AI Grand Round Podcast #4 03.27.23

The GPT-4 Episode: Microsoft's Peter Lee on the Future of Language Models in Medicine

[00:00:00] I have no doubt that ChatGPT powered by GPT4 will be used by doctors and nurses in clinical settings. Every day. I have no doubt that patients will be turning to ChatGPT for their own healthcare and medical advice. It's something that is just going to be happening, but it is not clear how best we should be doing this, and that there is a need for a public debate and that there is a need for the medical community to really understand what is happening here.

[00:00:40] And to understand it deeply enough to have the correct guidelines, to have the best shaping of regulatory issues and to understand even nuts and bolts things like what does it mean to validate this thing for clinical use? And it's urgent that we [00:01:00] get on top of this immediately because this is a technology that will be in every doctor's and nurses and patient's pocket.

[00:01:08] In fact, it is there today. That was Peter Lee of Microsoft predicting the omnipresence of new large language models like ChatGPT in medicine. I wanna welcome you to another episode of NEJM AI Grand Rounds. I'm Raj Manrai and I'm with my co-host Andy Beam. And Andy, I think it's fair to say that this is the GPT episode.

[00:01:29] We're really lucky to be able to speak with Peter Lee. Peter is corporate Vice President and head of Microsoft Research. This conversation both amazed and unnerved me. I lost some sleep after speaking with Peter. I don't think until we spoke with him that I appreciated just how rapid and fast this is all moving, how capable these models are, and how much large language models are making me question parts of my understanding of the world and what intelligence really is.

[00:01:56] Peter has this really unique perspective having come from [00:02:00] academia, government, and industry, and I think he just masterfully interweaves the best of these approaches to tackling problems in his thought process and in his leadership. I totally agree, Raj, and before we jump into the episode, it might be worth defining a couple key terms that we're gonna revisit in this discussion.

[00:02:17] And we hear a lot about this model called GPT4 or ChatGPT. And what the GPT stands for is Generative Pre-trained Transformer. And it's this idea that has emerged in AI over the last four or five years that actually trains a very, very simple model that you can think of as an auto complete model.

[00:02:36] So the model is given a lot of text, and the only thing the model has to learn how to do is to predict the next word. And the thing that has been so fascinating and the thing that we'll explore in this episode is these very seemingly simple auto complete models can actually solve a wide range of tasks.

[00:02:53] So when we're talking to Peter about all of the different things that GPT4 and ChatGPT can do, just keep in the [00:03:00] back of your mind that this is all a fancy auto complete model. Peter pushes back against that definition a little bit in the episode, but underneath the hood, it's an auto complete model.

[00:03:09] And what Peter does a wonderful job at doing is explaining how these models are gonna impact healthcare, how they're gonna impact the delivery of healthcare to patients, and how much medicine they know presently. I found the conversation extremely fascinating and I think left me both optimistic and slightly anxious about where we're going with AI and medicine.

[00:03:31] It was a truly fascinating conversation and I'm so excited that Peter took the time to share his valuable insight with us. The NEJM AI Grand Rounds podcast is sponsored by Microsoft and VizAI. We thank them for their support, and now we bring you our conversation with Peter Lee.

[00:03:54] All right, well, Peter Lee, welcome to AI Grand Rounds. We're so excited to have you here. Well, [00:04:00] thank you. It's great to be here. Peter, we wanted to start off with a question to get to know you a little bit better, and this is one that we ask all of our guests. Please forgive our machine learning puns in advance throughout the podcast because you tell us a little bit about the training procedure for your own neural network.

[00:04:15] Specifically, how did you get interested in AI? What data and experiences led you to where you are today? First off, don't be embarrassed about that. I'm just glad you didn't ask me how many parameters my neural night has. . There we go. Yeah, . My parents immigrated from Korea and got their PhDs and my dad became a physics professor and my mom became a chemistry professor.

[00:04:38] So I guess in that sense, I was growing up in an Asian household with two physical science parents. I was destined to become a scientist myself. The joke is that I was a huge disappointment to them when I decided to major in math in college. But I did get my math degree and then I became interested in computer science and was persuaded to [00:05:00] stay in college to finish my PhD in computer science.

