AI-GR Podcast 13 12.15.23 Zak Kohane

[00:00:00] Now we're in a very different realm and that's because of just empirical success. And not just empirical success in the hands of a few experts, but the fact that we actually have these tools. Let's get concrete. The fact that a mom enters the history of her child who has been having endless headaches, trouble walking, trouble chewing.

[00:00:30] No doctor is able to tell her what's going on. And out of desperation, she types in all the results of the reports and of the histories and physicals into GPT-4. And she's given a diagnosis, which she then goes to a neurosurgeon, gives him the imaging studies and says, GPT-4 thinks it's a tethered cord syndrome.

[00:00:51] And he looks at the image, looks at the child and says, yep, that's what it is.[00:01:00]

[00:01:00] Hi, and welcome to another episode of NEJM AI Grand Rounds. I'm Raj Manrai, and I'm here with my co-host, Andy Beam. Andy, this episode is a special one. We have our editor-in-chief, Zak Kohane, as our guest. We've both known Zak for a long time, but I think even we learned some new things about him during our conversation.

[00:01:19] He brought some nice scotch to the recording studio at the NEJM offices for our conversation, and as always, it was tremendously fun and insightful to listen to him. Yeah, it's hard to sum up Zak succinctly. I always like to say that he's one of my favorite high entropy personalities. And, by that, I mean that it's very hard to predict what he's going to say next, even for those of us who have known him a long time.

[00:01:44] So, I think he would break most of the language models. Even them trying to break what he would say. Would it be fair to say that his, his temperature has turned up? He's a high-temperature language model for sure. GPT-Z has a temperature of 1000. I think the other [00:02:00] thing that struck me is just what a huge impact he has had on me, and I'm sure that you would say the same, but also on the field as a whole.

[00:02:08] We've been to several events recently that Zak organized. And his progeny are just both prolific and numerous. He has really launched -- And still

in touch with him. And still in touch with him. And they're all still in touch with him. So, we talk about it on the episode, you know, the first student he ever mentored and his most recent graduates, they were all there and they were all talking about their experience with him.

[00:02:31] And I think, as a young faculty member myself and a scientist and a father, he sets like such an example across all three of those verticals. And I know I feel super lucky to have gotten to do a postdoc with him. He's one of the high-resolution simulations that I try and keep in my brain. Like what would, how would Zak react to this?

[00:02:51] Or what would Zak say? You know, again, given his level of entropy, that's a difficult thing to do, but it was really great to drill down on what's [00:03:00] going on in AI. But also what's going on at the journal with him. And again, uh, he brought some very nice scotch to this conversation. Uh, you know, Zak doesn't ever cheap out on things like that.

[00:03:09] So, this is one of the most fun conversations we've had this year for a variety of reasons. I totally agree. And I couldn't agree more that I think he's had such a big impact on the field and on the lives of his mentees. And we asked him about this during the episode too, you know, what's your approach to mentorship?

[00:03:30] We've seen it firsthand, the two of us, while in his lab and now as members of the journal, and as faculty members in his department. He said it very succinctly. The first and most important thing is to care, and so he cares about his students. And it shows in the way that he interacts, the way he runs his group, the way he runs the department and the journal.

[00:03:51] And I think it's a big part in addition to his, you know, being funny and entropic and everything. I think it's a big part of why people want to [00:04:00] stay in his orbit for a very long time. Stay around him, learn from him, and be close to him because he really is a special, special person. Completely agree, and if I also, my high-resolution simulation of Zak right now probably tells me he's so uncomfortable.

[00:04:14] He's so uncomfortable right now. He's so he's like, let's move on guys. Stop. Stop. Stop sucking up. Like let's, let's keep going. Yeah. So, that gives us a natural transition to talk about another important event. So, uh, season one, the first year of AIGR is in the books. We've interviewed some amazing guests.

[00:04:32] I have to admit that I never thought that Mark Cuban would be on an AI and medicine podcast. That was kind of a serendipitous event. And again, was really impressed with him. He's a fellow nerd. I walked away from that conversation really appreciating how deep and how nerdy he will get on a topic. He will go deep and actually understands the technical side of health care in a way that I found refreshing and surprising.

[00:04:56] Yeah, I totally agree. It's also extra funny to me [00:05:00] because we spent a serious amount of time during our postdoc days in Zak's lab just talking about basketball and Mark Cuban and the argument he had with physicians at the time. And now in a weird way, it feels like all of that time that many people probably thought we were sort of just wasting time together has prepared us for some very interesting and fun conversations with folks like Mark.

[00:05:23] I also think that we interviewed a lot of really amazing clinician scientists who've really put two sets of both very demanding and difficult skills into their minds at the same time. And, so, they articulated this skill set very well. I'm thinking about the conversation we had with Ziad where he talks about both the medical knowledge and the sort of technical economics and machine learning skillset that he developed to be able to do such creative and interesting work.

[00:05:55] Lily Peng is another one that comes to mind, one of our early conversations, really amazing [00:06:00] clinician scientists. That was a huge trend for me from this season was the rise of the clinician scientists in AI. So Ziad, Lily, Atul, Ewan, like so many folks who are bilingual in both medicine and AI and that has proven for them to be like a really powerful combination.

[00:06:18] I think another big trend, our first episode launched December 15th, 2022. So, we were still one month, not even a month out from the launch of ChatGPT. And I think by any measure, the conversations this year were dominated by large language models in ChatGPT. You know, we were pretty lucky. What's a large language model?

[00:06:36] Uh, well, listen, yeah, I think we all, we all know what that is by now. I think we were really lucky to actually get to talk to Peter Lee about GPT-4 before it was made public, that I felt like really privileged to get to hear his thoughts before the madness of GPT-4 had been unleashed. And that momentum carried through the rest of the conversations that we've had this year.

