Obviously you, you broke the path, um, to get a reimbursement model for that. But I'm curious what you think about. Other areas that people are excited about medical AI, where at least one of those conditions doesn't hold. So for example, a lot of people are very excited about risk scores.

[00:39:33] So you can have an AI that pulls in the patient's entire EHR data and predicts things for like mortality, decompensation, sepsis. Um, and to me, those business models have always been very tricky, but sort of given your deep expertise in this field, I wonder if you think that other areas of AI that are kind of more amorphous, if there's a sort of a sustainable reimbursement model behind some of those risk scores or predictive models where there isn't this sort of [00:40:00] natural billing procedure built into it?

[00:40:03] Yeah, so indeed, I, I focus right now on the preparation because it's about the evidence, it's also about the regulatory approach and uh, the safety aspects. And so as we have seen with Dr. Obermeyer's paper in science a while ago, you're probably aware of, um, he's a friend of the podcast. Oh yeah. Okay. Well, even better.

[00:40:24] So I don't have to tell you anything new. Um, but clearly the problem there that it's so easy to find a trap to use proxies for what you're, you think you're measuring or optimizing as an objective function rather than what you should be optimizing for. And there is no regulatory offsite whatsoever. And so I think, uh, yes, there is a potential benefit, but show that, in improved patient outcomes.

[00:40:50] It may be harder to get definitely someone like CMS to pay for that, but I'm not, CMS payers may be more interested, but I'm very [00:41:00] worried about the backlash that we already saw from HISA research, which now has led to, uh, the Office of Civil Rights in HSS looking into this and essentially saying, Any racial bias or any other bias in an AI is the responsibility of the provider using it rather than the liability being taken care of.

[00:41:18] Because the FDA looked at racial bias, and in, in our case, we proved that there was no racial bias in AI. So I think I would be careful until the, the safety aspects are better understood, and, and especially when I work with the Federal Trade Commission, they're looking into this OCR is looking into this.

[00:41:41] I don't want to damage the field, or I have, have to feel damaged by Congress saying, let's shut down AI because we worry too much and there's too much bias and harmful effects going on. And as you know, there were actually harmful effects of this AI. Yeah. So, so first, do no harm applies to medical AI, just [00:42:00] as well as it does to, to human intelligence.

[00:42:02] This is back to the ethical framework. Patient benefit or do no harm normal efficience. Uh, justice or equity? Do we do it equally well for everyone? And are we leaving the patients autonomous? Meaning what about their data? Right? Who, who's the ownership of the data? Are we taking care of privacy? Then can they make their own decisions?

[00:42:20] Absolutely. Uh, it starts with ethics. Great. So I think we will transition to the Lightning Round. Now.

[00:42:33] The Lightning Round is a series of short questions that we're curious of your thoughts on. Also, just an opportunity for the listeners to get to know you a little better. So some non-medical, non-technical questions. The goal is to keep the answers as brief as possible, but no briefer. So, uh, a couple sentences should suffice and we'll sort of knock these out in quick successions.

[00:42:52] Are you ready? I'm ready. So the first question is, if you weren't a medical AI researcher, what would you be [00:43:00] doing professionally? My three kids are all, uh, uh, well, one is becoming an engineer. They're all engineers, they're all this service academies. Uh, I'm really excited about the potential for space exploration and especially asteroid mining.

[00:43:15] So there's, there's a lot there, but in the medical field and AI, there's so much to do. There's so many exciting opportunities right now. It's, I'm just amazed. I wish I was 30 years younger and I, I admit that I did not have, I admit that I did not have asteroid mining on the Bingo card. So that was an excellent Lightning Round response.

[00:43:35] Yeah, I love it. I love it. Michael. Will doctors still be responsible for documentation in five years or will generative models like ChatGPT have taken over that task? It was interesting the way with went with electronic health records where, I was a big fan originally when this, you know, was at the short lift level, remember the sixties and, and mycin where it was, the excitement was about the doctor [00:44:00] type stuff, not about the AI behind it.

[00:44:02] And so I was excited about this. The money or the time saving aspects, the more efficient and, and like I mentioned, we lose productivity not to gain it. And then what we'll turn out is that if I don't type in my own finding, there's no reimbursement on the backend. If I copy a note, reimbursement falls apart.