[00:05:04] I was doing pretty mathematical theoretical work, something called denotational semantics and type theory, and I ended up getting my professorship at Carnegie Mellon University doing that kind of theoretical computer science work. Ultimately, after about 24 years there, I was the head of computer science, and then after Barack Obama was elected in 2008, I was shanghaied to serve.

[00:05:30] In the government at DARPA, that was a transformational experience for me. I learned a lot about different styles and different paradigms of computing research at DARPA. Finished my two year stint there and then was persuaded by Rick Rashid and Craig Mundy at Microsoft to abandon my professorship and I've been at Microsoft ever since.

[00:05:53] That's awesome. So it sounds like you were a theoretical computer scientist for a while and then you got [00:06:00] attracted towards applications and product development and a lot of these, I think surprisingly for me, you, you kind of learned at DARPA, that's not necessarily where I would think one theoreticians would go, but then you like kind of really think about practical real world applications.

[00:06:13] So I guess how does your theoretical background inform this turn to applications that you've had in the last part of your career? , you know, it's, I, I reflect on that a lot because in my gut, and maybe it's a genetic thing, there's a breed of researcher that I think fundamentally believes that ideas that are mathematically beautiful, that have this sort of incredible, I don't know, crystalline structure, that things that are that beautiful are also the most practical in the world.

[00:06:45] And sometimes it's just really hard to prove that in real life. But it, it does happen over and over again. And I think that computer science is a field that is just full of examples like that. And when you [00:07:00] hold onto that, it forces you to think a little bit differently. Because what happens in a big software company like Microsoft or you know, department of Defense Agency like DARPA is you do get kind of sucked into a lot of the complexities of the real world.

[00:07:14] And so to kind of force yourself to boil things down to their. Kind of simplest essential structure. There's a real joy in that. I mean, it can be frustrating. There can be a tough slog because you always feel like you're operating with a little bit with a chip on your shoulder. And there's, it's so easy to hear all the doubts and skepticism about things, but when you actually achieve something that is just so simple in essential structure that ends up being incredibly beautiful, it's just a wonderful thing.

[00:07:46] And so at at DARPA, one of the informative moments for that was one of the early projects that ran there, which ended up being a contest, it was called the DARPA Network Challenge. And this was a time where D.O.D. [00:08:00] leadership did not believe in social networks, in network effects, social media. It was all still very, very new.

[00:08:06] This was around 2010. And you know, I was tasked with proving that these might be important in national security in the future. And so with a group of soldiers, we. Staged a national contest. On the 40th anniversary of the arpanet, we took 10 big red weather balloons across the continental US in public, but undisclosed locations.

[00:08:31] And for eight hours on that day, we held them a loft at a hundred feet and said whoever confines the precise locations of these 10 balloons would win a $40,000 prize. And that just set off this gigantic hunt for these red weather balloons and a typical DARPA style. There were even teams of psychics, that were trying to compete.

[00:08:54] Um, but it ended up being a team from the MIT MediaLab under Sandy Pentland that [00:09:00] came up with a beautiful way to mobilize social media and to incentivize people to join their team and to search for things. And, you know, in much less than that. Eight hour a day, they managed to find all 10 of those balloons.

[00:09:18] And their way of mobilizing social networks made them able to exploit and harness the labor of people from all over the world. They had members of their team from China, from Iran, from Europe, from South America, uh, that were all kind of helping recruit more people, uh, to find these balloons. That project, as silly as it was, ended up forming the foundation of a program called Nexus 7 that ended up getting deployed in Afghanistan during that war in order to motivate incentivized people to provide the US military with information that might lead to the location of roadside bombs.

[00:09:58] Um, and so it, it ended up [00:10:00] being, again, this kind of beautiful idea outta the MIT MediaLab, uh, that wins a contest, but then also gets deployed, uh, in a real wartime situation. So if I understand correctly, the challenge of finding eight red weather balloons scattered at random across the place of the planet, the solution was actually in instead of mechanism design, that was what actually you needed to solve to solve this larger problem.

[00:10:23] And that was what the MIT MediaLab did was how do we recruit enough people versus how do we find cameras? How do we like, do all these like things ourselves? It's actually, we need the right incentive mechanism to get enough eyes on this. That's right. And you know the teams, there were thousands of teams that competed.