[00:06:59] I totally [00:07:00] agree. All right. Should we jump into our conversation with Zak? I think so. And on to season two, next year. The NEJM AI Grand Rounds podcast is sponsored by Microsoft, Viz.ai, and Lyric. We thank them for their support. And now we bring you our conversation with Zak Kohane on AI Grand Rounds. All right, Zak.

[00:07:22] Well, welcome to AI Grand Rounds. We are thrilled to have you here. I am truly honored. You two are the most famous part of *NEJM AI* at this point. The podcast exceeds the journal. I hope not for long, but for now, it really does. Congratulations. Thank you for that, Zak. So, Zak, this is a question we always like to get started with on AI Grand Rounds.

[00:07:46] Could you please tell us about the training procedure for your own neural network? How did you get interested in AI and what data and experiences led you to where you are today? Well, first, let me just say, feel free to cut me [00:08:00] off because I love to talk about this topic. I remember distinctly in high school in Geneva, just being flabbergasted that programming a computer actually worked.

[00:08:15] And so I had an HP 25C, and I programmed it. And I could run a little moon lander on it, and it was, it seemed miraculous to me. And using retrospective editing, I'm going to claim that it was just the beginning of the ever-moving goalposts for AI. Having a machine do some intellectual activity, systematically, again and again, amazing!

[00:08:47] And so, that was also informed by my love of science fiction. Growing up in Geneva, I, I'm not exaggerating, I probably read about 50 science fiction [00:09:00] novels a year. And Arthur C. Clarke, and Asimov, and Stanislav Lem were three of the writers who were writing a lot about robotics. And so that gave me perhaps the long view, although it seemed very long.

[00:09:19] So, when I first came to this country at Brown University, we still had mainframes and there was a great social nexus around the different terminal centers around campus. I went to one such, and I saw for the first time a, what's called a detail terminal, which had one of these IBM Selectric type balls, moving along and typing furiously, and I was looking at it.

[00:09:42] It was a lot better than my HP 25, and I asked one of the students who were super, who was supervising that center, how did these things work, and explained their computer terminals linked to the IBM 370. And, I said, how do you control them? He said program. What's a [00:10:00] good programming

language? He said APL, literally, a programming language, which turns out to be a really good language for doing linear algebra.

[00:10:10] And so I picked up the book, taught myself how to work in APL, and I said to myself, how in the heck does APL get executed. So, I asked someone, they said, machine instructions. So, I took a course in Assembler and built higher-level abstractions out of Assembler. And I said, that's great. It's no longer a high-level language, but how does that work?

[00:10:34] And so I took another course at Brown on building electronics. And they gave us a book, *Zen and the Art of Motorcycle Repair*, as the required text. And the rest of it was building circuits. I think the final project was a 4-bit CPU built out of flip flops. And so, literally, every time I would flick a switch, it would be one clock cycle.

[00:10:58] And I remember I needed a [00:11:00] carry bit to go from one set of 4-bits to another, and it was getting there too fast, and so I could actually just lengthen the wire, and it would arrive there at the right time. And so, that was one way of thinking about how these things work. Then I took a course with, uh, Eugene Charniak, who was an early AI investigator.

[00:11:21] And among the many projects we did was something called a Constraint Propagation Network to do visual recognition of edges. And that was my first hint of how very simple operations could result in recognition. And then I did what I thought I was going to do anyway, which is I went to medical school. So, as an immigrant, you'd say that's a natural.

[00:11:41] All immigrants are told by their parents to go to medical school. That was not the case. What had happened was, I thought in my ambitious, naïve self in high school, I was going to get a Ph.D. in biology. And when I went with my father to visit colleges in the United States, much against his wishes, I ran into someone.

[00:11:59] [00:12:00] It was just a two-minute meeting at Yale, and he told me, he asked me what I wanted to do. And I, like a naïve young man, I said, I want to get a Ph.D. in biology. He said, don't do it! The doctors will treat you like dirt. Get an M.D. Otherwise, they're going to get much more resources than you. You're not going to be happy.

[00:12:17] Get the M.D. And I turned around and told my dad, and he said, makes a lot of sense to me. And so that made me go to medical school. And, but I was the first in my family to go into medical school. So, I arrive in medical

school, and at first I'm just overwhelmed by the amount of data I'm being force fed.

[00:12:36] And then I realize – So you're in Boston at this point – In Boston at this point. And I tell myself, uh oh. This is not science. I'm in deep trouble. I've always wanted to be a scientist, and instead I'm in a noble trade. What am I going to do? I started thrashing, and looking around, and fortunately, I had some [00:13:00] very nice mentors.

[00:13:01] And mentorship for me is a very important theme. And they said, go meet this person, Rob Friedman at BU. And he said, what about this guy, Pete Szolovits, at MIT? And I met Pete Szolovits, who was and is a professor of computer science at MIT. Then he was fresh out of Caltech and had just inherited a clinical decision-making group.

[00:13:20] And he was taking very seriously the operations of diagnostics in medicine. And all of a sudden, I realized this was going to be the science I was involved in. And I was thrilled. And that's really where I got bit hard by the AI bug. And frankly, a lot of lessons that everybody's learning now about AI, I learned back then.

[00:13:48] So even though we had, as it's called, an AI winter, a huge and well-earned disappointment, a lot of the fundamental lessons about decision making, [00:14:00] utilities, probabilities, bias, were all there in, um, the Ph.D. work I did back in the day. One of my favorite pastimes is to go back and read papers from medical AI of that period because in many ways, obviously we're using different methods, but the messages are still fresh.

[00:14:18] You will see in many papers, like the promise of computer-aided decision support has been long promised but has never been delivered on. What do you think about what's similar now? Like, what's different? The era of expert AI has always been so fascinating to me, but like, what's different now and what's the same?