[00:44:22] And so yes, you may become more efficient and there's probably going to be a way in where that is, depending on how you see the payer. They want to pay for what a doctor actually does rather than, and what, and hopefully what, you know, this is the value-based care part, but what, uh, what benefits the patient not for the time that the doctor spends, that the doctor's now saving because they're using AI, right?

[00:44:47] And so, It's going to be interesting, but I'm not expecting five years that this will be widespread. Got it. Alright, so, uh, next Lightning Round question is, what is your favorite band or musical artist? [00:45:00] Oh, I wasn't expecting that. Um, right, that's the point of the lightning round. We saw the sh Chicago Symphony, uh, with Beethoven, which was great last weekend, so this was wonderful.

[00:45:10] And I've been listening to Lana Delray, which where I love, uh, singing. Oh, nice. Yeah. It's, it's all over the place. That's excellent musical breadth. Lana Delray and asteroid mining are fantastic answers to the, to the light room. All right, and next question. Uh, do you think things created by AI can be considered art?

[00:45:30] Ah, that's a cool one. Do artists consider it art? I would say that, you know, the, the experts and we are seeing some discussion there, right? If you look back at healthcare, what patients want is not what the doctors agree or disagree or ai agree or disagree. They want the best patient outcome. Whatever gets them to the best outcome in a ethical way is best As a user of art, I care about the quality of the art and you know, that can be newness, it can [00:46:00] be what looks visually appealing or emotionally appealing.

[00:46:03] I don't care about the process we got there. So as a user, yes, uh, as an artist, I would say no. So fair. Okay. And so again, Lightning Round questions. So your, uh, responses. Yeah, yeah. But these are complicated. Yeah. Well, um, so this one, I know that you're going have a long answer to try and I think compress as much as you can.

[00:46:31] So will AI in medicine over the next decade primarily be driven by computer scientists or clinicians? And this is a forced binary response. I'm not answering, but I, I need my time. So either I have a long answer. I have, I need time to think. It takes time. It takes time to run the compression algorithm. I understand.

[00:46:53] Yes. Take some time. Wow. I think computer scientists are [00:47:00] the rate limiting step of this boat. There's many great limiting steps. Many, many, many, many, these are just two of them. But then I, I would fear towards the, you know, engineering, computer science aspects. Got it. Great. Last Lightning Round question. If you could have dinner with one person that are alive, who would it be?

[00:47:21] I had an interview recently. Uh, I said Elon Musk because there, I'm interested in autonomous cars and I think having an ethical framework to start with. We, we really help get this across. Where now it's sort of, you know, who's liable. Um, do we want this? I I I think it can be easy if you just start with ethics and that's another AI feels Well, I mean, I, I don't even know the name of the CEO of John Deere, but I would love to have dinner with him.

[00:47:51] He's probably somewhere around here in Iowa, but I've never met him. We'd love to meet him. I'm guessing it's not John Deere, right? It's not John Deere. It's not John.[00:48:00]

[00:48:03] Alright. So I think that that was a, um, top five, uh, Lightning Round performance. I think that the, the entropy of the responses was excellent and we got, uh, we got a little bit more of Michael Abramoff than I think we had going into that. So I think that was an excellent performance. So we're gonna, we're just gonna wrap up with some big picture and some conclusions here where we kind of zoom out from the work that you've done so far and try and, uh, take, take stock of the AI medical field as a whole.

[00:48:29] One question that I'm curious about, I was a, a snotty computer science undergrad. I have this story about when I, you know, I met my wife who's a clinician that I told her I was gonna replace her with AI and computers. That's what's over. Well, I'm sure, yeah. Is as far as pickup lines go, I would say that's bottom five.

[00:48:46] Um, nonetheless, so I, I have some sense that there was deep skepticism towards AI in the medical field. Certainly not to the extent that you sort of lived to the experience. So the, the Retinator, I think you sort of [00:49:00] hinted at was not exactly something that you were flattered by. So what I'm curious about is, over the course of your very long and impressive career, do you think the medicine as a whole has become more receptive to AI?

[00:49:13] Or how has that sentiment changed? Absolutely. It, it changed entirely. Part of that is just a generational change where people are just more comfortable with computers. When I started medical school, I was probably the only one in five City blocks with its own Apple two, which I built myself. And so that, that has entirely changed.

[00:49:35] And like I said, I think healthcare should be proud that they were the first to use autonomous AI widely in a, in a very ethical, sustainable manner. And now I think it's, it's a wonderful environment for where we can innovate, but we need to do the right way, not innovation for innovations sake. And, and that's right.