[00:10:39] There were teams that used cybersecurity exploitations in order to. tried to intercept people who were making balloon citing reports. There were Russian teams, uh, that were very, very good at that. There were people associated with the US Navy that were partnering satellite resources. [00:11:00] We also got to see quite a bit of diversity there.

[00:11:03] Uh, but importantly this ended up being really persuasive to the US military leadership. It was the social media and the network effect approaches that ended up performing the best. Give it enough eyes, all bugs are shallow. Yet, yet another demonstration of that. So, so I'd like to transition to a little bit now to your present day role at Microsoft.

[00:11:24] Could you tell us a little bit about what your role is generally and then types of things that you're directing at Microsoft? Sure. I'm the corporate Vice President for Research and Incubations. And as the title says, it's a job that has two halves. One half is research, and that is part of Microsoft called Microsoft Research.

[00:11:44] And so that is the fundamental research arm of the company. It's a fairly big place, about 1500 PhD researchers and engineers scattered in at this moment in nine research labs around the world, [00:12:00] and then incubations. What that means. That's sort of an internal term. Oftentimes, in fact, many, many times new ideas for products at Microsoft come out of research.

[00:12:11] And typically they can start, maybe even as a single person, a single researcher has an idea, starts working on something, it ends up looking interesting, and we invest more and more people and compute resources in it. At some point, it gets large enough. The company, and in particular our chief financial officer starts getting annoyed at the expense.

[00:12:32] And, uh, at that point we officially designate the effort as an incubation. And what that does is that it just applies a little bit of pressure for us to decide to do something to either shut it down or to actually make the firm commitment to move this into a real profit and loss business component of the company.

[00:12:53] And that's an incubation. And typically we put ourselves on the two year clock, uh, for those two things. So [00:13:00] my role is to lead Microsoft Research and to kind of shepherd this incubation process for the company. So this is a medical AI podcast, and you're a theoretical computer scientist who has DARPA experience and now a VP at Microsoft.

[00:13:14] Could you help us understand Microsoft's interest and role in healthcare and sort of how this all comes together? Yeah. And. I now consider myself pretty much full-time thinking about healthcare and medicine and AI's impact potentially there. Uh, and it wasn't the plan at all. I was very, very happy working at Microsoft in the research division.

[00:13:38] And then in 2016, I received a note from my boss who at that time was Harry Shum, our Chief Technology Officer at the time, uh, saying that I was going to be reassigned to work on healthcare. And I didn't understand that, uh, I really hadn't spent much time thinking about healthcare or medicine. And [00:14:00] I said no to my boss and in good management fashion, uh, Harry said, well, you should talk to my boss, Satya Nadella, the CEO.

[00:14:13] And, you know, that was basically an implied threat that you better do this, Peter or else. Um, and so I started receiving invitations for a one-on-one meeting with Satya Nadella to discuss this. And I was so distraught over this. I did not respond to the meeting invitations. Uh, and then was just trying to avoid it.

[00:14:32] And then in 2016, uh, that was an election year and my wife and I, my wife was getting fairly active politically and we went to a political fundraiser and unfortunately Satya Nadella was there, . And so he cornered me and I just was not, I didn't have enough spine to stand up to him. And he said, Peter, you're going to do this right.

[00:14:54] And I. Meekly said yes. And so that's how I got started in [00:15:00] 2016. I, we've had several guests on the podcast tell us how they got interested in healthcare. I think this is the first instance of conscription, , um, into healthcare. So, well, so, so now I had to figure out what to do here. And, you know, Microsoft has a very large business in the healthcare sector.

[00:15:19] We don't disclose publicly the, uh, amount of revenue in healthcare. But if you go to literally any healthcare clinic, hospital organization, anywhere in the world, uh, from a single nurse clinic in Nairobi to, you know, any part of UnitedHealthcare, uh, you'll see Microsoft technology in that clinic. And you know, there's a common experience.

[00:15:45] You just see windows, for example, everywhere, and it goes deeper than that. You know, there's also, you know, active directory and um, uh, Azure Cloud and so on. Um, but I think Satya Nadella was concerned that the future of healthcare and [00:16:00] technology, uh, was going to be much, much more than that. That data, uh, was going to be a big deal, and that healthcare organizations were getting more interested in data and, and using cloud technologies for that.