[00:14:36] So, it was very strange being a graduate student in the 1980s because you'd hear some luminaries like Marvin Minsky and they'd paint verbally huge edifices. Describing how AI was going to emerge. And it was like religion. But I never understood how it was actually going to happen. And it may be that none of us were smart enough to fully understand his [00:15:00] vision, or maybe it was a little bit too abstract.

[00:15:04] Like society of mind stuff, like, yeah. Exactly. Society of mind stuff. And so, I mean, he did some very concrete stuff in the day, including slowing down neural network advances for a few years by his XR result. But now we're in a very different realm, and that's because, frankly, of just empirical success.

[00:15:26] And not just empirical success in the hands of a few experts, but the fact that we actually have these tools. Let's get concrete. The fact that a mom enters the history of her child, who has been having endless headaches, trouble walking, trouble chewing. No doctor is able to tell her what's going on. And out of desperation, she types in all the results of the reports and of the histories and physicals into [00:16:00] GPT-4.

[00:16:00] And she's given a diagnosis, which she then goes to a neurosurgeon, gives him the imaging studies and says, GPT-4 thinks it's a tethered core syndrome. And he looks at the image, looks at the child and says, yep, that's what it is. So, this is, we're no longer in the theory or promise. It's the present.

[00:16:20] That's why I'm so much more confident. I think that model, so that example that you just highlighted, also indicates an interesting model for the safe or responsible use of AI in medicine, right? Which is a way to prompt a new type of conversation with the patient and the doctor. Um, I'm curious, so going, you know, if we can dig a little more into your Ph.D. with Pete Szolovits in the 80s, um, can you compare, and maybe contrast, the intellectual climate of MIT in the 80s with today? With how a Ph.D.

[00:16:54] student would approach this today. Is there, do they have the same intellectual freedom to pursue [00:17:00] AI in medicine today that they, that you did when you were in, you were doing your Ph.D. So, clearly, you're trying to get me canceled. Um, so. Intellectual climate. The realm of intellectuality is broad.

[00:17:15] It's broad or can be defined in ways that are limiting. And so, um, I can't speak for all of MIT, but certainly in Pete's group, you were free to do whatever you wanted. He never told you what project to work on. For the better or worse, by the way, that's become my modus operandi as a mentor as well. But we were free to think thoughts like, you know, I was thinking about the fact that I was seeing that history, temporal history, and the partial order of events in a history were just as important as the actual events.

[00:17:51] In other words, to diagnose someone when they first have joined us before the transfusion, as opposed to joined us immediately after the transfusion, as opposed to [00:18:00] joined us 45 days after the transfusion.

Could be the difference between irrelevant to transfusion reaction to hepatitis. And I decided I was going to focus on representing temporal histories and alternative temporal worlds.

[00:18:15] My next-door office mate was working on qualitative physics and how to model physiological relations using the qualitative version of differential equations. These were very, very broad topics. But you can see both of that, that range still happening today, right? In Pete's group, with the right mentor.

[00:18:34] With the right mentor, you can. Yes, you can. But I think the big difference between MIT then and MIT today is MIT today is kinder and gentler, both to its benefit and to its detriment. It's still less kind and gentle than Stanford, and I'm told this by graduate students who visit Stanford, [00:19:00] but back in the day when you presented in Pete's group, there was always someone who was just very eager to take you down.

[00:19:07] And that turned out to be incredibly useful in mental hygiene because it saved you months of sterile, irrelevant explorations. It was a bit egotistonic at the time, but incredibly helpful. So, I think that the notion that for your Ph.D. you get to explore an entire world and find out where you want to make your particular contribution is a huge gift that's very rarely given to us.

[00:19:36] I think we should be doing more of it today. I also wonder too, like outside of like institutional climates, if the effect of large language models has reduced intellectual diversity of like what we're working on. Because it's hard to imagine grad students sitting next to each other now. Like in generative AI, there's this concept of mode collapse, and it's when the influence of the prior is too strong, and everything just collapses on like one thing.

[00:19:59] Do we [00:20:00] have societal mode collapse now? I think we have intellectual mode collapse within AI, and like just hearing you talk about like the different range of things that were going on in the 90s. I mean, obviously, like large language models are having their day and doing a lot and we'll cover that. But I wonder like what we're missing and the tales of the distribution since we're all focused on the mode right now.

[00:20:17] So, it's interesting. There's both more science and less science, both then and now. So, every one of us had theories, but they were not well-grounded in empiricism about what kinds of methods would work better for the diagnostic problem. And these theories were more intuitive and based on cognitive psychology and perhaps pragmatic insights than anything else.

[00:20:49] But they were fundamentally different methodologies. Now, at the base layer, these large language models fit a very well-oiled apparatus. But at [00:21:00] the top level, we're treating them like psychology. Most of the papers that you read, even in very technical conferences, are about different effects on the performance and behaviors

[00:21:13] of these models. And very few of them really try to dissect out how specifics of the model are actually related to that. It's like behavioral machine psychology. Yes. Robot psychology. Yes. I've had a discussion with Raj. If you look at the iRobot book collection by Asimov, it's about the three laws of robotics, but most of the stories are apparent breaks of the law, and the whole story is how, in fact, they're not breaks of the law.

[00:21:44] So, what logical pretzel twist got the robot to do what it takes? But the person who is there throughout, there's multiple robots, the person who's out there in every single story is a woman, Susan Calvin, who's a robot [00:22:00] psychologist. Uh, this, continuing with the theme of

your career, Raj, do you want to go on to the DBMI portion of the questions? Um, I think it makes more sense. We'll come back to that. Can I interrupt as I usually do? Yes, go for it. There's another thing that really motivated me about AI and medicine. It's when I came back to medicine and I could not believe how primitive it was and how it was a telephone game, someone told someone, someone told someone, and we were, you know, working off of written notes and hearsay and how this is all being turned into practice on very sick patients.