[00:49:54] By the way, about the Retinator, and sorry about these long answers, it was an interview with me in the American Medical [00:50:00] Association on the website. This guy's doing AI the right way, so one can turn it around also. And what doesn't destroy you makes you stronger. And in this case, the Retinator absolutely made me stronger and, and the field of all, I think stronger got, yeah, I, I have seen that too, that just having seen, um, my wife go through training that medical students now and residents kind of just take for granted that something like this is going to be integral to their practice.

[00:50:28] And, you know, I, I don't wanna speak out of turn here, but I think that there's less of a mystique around the clinical reasoning process. Certainly safety, patient benefit, all of that's always top of mind. But I think that they're just used to working with computers to make all kinds of decisions in their life.

[00:50:44] And so it almost would seem unnatural if that weren't also part of their patient interaction in clinical practice. So I think that you're right that there's like a generational thing going on here too. It's still interesting. You're so right, but that, you know, some, you still see so often that [00:51:00] people think the doctors are always right.

[00:51:02] And that when there's a disagreement between a doctor in ai, it's therefore the AI is wrong. The whole continuous learning debate that we didn't get into is, is, is part of that. And then you have the others and they suddenly think that the computer's always right and let, let's try to find a middle ground here where we know that they're both perfect, which they are.

[00:51:19] Right, right. Thanks. So I, I, I think you've touched on this, uh, question quite a bit, but maybe you can give us some just concluding pearls. Um, you've given everything happening in the field, given your work, how do you think clinicians should think about the impact of AI on medicine? So I discussed it with the residents and the Fellows and, and even medical students all the time.

[00:51:43] Of course. I think there's going to be a lot more of it now that, uh, the, the part is sort of. At least it's been demonstrated. It can happen. So there, there's a lot more going on right now. There's a lot of ais and development have reached commercial [00:52:00] success, not only in direct patient vision, but but elsewhere.

[00:52:03] Elsewhere as well. So there will be a lot more of that and they have to come to grips with that. But don't forget, in most fields there's so much underserved patients. 50, 60, 70% of patients do not get appropriate care that they need and deserve and will have better outcomes. So there's so much work to be done.

[00:52:25] So we should see it not as well. There's limit is my opportunities. It expands my opportunities. We can now reach more population and I actually can be of more benefit to these patients and to this population. So part of it's death. Will there be less diagnostic and more interventional? Yes, because I see autonomous surgery, the real autonomous surgery where there's a full uh, surgical procedure.

[00:52:49] As something a little bit harder to validate and prove that it's safe because you literally make a decision every microsecond a medical decision and so do you need to validate it as a [00:53:00] series of decisions or FDA and and, and the field hasn't decided on that. I think. I think that's more decades away before we are feeling comfortable with that.

[00:53:08] There's definitely. Especially in ophthalmology, which is always leading like in gene therapy, but also in, in automated surgery, coronary surgery, that's almost fully automatic, but that's a very, very specific, very narrow field. Um, so anywhere there's interventions, it's good to have those skills more than the diagnostic skills, which are slowly, computers are typically better.

[00:53:31] Great. So I think the next question is an area of disagreement that we might have, but the disagreement might be on semantics versus on the foundational thing. So I've been on the record before as being somewhat of an explainability skeptic. We wrote a paper a while ago trying to lay out why we think that explainability for AI is not a good trust mechanism, cuz I'm sure that you hear this a lot too.

[00:53:55] One of the things that I always get asked is, how can I trust this algorithm if it can't [00:54:00] explain its reasoning to me. So I'm curious if you believe that explainability is important for winning trust of the healthcare workforce for a particular AI model or algorithm. So I'm going to give a, a complex, uh, this one deserves a complex, so go for it.

[00:54:19] Oh, polychromatic answer. So a physicians, some specialists do not know how they come to a diagnosis either. It's literally a black box, uh, skin, melanoma famous example. There is no biomarker for skin or let alone a set of biomarkers for skin melanoma. It's literally gestalt. Um, does that mean we cannot figure those out?

[00:54:41] Active research right now? Are there needs of the population where we may need, and, and I will go into the semantics of explainability, but where we don't know what it's doing and it's literally learning from examples and we validate it enough that we trust it because we wouldn't know what to look for if, if we, we said it [00:55:00] explainable.