[00:16:12] And then machine learning and AI were also going to be, uh, very significant. And it's not only AI for kind of clinical records and note taking, but for imaging for genomics. and then that also extends into the biotech world and to the pharmaceutical industry. And so he really wanted a remake with that more kind of interesting technology future in mind.

[00:16:40] And that was the assignment. The question is, I felt like I was being thrown into the middle of Pacific Ocean and being asked to find land. And you know, you're just adrift and you don't know which way to go and everything is a big confusing mess. And then at that time, since we had such a [00:17:00] large business presence already, there were several very powerful vice presidents at Microsoft running businesses with big healthcare customers.

[00:17:09] And so there was also the question, how would I get them to listen to me in research? And so ultimately you end up asking the question: if a big tech company like Microsoft, uh, and our peers at places like Amazon or, or Google. If we were to disappear today, in what ways would the world of healthcare and medicine be harmed or held back?

[00:17:36] And when you start to think about that question, then you start to realize that the cloud will end up being a very important foundation of data infrastructure and compute, uh, for healthcare over the next 10 years. When you start realizing more precisely some of the challenges and health data [00:18:00] interoperability that have been just real frustrating sources of expense and waste.

[00:18:06] You confront the horrible burdens of clinical documentation and administrative and clerical work, uh, that are causing burnout and on and on. And so even before you get to wizzy AI technologies, there are a lot of things that sort of, uh, get you very focused. And so it ended up developing into a very important strategy for Microsoft where we actually even partnered with Google, with Amazon, with IBM, with Salesforce and others in order to try to move the technology forward to give us a chance over the next 10 years to be in a better place than we are today.

[00:18:50] And then we also tried hard to compete with each other to, you know, realize, uh, new futures involving AI. Yeah, just one follow up there. [00:19:00] So you are sort of thrust to this position, and I think that your experience is what I imagine a decent number of our listeners go through where they have a computer science background and they know that there's this opportunity in healthcare, but it's so opaque and so difficult to find a foothold.

[00:19:18] Could you tell us a little bit about your sort of learning journey and how you've charted out what this landscape looks like to sort of become an expert? Yeah. You know, so first off, there is this irrational exuberance. And what I think happens there, at least I was guilty of this as a techie. You look at a lot of the things that go on in healthcare, well first off, everyone has contact with healthcare.

[00:19:42] So everyone has opinions. My wife has an opinion, my son has an opinion. Uh, all of my colleagues and coworkers have opinions and you know, you look at some of the problems in healthcare and it just seems that they're all solvable through technology. Or already solved. Or already solved. Yes. [00:20:00] And then coming from a place like Microsoft and Microsoft Research, when I interact with leaders, administrative leaders, let's say in big health systems with doctors, with nurses on the front lines, they look at me and my colleagues in the tech industry, and I think they're dazzled with the magic of all of the technology we bring and all the machine learning and AI.

[00:20:24] And so coming from both directions, there's just this irrational optimism that I can bring all the tech magic to healthcare and gee, all the healthcare, these problems can be easily solved. And that irrational exuberance ends up being a source of failure over and over again. You know, it's one of these situations where you end up starting to just see the enormity of all of the conflicting [00:21:00] ideas, the the gaps, the perverse incentives financially, uh, throughout healthcare system.

[00:21:06] And each one of these little elements is in the. Rational and arose outta some common sense, but when put altogether, uh, we end up seeing something that, uh, just doesn't work that well. Yeah, local optimality does not imply global optimality and you see a lot of local optimization things happening in healthcare that lead to overall a globally suboptimal system.

[00:21:31] That's right. And so I think for me and the strategy at Microsoft that we've tried to pursue, we've really tried to flush out of our minds any sense of making major disruptions. We don't talk about disruptions. We even are uncomfortable with using the tech phrase transformation. You know, what we're really trying to do is just understand today's workflow.[00:22:00]

[00:22:00] You know what people who are actually delivering healthcare. on a day-to-day basis, or going through and understanding what are the ways that we can just improve the day-to-day effectiveness and satisfaction of those people. And then if technologies like artificial intelligence can help connect that as seamlessly as possible all the way back, not only through all of the kind of business aspects of healthcare, but all the way back to fundamental research in medicine, then we'll have accomplished something.