[00:22:29] And I felt very much like there's a Star Trek movie in which McCoy, they have to go back to get the whales back to the future. It's a long story, but on the way, they go back in time and McCoy, the doctor is going across the hospital wards and he sees a patient and he's saying, why are you here, dear?

[00:22:50] He says, I'm in dialysis. He says, dialysis? That's incredibly primitive, barbarian, and he gives her a pill instead. And I [00:23:00] felt that we were in that primitive state, and that we were not learning, and that we were not using optimal decisions. And so that was incredibly motivating. And let me tell you that anybody who does not get calloused, if you stay in medicine long enough, you realize so many suboptimal decisions are being made, uh, both at the data acquisition, the decision theoretic level, and then at the, uh, action level.

[00:23:26] And so that has become incredibly motivating for me, and it's always, I think, a slight, a hint of somewhere between urgency and a teensy bit of anger that we're not doing better by our patients. I'll say like, one of the great intellectual revelations of my life was when I did a postdoc with you.

[00:23:45] So, I was dating a med student, married a med student, and I was having the exact same thoughts that you did. We're like, these problems have been solved for 30 or 40 years in computer science and other areas, but I would talk to med students and medical people, and they would look at me like I was crazy.

[00:23:59] And [00:24:00] then I came up here for a postdoc and I was like, oh, okay, there are actually, these are my people. Like, there are actually other people who are thinking the same thoughts and trying to solve these problems. I think that's a credit to you that you've established this like community of folks who want to improve medicine through quantitative methods and there's like a real Zakosphere of people who think like that.

[00:24:19] Well, thanks for the compliment. But that's, in the end, why I created my first lab, why I created a center, why I created a department. I needed to have my people around me. I had a community. I needed a community around us. So, when you go into a fancy academic health care system and you want to have an academic career in the 1990s, probably what you want to be working on is on gene knockouts.

[00:24:43] I had zero interest in working on gene knockouts. But, there was no great example of career advancing in an academic mode, so I wanted to create that community that would allow it to be not an exception, not a psychiatric outlier. An [00:25:00] archetype. Yeah. And it would create the archetype of advancing.

[00:25:04] Even without an M.D. So, I think that's a perfect transition to our next topic, but I want to ask one more question about your career trajectory before that. So, I'm jumping ahead a little bit because we're going to ask you this as a sort of concluding question as well. Mm hmm. But, you did an M.D., you did a Ph.D., we've talked a lot with our previous guests about the training journey and the skill set that they need to have impact in medical machine learning or in an inherently interdisciplinary discipline like biomedical informatics.

[00:25:35] And so Ziad Obermeyer in particular articulated an extremely, I think, cogent case for both skills needing to coexist within the same mind. So,

the technical and the domain expertise are the medical skills. This is as opposed to what seems like a tempting narrative or tempting approach, which is to take a deep-clinical expertise and pair them with an AI researcher who doesn't know much about medicine.

[00:25:58] And the problem as Ziad pointed [00:26:00] out was that they often don't have much to talk about. They don't speak the same language. They can't see the connections. They can't whittle away. And they're like effectively lobotomized. They're lobotomized. They can't eliminate the bad ideas, right? What is good research?

[00:26:10] It's sort of eliminating lots of bad ideas or inferior ideas, inferior approaches to the same question. So, you have an M.D., you have a Ph.D. You were the first person, I think, to convince my mother that I don't need to go to medical school on my defense day of all days for the Ph.D. Do you believe, so I know that one of the answers, which is credentials, you know, you don't need to have both degrees to have impact.

[00:26:35] But do you believe that both of these skills need to coexist within the same person to have impact in biomedical informatics and in medical AI or medical machine learning? So, the short answer is yes. But you don't have to get it professionally. Peter Szolovits used to have in his group a challenge that he had done himself, which is read the entire Harrison's textbook of medicine.

[00:26:57] And I do think the thousand pages. The [00:27:00] thousand pages. Yeah. And he did it seriously. And I can have with Pete as good a medical conversation as I can have with any doctor. And, so, I think that's what it takes. It's not necessary, but it's much better. And it's very simple from the architectural perspective, having a relatively fast interchange in your own head than between two individuals

[00:27:27] makes a difference. You know, we're talking in minutes. It's latency. It's back to latency. And by the way, for me, latency is a huge issue across all of science. Can you just say your mantra here about a knowledge processing discipline? Yes. So, I do think that medicine is, in fact, fundamentally a knowledge processing discipline.

[00:27:50] And I think the understanding and misunderstanding of that relates to both a lot of the opportunity that we see here today and a lot of the despair, [00:28:00] dysfunction, and resentment that we are witnessing today. And I do think that the latency, even between individuals, as it's encompassed in institutions, where we have to get a research project going, there are so many

hop, skips, and jumps that have to happen, that shortening them by a factor of 10 would be tremendous.

[00:28:29] There's a story, which I'm embarrassed to say I wrote in this book that I wrote with Peter Lee and Kerry Goldberg. But Peter Lee was smart enough, smarter than I, he actually cites it. And I was never citing it until I saw him citing my story and I said, I should be citing my story. And here's the story.

[00:28:47] My first day as a doctor. I go into the basement of the Brigham Women's Hospital, I'm in the NICU, the Newborn Intensive Care Unit. And I never, I didn't know I was going into pediatrics until the last moment for a [00:29:00] variety of reasons. And this is the first time I ever saw these little translucent creatures.

[00:29:06] And that was scary. They were all in vents, and these kids are so immature that if you don't, they sometimes forget to breathe. And then the crazy thing is you flick their toe. You flick their toe, which is like, you know, I went to medical school and now the nurse is telling me flick their toe. I said, what the hell is this?