[00:55:00] It cannot tell us, well, there's a, a hemorrhage here or a pigmentation there, or the, the, the border is irregular or there's a calcification. So, but then I think there's different aspects and, and I try to define it in the paper I wrote with with FDA on this because it is about metrics for ethics, right?

[00:55:16] And one of them is how do you measure explainability or at least quantify it? I think there's explainability after the fact where it's about. Well, these are the steps I took as a decision maker to come to my conclusion. And here's the reference in the literature maybe, right? That's there's one way of considering explainability.

[00:55:34] If I have a colleague and I ask them, what do you think that typically that's the type of answer. There's another more algorithmic level, which is at a unit level, can you prove that different parts? That's how we validate typical computer codes, right at the unit level testing and then larger aggregates, uh, at the system level.

[00:55:53] That's another type of explainability, more at the algorithmic level, and then ultimately, [00:56:00] almost at a pathological level, do we know what it is doing in terms of the, the pathophysiology of the disease we're trying to detect. And now I'm talking about my markers. These are three very different types of explainability.

[00:56:14] Uh, we call them fallibility, explainability and transparency in the paper, and we differentiated them. So I think it's, uh, a little bit more complex, but I think. You refer mostly, well, actually I don't know what you're referring to, but I think you're referring mostly to this biomarker versus non biomarker, where I think if you have biomarkers, these are priors, why not use them?

[00:56:37] One of the biggest limitations in AI in healthcare is going to be the sparsity of data. Unlike self-driving cars where you can just have a car drive around that you have millions of images to train from. There's all these ethical aspects with getting images from patients. There's radiation. It's dangerous, it can harm them, let alone normals, which is actually more of a challenge, as you probably know, [00:57:00] because normals, why would you expose 'em to radiation just to have normal, and you need many normals to make a good ai.

[00:57:06] And so we will always, in my view, be data limited. Especially there's a problem for more rare diseases where there's only 3000 melanomas in the world in the eye per year. So if you want 50,000 examples to train your AI, you know, there's a, there's an issue here. So if we know that our well-established priors that we can build in and build tactics for that are racially in variant, for example, probably a better way to go.

[00:57:32] But if you don't have those, we will have to do, you know, where it's unexplainable. But that's, that's from the, the bio pathophysiological algorithmic aspect that helps you, by the way, at the unit level testing that I mentioned earlier. Because now you have detectors, you have a combination of detectors, you can test each of them.

[00:57:51] And if in our case the algorithm fails on, on some patient, on some image because, oh, this detector didn't fire, there's probably, you know, we, we need to [00:58:00] see what is flaw there. It makes it easier to tune it to new populations or new dataset. Not that we are doing that, but with different applications than the diabetic eye exam.

[00:58:11] So I think there are uses for this, but it shouldn't prevent us that, oh, it's something is not explained and we should never do it because ultimately it's about patient benefit health equity, not about, And if you can prove it's safe, So as I suspected, we actually don't disagree. I think that I'm totally on board with you as verification and validation as a really strong trust mechanism that if you give a doctor evidence that these parts have all been battle tested, we know how it works.

[00:58:41] We do that in other areas. It's usually called an RCT for for drugs where maybe we don't understand the mechanism of action and things like that. And then I really liked your other part of that, which is I'm gonna take. What computer scientists would call an end-to-end task, which is where you put in the raw images and the label for the disease and have the [00:59:00] computer sort out all of the different components that may or may not be present in an eye that gives rise to a disease.

[00:59:05] And I'm gonna decompose that into a set of biomarkers that I know are more reliably captured, don't have all of these other noise artifacts, may not be influenced by gender or racial biases and things like that. So again, like I think we actually totally agree with that. You like a smart decomposition of the task that you want the algorithm to learn.

[00:59:25] This is something that I, I certainly don't find controversial and I think is kind of this secret sauce in a lot of areas where Deep Board has had success, where we actually sneak in some domain knowledge to give the algorithms kind of a headstart on what we actually want it to learn versus having it learn everything Tabula Raza from scratch.

[00:59:44] And so indeed the, the, it's, it's gray. There's not black and white, I would argue. Why waste so many training samples on, you know, testing edge, right? I mean, why, why does it learn? Mm-hmm. You, you waste 50,000 samples to learn at an edges and edge. Well, I [01:00:00] get to told you that, right? I mean, and there's ways to do that.