[00:29:23] A little shake. But then I get my first patient admitted. A newborn full-term baby, big, chunky, a baby, but his lung had collapsed. And what had happened it was pneumothorax, and they quickly stuck a needle in so it reinflated, but it convinced his blood circulation in his lungs that it was back in the womb because the oxygen saturation went down.

[00:29:48] So, the resistance in his lungs are very high. And so even though we were giving a lot of oxygen via bag, it was not getting into his blood, cause the bloods. Was [00:30:00] trying to go to the placenta, which no longer was there. It was not going through the lungs. And, so, I was hand bagging this child for about 12 hours and then I gave back a dead baby to the parents at the end of it.

[00:30:15] My then girlfriend picked me up after my, this first night on call, and I cried. And it was the only time I ever cried in residency. And it was not just out of sadness, it was out of frustration and feeling of uselessness. This was compounded when I found out that next door, two months later, a therapy was approved called Extracorporeal Membrane Oxygenation, ECMO.

[00:30:41] They had just finished the trials to show that for this disease, it worked. So, if that trial had just completed two months earlier, that baby would have lived. Or if you could have enrolled them in the trial. In the trial. Yeah.

Actually, it turned out the trial had been stopped and in fact [00:31:00] Because of dirty, dirty, dirty laundry.

[00:31:02] Yeah. It's, ECMO is expensive. The hospital was trying to figure out, before insurance funds it, are they going to do it? Right. Because you need to have two techs on it, and it's, so it taught me a lot about medicine, but, and I frankly, I was in fear of running into the parents. Because I was, I didn't know what to, if I, what would I tell them if we had that conversation.

[00:31:27] But that told me, a weak delay on the IRB. A piece of equipment that was not debugged correctly. Getting the right respiratory tech. All those delays. If you eliminated them, that baby might be alive today. We tend to focus only on diagnosis, right? Right. On that sort of, the encounter itself, getting the right diagnosis.

[00:31:48] But not on latency, on all the sort of slowness and the hurdles and the delays. Lower the coefficient of diffusion. Yeah, and that's where it's at. And it's the asymmetry too, from institution-to-institution or room-to-room. Yeah. Where [00:32:00] the knowledge is not flowing or being processed. And I can tell you, thousands of Americans die every year because of that.

[00:32:06] Yeah. That would not have to die otherwise. And, so, that is heartbreaking. Yeah. The future is here, it's just not evenly distributed. It's not evenly distributed and it could be accelerated. And I am an unapologetic accelerationist, before the term ever was conceived by certain crypto traders, um, because of the understanding that we already had a lot of the good ideas, but they have to be tested, and we have to be able to engage in science all the time in the practice of medicine.

[00:32:43] It's not, I don't think it's understood deeply enough. So, I think that's also a good transition point. So, Zak, I realize this is an amusing question for me to be asking you as a member of your faculty, one of your faculty members. But Zak Kohane, why did you start a department at Harvard focused on biomedical [00:33:00] informatics?

[00:33:00] What is your vision for the department? So, given that we agreed that I have been repeating endlessly that medicine is a knowledge processing discipline, I asked myself, how will I get a community that not only understands that, but is able to implement the experiments that will be convincing to industry to make this happen?

[00:33:25] And creating a critical mass of such individuals at Harvard seemed important because by the time I was in position to have a department, Boston very clearly was the epicenter of a lot of activities: pharma, biotech, hospitals, academia. And it seemed like if we could demonstrate that critical mass here, it'd be a great model for other academic centers.

[00:33:57] And so, once [00:34:00] I'd convinced the dean, then Dean Flier, that this made sense. Things went relatively rapidly, and it turns out that every single one of the faculty that we've recruited is a star, even the ones like you, Andy, who we did not recruit, or we tried to recruit, is a star and are making, huge inroads at this very special time where you do have to understand not only that medicine is an information processing business, enterprise, but what works and what does not work.

[00:34:36] What are the limitations? And I remember well the various conversations that we've had. For example, right when IBM, Watson started going to medicine, we were all very impressed with its performance in Jeopardy, but when they started telling us about medicine, maybe we spent a month or two saying, what secret sauce did they have?

[00:34:55] And we said, this is not real. This is not right. [00:35:00] It took us several years, but I think we have a very good bead on what is real by virtue of addressing the challenge, not as a purely academic challenge, but how are we going to transform medicine? And so that is a research mission, but it's also an educational mission, right?

[00:35:18] And, so, you recently announced a new Ph.D. program in artificial intelligence and medicine at the department. Could you tell us a little bit about that Ph.D. program and what you're hoping to achieve? So fortunately, because of the success of these large language models, there's even a further increase in the number of qualified undergraduates who have machine learning skills.

[00:35:43] What they don't have is appreciation of what the challenges are in medicine, what is medicine, and the practicalities of application of these various technologies to medicine. And, so, what we hope to do in this Ph.D. program, [00:36:00] which we just got our first set of applications, and we'll be matriculating our first group in the fall of 2024, what we aim to do in this program is very much modeled on a much older program that was started between MIT and Harvard called HST, where we take very strong, quantitatively trained individuals.

[00:36:20] Ph.D.s in math, computer science, physics, mechanical engineering, and then immerse them in medicine, medical applications, while deepening their methodological knowledge. And, so, the goal is just to create, I hope this doesn't sound overly arrogant, a new set of leaders who can actually lead us in this new world where machines are going to be at our side in medical decision making.

[00:36:49] So, we could spend literally 10 hours talking about all the different things that you've created. But I, I now want to talk about something that we haven't had actually talked about on the podcast yet, which is *NEJM AI*. So, we're all sitting [00:37:00] here enjoying some very tasty scotch, courtesy of Zak. And if you roll the clock back, cheers, two years ago, we were sitting in Zak's backyard, also drinking some tasty scotch on a very cold winter night in Boston, talking about creating a new journal called *NEJM AI*.

[00:37:14] So, could you talk to us about the genesis of the journal, what it is, what its mandate is, and why we need it? So, it's an odd thing for me to be even talking about *NEJM AI*. If you talked to me as a grad student that there would be such a thing, I would think it was an elaborate prank and, but I'll give credit where it's due.

[00:37:36] The current editor in chief of the *New England Journal of Medicine* and the prior editor-in-chief, Eric Rubin and Jeff Drazen, thought this was a good idea about five or six years ago. And then they approached me two years ago about it, and I told them absolutely not. There was just not enough good articles.

[00:37:57] There were a lot of great AI articles, [00:38:00] but not enough good articles that would interest a clinical audience. So, then they approached me two years later. Of course, it was deeper now in the use of convolutional neural networks, particularly for image recognition. So, I felt there was going to be enough interest.

[00:38:13] And then, using my other favorite adage, better lucky than smart, large language models erupted on the scene. Not when they were first written up in 2017, but when they started being released to the public in 2022. And all of a sudden this seemed like perhaps the most momentous occasion in medicine for many decades.

[00:38:36] And, so, I was very lucky to be able to then recruit you two guys to be among our deputy editors and an amazing group of members of the editorial board. But what's our mission? Our mission, first and foremost, is something

very practical, which is which of these widgets is clinical grade? It turns out there's no there's a [00:39:00] lot of authority out there that's going to tell us that this AI program is better than the other AI program.

[00:39:08] There are some FDA processes which don't generalize well to multiple hospitals. Or don't specialize well to multiple hospitals that don't work well with the way medicine changes over time or with the way practice is varied by location. It doesn't involve comparing one to another. So, there really is a huge gap between this promising technology that we're already all using and measures of clinical quality.

[00:39:40] And so that, and its embodiment in clinical trials is the bread and butter of the journal. But in order to make it a little bit more educational and palatable, in addition to this, we're going to have, and we already are publishing in our advanced online issues, [00:40:00] perspectives, case histories, and very importantly, because data is so important to the performance of these models, benchmark data sets.

[00:40:11] And, so, we want to both advance the field by evaluating these programs, but also advance the field by knowledgeable commentary, knowledgeable reviews, but also creating benchmark data sets and benchmark questions around those data sets to allow others to move the field even faster. Yeah, I think that's awesome.

[00:40:33] And when you like approached me and Raj about this two years ago, like it was like, I could not have said yes faster. It's like such an obvious thing to like want to do. And I'm super excited, how the journal has grown, I guess, so that Charlotte doesn't get mad at us. Yes. Say that you're an author out there.

[00:40:49] Yes. And you have a great AI paper. Yes. Why should they submit it to *NEJM AI* over a general interest journal or some, or like a medical specialty journal? So, it's like [00:41:00] the old army commercial. We're going to make your paper be the best it can be. Because we really care, in the spirit of our training, that these articles be relevant and safe and useful for our readership.

[00:41:17] And that means we're going to put a lot of effort in working with you, the author, to making sure it's maximally relevant. And so, yes, we love, as much as anybody else, a splashy finding. But we're going to make sure that the splashy finding is, first of all, has substance. But also, that we're clear about its limitations, and we're gonna make sure that the message in the paper is neither smaller or larger than what the actual data and analysis can show.

[00:41:50] We really want this to be usable by patients and doctors and technologists as a reliable measure of, [00:42:00] frankly, the real thing in AI in medicine. I think that's a good answer. I think, I think we're ready for the lightning round. What about you, Raj? Let's do it. Okay.

[00:42:12] Okay, so just, this is a little bit of a Zak deep cut, just because I know you. But lightning round question number one. What is your favorite Radiohead song? Oh my God. I'm gonna have a meltdown. Okay, uh. What's the one, uh. The end credits one is one that I hear you, songs, or it's the one that comes at the end of *OK Computer*.

[00:42:34] Yeah. I'm having a senior moment. You bastards. That's a good one. I am truly in awe. That's good. "Lucky"? "Paranoid Android"? "Paranoid Android." "Karma Police." "Karma Police" is one. "India Rubber," "The Bends." Actually, because of this, because of how well it got rendered by… okay. Hang on. Radiohead, [00:43:00] *Social Network*.

[00:43:01] Oh. Creep? In social, the, the Yes. In the movie? Yeah. Well, it's like a different version. It's like a, it's a different Facebook movie? It's like a female version? It's not female. It's um Okay. It's like Quarrel or something. Yeah, it's a Quarrel version. Let's, let's, let's get this right. Yeah. Radio… Which is funny, because for like Radiohead nerds, "Creep" is like, not the one.

[00:43:23] I know, I know, I know. But I, but I love this cover. Cover. Um. Can we sample that in the podcast episode? Yeah, we can be the lead in. Alright, we got the thumbs up. We have Zak Kohane and AI Grand Rounds. I'm a creep. Okay. Yes. So my favorite, it's not my favorite Radiohead song, but it's my favorite song because of the cover by Scala and the Kolachny brothers.

[00:43:51] It's the version of "Creep" that appears in *The Social Network*. Excellent. That is a deep Radiohead, uh, cut. So, that's a good answer. [00:44:00] Zak, if you weren't in medicine, what would you be doing? Writing science fiction. Well, that's a very, okay, so that tees us up nicely. Um, so the next lightning round question is, what is your favorite piece of science fiction book or movie about AI?