[01:00:02] So that's why I mentioned prior, so I couldn't agree more. Got it. Yeah. Yep. Awesome. So we have a related, uh, I think related question, will machine learning in your view, exacerbate healthcare disparities? I think I know what you're gonna say, but, uh, but please, please go. Well, so we have to proof and hint that it, it can improve them and literally eliminate them.

[01:00:21] So that's all I want to share because, you know, read the paper when it comes out. I wish they would just, you know, give us, accept it, but yeah, it'll come soon. And, and so, uh, yeah, it absolutely has the potential to, um, to increase them as, uh, Obermeyer, uh, paper shows, right? That, that, and so it is again, The right use at the right time and for the right group of patients, uh, you know, validated in the right way.

[01:00:47] And it can absolutely randomized clinical trial proof. It's only one for the diabetic eye exam, but it can be done. Great. Um, so I, I'm [01:01:00] really happy that I get to ask this question, um, because I, I have no idea what you're gonna say. Um, but what is your most controversial or contrarian opinion on medical AI?

[01:01:13] Um, that has shifted a lot over the years I think. So it used to be, you know, this can never work. You're just a doctor. You dunno what you're talking about. So, um, this can work, but you know, never going to be, people are still saying FDA will never quote, uh, approve this. Mm-hmm. Literally. Um, and, and so, um, so these have been controversial.

[01:01:39] Takes never, reimbursements never will improve. Health disparities will never improve outcomes. So been through all of that. So it seems that it's not controversial anymore. I'm still running into a lot of interesting debates about autonomous versus is just of AI and why we need that. But where it seems so obvious to me, and I can explain it in [01:02:00] terms of productivity and, and what economy is all about.

[01:02:03] But, um, so that's interesting. What I love is that you've had controversial opinions your entire career, but you keep proving them to be non controversial, I'm afraid. Um, but, but one thing is what is interesting, what you run into is, is what we call incidental findings. Mm-hmm. And so, and it's interesting when you go into a healthcare system, And you propose to use AI because it's better for the patients, better for the populations they serve, it's it's billable, et cetera.

[01:02:38] So everything must being taken care of. And there's still resistance because clinicians may object and say, well, this doesn't replace what I do, which is a full eye exam. And it doesn't, it's for a specific disease. But what is interesting, and that is maybe controversial, I say, let's look at what is beneficial for patient outcome.

[01:02:59] They should be [01:03:00] doing this. Mm-hmm. Why are we paying for this? And that's a, that's a very interesting debate and I've somewhat controversial take that. I think we should start and work our way backwards from patient to population benefit, meaning both health equity, you know, equally uh, applied, but also patient benefits.

[01:03:15] First do no harm, and then second do actually good. Great. Right. Got it. Thanks Michael. This has been a lot of fun and very illuminating. I have one last question for you. Uh, what are you most excited about for the next five years in medical AI?

[01:03:32] The scalability of, of what we have done. Meaning like, uh, you said so rightly, uh, it's only one specific disease. Can we apply to many other disease? We have AI's in our fields. Many other groups, research groups, companies have AI's in other fields. Do we, there is a pathway, but do, can we scale it quickly enough to benefit a lot of people [01:04:00] on the short term because we need to, healthcare is too expensive cause in too much, it's not reaching the people we need it in many cases.

[01:04:06] Uh, can we better? I mean, equality is awesome here in the us. I mean, let me, the reason it came to the US is because healthcare is so much better than anywhere else in the world, but it's an equally distributed, and so can we use your AI to equally distribute it in this short period of time with, you know, 5, 10, 20, 40 AI's that is.

[01:04:25] You know, let many, uh, flowers bloom. That is really what needs to happen here and what's, what's what I hope will happen sooner rather than later, if only for my kid's sake. And your kid's sake. Right, right. Well, I think that's all we have today. I would just like to say thank you for being on and sharing what I think is just an amazing amount of tenacity and grit and foresight, uh, as one of the pioneers in this field.

[01:04:49] And I think the entire community will have benefited from the pathbreaking that you've done and kind of showing us how to go from a real unmet clinical need to an actual product that can help [01:05:00] patients. So Michael, I just wanna say thanks for being on AI Grand Rounds today. Was great having me and thanks for having me.

[01:05:06] Thank you, Michael. Have a good day.