[00:44:16] About AI. Yeah, *2001: A Space Odyssey*. Yeah, that's a good one. Yeah. Yeah. Alright. The second best one is, a short story by Asimov where this computer is asked to solve a bunch of things. The last answer, or the last question? Yeah. Yeah, that's a good one. I'm amused that I get to ask this question today. How much time do you spend on Twitter/X a day, and true or false, did your daughter buy you a mug that commemorates your addiction?

[00:44:46] I recently been trying my best to cut down on Twitter and I'm ratted out by my iPhone and I truly believe I have cut down. It's still 1 hour 30 minutes a day. [00:45:00] But I would not be surprised if before my lean diet, I was doing 3 hours a day. Alright, so 50% improvement, good. And yes, my daughter did buy me

[00:45:11] "You are a Twitter addict," and it's getting to the point that my friends are beginning to try to shame me as Kerry Goldberg just did on Twitter, X, to say that I'm spending too much time there. This is one of our highlights. I will say the most reliable way to get an answer from you is to send you a DM on Twitter.

[00:45:27] Actually, I think we sent you the prep questions for this on Twitter DMs. Yes, yes you did. You're about to get spammed from Twitter DMs. You better close the DMs. This is, here is the, a good AI comment, but in the end, it's just a, an excuse. When ChatGPT came on the scene, LLM progress became so rapid that I truly believe the only way to stay up-to-date,

[00:45:56] even with casual conversation with other computer scientists, was to look at the [00:46:00] preprints being referenced. I still believe that. Yes, so that's convenient. I wish I could say that all my time on Twitter was spent looking at which are the right preprints to look up from Twitter. I'll say it is right now the fastest way to get access.

[00:46:17] So, our last two lightning round questions are an attempt to get you into a little bit of trouble. Okay, go ahead. So, this one I'm ad libbing. Yeah. Um, but in 20 years from now, will *NEJM AI* have a higher impact factor than the *New England Journal of Medicine* itself? There will be no *New England Journal of Medicine*.

[00:46:33] *NEJM AI* will have subsumed *New England Journal of Medicine*. It's going full board. Alright, our last question, Zak. Who will win the 2024 U.S. presidential election? Oh my God. And do you have any bets attached to this? Yes, so. With your deputy editors. Yes, so I had bet, and I really regret this bet. Because it was way too optimistic.

[00:46:57] I had bet with my deputy editors. [00:47:00] That Trump would not win the primary. So, unless he decides to go in as a third party candidate or something else happens. This was in 2016? This was in 2016? Yep. No, the Trump's first presidency. It was 2016. No, no. But bet against. We lost. We bet. On the other side. Right, right.

[00:47:21] I'm talking about the bet right now. Right now. That bothers me. Yeah. Yes. The bet that bothers me right now is that I took on a bet with you after Joe Biden was elected that Trump would not be the nominee of the Republican Party. It seemed to me at the time a reasonable bet. At this point, it no longer seems like a good bet unless either he runs for third party or something else happens.

[00:47:53] Or Sam Bankman Fried pays him 5 billion dollars not to run or something crazy like that. I don't think, well Sam Bankman Fried, [00:48:00] can I just point out, yet again, Another way to pronounce Bankman Fried's last name is Bankman fried. Yeah. Nomen est omen. Nomen est omen. Can you tell our listeners what that means?

[00:48:12] That means, it's a Latin saying, which is, Your name is your destiny. And, yes, I don't think the 5 billion from Sam would do it. I think there's an issue of pride. But at the same time, I truly do not know who's going to win the election. Because I think that the number, it's like a truly chaotic process, both in the colloquial sense, but in the technical sense, because there's so many initial conditions that we don't know of that may affect that process.

[00:48:45] It could be Trump. It could be Biden. It could be someone that we didn't see coming. Fantastic. All right. Well, you survived the lightning round. Cheers. Congrats on surviving the lightning round. And [00:49:00] congratulations on a great year of podcasts. I was joking about it at the beginning, but I've gotten so many compliments, around such that I've almost convinced myself that these podcasts were my idea.

[00:49:14] Congratulations. The ultimate, the ultimate compliment. You deserve a lot of the credit because you put me and Raj next to each other in cubicles for our postdocs. Well, we were eventually separated because we were having so much fun and there were multiple noise complaints. And so actually we've been training for this for a long time, so you actually do deserve indirect credit.

[00:49:32] I think we spent a day of our postdoc as postdocs through discussing the Mark Cuban debate with physicians. Yeah. And so, it was perfect trading for, uh, for this, Zak. So, thank you for creating the setting for Disrupting Countway Library and training for this. And so, if you want us to mention you in next year's, we'll have next year, we'll do this again in a year.

[00:49:57] We're going to pick out two [00:50:00] articles that we want to commemorate after a year. So, here's my challenge. Whether it's through

drama, science, excellence, or impertinence, submit that article that will get you mentioned in the December episode, in the December of episode of 2024. All right. So, Zak, we just have a couple of big picture concluding questions for you.

[00:50:26] I'm going to change this one that I was going to ask because I think you addressed it in your career arc already. So instead, I'm going to ask about your mentorship philosophy. So, I have to say, I was at this meeting that you organized in Maine. The RAISE meeting, very important, excellent group of individuals that you recruited to the event to discuss safe and responsible use of AI in health care and medicine.

[00:50:52] And I was struck by the agenda, the level of discourse, the investment that everyone had in the [00:51:00] goals of the meeting. But I was almost just as struck by the fact that you had some of your yeah, I actually think your first trainee, Dan Nygren, uh, and your most recent trainees, and me and Andy, somewhere in the middle there.

[00:51:15] At this event, and we were all connecting with each other, we were all speaking, we were all on A, excellent terms with you, and B, we were all sharing our positive experiences with our time in your lab and in your orbit. And I think it speaks a lot that you're still at 2020. This was a couple of weeks ago that you're still connected to all of your mentees.

[00:51:38] I hope you can maybe distill for especially our new PIs who are creating a lab culture. So, I think still includes me and Andy. We're still building our labs, developing a culture. We've tried to, I think, emulate aspects of how you run your lab. Can you give us your philosophy for mentoring, how you approach running your lab and

[00:51:58] nurturing [00:52:00] scientists at very critical transitions in their Ph.D. and in their postdoc, their clinical training at these points in their career. Well, I think there's one aspect which you take for granted, but I assure you, you don't. You two already have it. But it's actually caring about your mentees. Turns out it's not that common.

[00:52:24] And I think it's something that you pass on, just like child abuse is passed on generationally. I think good mentorship and bad mentorship is passed on generationally as well. And I think it doesn't take that much caring. Just ask yourself, is this person heading in the direction that's going to make them happy?

[00:52:43] And you don't have to think about it in extremely altruistic ways. Just ask yourself dispassionately. Just by doing that, you're making yourself three or four standard deviations better mentor than most people. So, if you have that edge, then the next comment becomes irrelevant, which is…
[00:53:00] That's the first order effect.

[00:53:01] That's the first order effect. The second order effect is let them explore. Yeah. And let them fail. Yeah. And let, absolutely let them fail. And, just shrug if they fail. Don't try to make them feel better about it. Don't make them feel bad about it. Just shrug. It's what happens. On to the next thing.

[00:53:18] Yep. Alright, so I think we probably have like 10 minutes left. So, I think this is the December episode, so I would like to look back at the year that was in *NEJM AI*. First, I would love your thoughts as editor-in-chief of any big trends that would be worth pointing out to our listeners. And, also maybe, I think, tease the editorial we have on the use of LLMs that will be coming out.

[00:53:41] I think that might be an interesting discussion point too. Right. So, I think, well, one trend that I just want to note, it's not a trend, it's because it's a flat line, it's essentially all our submissions involving large language models involved one model. [00:54:00] ChatGPT 3.5 or 4. I think, I hope, that's going to change.

[00:54:08] And so, therefore, there was not much in the way of a comparison of other modes. It was more like, what can this do? And so, there was a lot of, isn't it amazing? It's like the dog in the opera. It's not how well it sings. It's the fact that it can sing at all. And so, we were seeing GPT-4 doing a lot of tasks that frankly, most of us would have been skeptical that any AI program could do.

[00:54:36] And even if there were a lot of problems, as people had pointed out, the fact is it was doing it. In some semblance of something competent.

[00:54:45] In terms of trends, I think what's delightful, I think you just shared with me a clip where our associate editor, editorial staffer, Morgan Cheatham, was hosting [00:55:00] Daphne Koller. And she was commenting about the fact that we're on the steep exponential. And just to make it real for our readers, the fact that we went from a GPT that scored lower on the national medical boards than any human, to a GPT or a MedPalm that scored better than 90% of those taking the boards in under three years. I mean, you heard me in your office as a postdoc for years talking about trying to get an AI to pass USMLE. And I would bang my head against the wall and we couldn't get higher than 40% on actual like USMLE questions. Turns out that we were just a couple hundred of billion

parameters short. And the fact that this went from like completely intractable to now almost completely trivial.

[00:55:48] Like we take it for granted, right? Over a year, every day I wake up, I'm just, I'm shocked. So that's great. So that's what Daphne was thinking about exponential. She also said we don't know what the [00:56:00] exponent is. We don't know the exponent. And so, therefore, we should be extremely reluctant to make any prediction except to say A, whatever the problems that we have are not the problems that are going to be relevant next year.

[00:56:13] So, I think hallucination and up-to-dateness will be less relevant next year than they are now. But there'll be other problems. I think one of the bigger problems will be, what if this thing ends up saying things that are right and we don't know how to prove that it's right. I think the bigger problems, from my perspective, is how do we get our educational system and our medical system to keep up with this crazy approach increase in capacity?

[00:56:45] Because all of a sudden, we actually have the potential now, next year, to give superlative performance. And it'll become a legitimate question, just like my baby, my first patient baby, why are we [00:57:00] not giving that patient? Where it's uncertainty about what to do next. The benefit of a second opinion. And I think that question will not be a ten year, from that question, it will be a next year question.

[00:57:13] Yeah. We also think about equipoise a lot in medicine. And not using AI will not have equipoise in the very near term, it seems like. I think that's it, Zak. Are there any concluding comments that you would like to leave us with before we? Yes. I think that, you're talking about the future, and I think I referenced this before, which is, what a strange era we're living in.

[00:57:37] Doctors, more than ever, are discouraged and stressed, and yet we have this exciting technology. That's going to revolutionize medicine. And there are other technologies that are going to revolutionize medicine. But medical doctors have never been more stressed, depressed. I think a survey showed that young doctors in the U.K., somewhere like 50% were looking [00:58:00] for careers other than clinical.

[00:58:02] In the United States, I think it's about 30%. At Harvard, I think it may be closer to 50%. And it's a sign that the best and brightest don't feel that they can change medicine. In the standard process, and so I really do think this is a wakeup call for those of us who seek to be leaders to say, how can we take all these pieces, these new biotechnologies, new AI technologies, the current

broken ways of paying for health care and turn it into something very, very different.

[00:58:39] I think it's a real challenge, and you have to ask yourself, why is it that still today, the most recognizable sign of a doctor is the filthy stethoscope that they're hanging around their neck that's touched so many other people and carries the [00:59:00] virus from one person to another?

[00:59:01] Why is that the symbol, this very old acoustic technology, the symbol of medicine? Yep. Well, Zak, on behalf of the entire staff of *NEJM AI*, it's been a great year and we look forward to a productive 2024. Thanks for being on AI Grand Rounds. Hear, hear. Thanks, Zak. This was great.

[00:59:27]