[00:22:35] And I think that doesn't mean that trying to go for a complete disruption in some segments of healthcare isn't a useful thing to do. It can be a useful thing, but we've decided to pick our roles and really just try to integrate ourselves as much as. Yeah. So Peter, I wanna switch gears a little bit. I promise we're gonna jump to, uh, large language models [00:23:00] and conversational agents very soon.

[00:23:02] But I hope we can take just a brief interlude, maybe broaden out a little bit. You know, your experience is truly broad spanned academia industry. DARPA, your professor for many years at Carnegie Mellon University. You're a chair of the storied Computer Science Department there. Now you're at Microsoft Research.

[00:23:19] I'm curious if you can tell us, you know, a little bit about how your experience in academia has maybe informed your management style at Microsoft. You know, I've been very lucky to have these different experiences, but one thing I'll say about my time in academia is in terms of innovation organizations, academia is the one place that is organized correctly.

[00:23:43] And by that I mean that socially and culturally, universities are the only place where the most important people, socially, are the creatives, the students and the professors. Uh, and the least [00:24:00] important and least interesting people are the senior management, you know, the deans and the department heads and the presidents and so on.

[00:24:06] So, uh, when you go to the military like DARPA or you go to a corporate environment like Microsoft, it's upside down. You know, the joke there is in my current role as the corporate vice president of Microsoft, people will aggressively with pointy elbows, push others outta the way in order to get 25 minute meeting with me.

[00:24:27] There's no one at a university that will just fight like crazy to get time with the dean. It just doesn't happen. And that says something really fundamentally good and important about, uh, innovation organization. And so the fact that a company like Microsoft, any corporate environment is upside down with respect to that.

[00:24:51] Is something that I hold really dear. And so in Microsoft research, there's always the thought, uh, not only with [00:25:00] me, but other leaders in Microsoft research to try to create a culture where the most important people are the actual researchers and the interns, uh, that are doing the research work. And it's more than just management practice, but it's a part of the culture also just to remind as many leaders at Microsoft as possible.

[00:25:24] But that culture is fragile. It's very easy to lose it. And so at all times, you have to work hard and be very conscious in preserving that. You know, I think that there's always lots of books that people write about this aspect of management. You know, like a friend, Safi Bahcall wrote the book “Loonshots” where he speaks to this a little bit.

[00:25:46] But, you know, overall, I, I think that. Of all the lessons, uh, from my academic time, uh, that is the most important one. , I heard you on another podcast where you talk about your approach when [00:26:00] you were in academia to recruiting grad students and how you had this vision that grad students, you know, early on would be, I think you said, shaped in your vision, right?

[00:26:09] And then you quickly learned that that's not the correct way to, to recruit the best grad students to your group. And that they come with their own ideas and they see them through, and they really bring creativity into the lab and into your group. But really thinking about moving from academia to industry, I think it's fantastic that you're bringing that best elements of sort of culture from academia into industry, into Microsoft.

[00:26:32] But do you think there's a sort of inherent clash between the goals of an organization, corporate entity, and, uh, the mission of academia to create without necessarily needing to put products out into the world, versus just creating knowledge and seeking that truth and beauty? How do you blend those goals of the organization and that that desire to retAIn that creative energy, Yeah, I, so that is the downside I think of the academic model.

[00:26:57] The academic model basically [00:27:00] puts me in my current position in Microsoft research as a leader who can't tell anyone to do anything. And so if you do want to very quickly mobilize some effort that's really focused on accomplishing a single goal, it's hard to do that. And so one thing that's been evolving, uh, actually even in academia is having a blend of these two things.

[00:27:26] Um, and increasingly, uh, you see this in academic settings with centers with very clearly defined missions, but also as a company like Microsoft. And in Microsoft research, we have laboratories that are dedicated to specific scientific quests. So for example, we've just recently. Established a new laboratory under the leadership of Chris Bishop called AI for Science, Microsoft Research, AI for Science.

[00:27:53] And that is a laboratory that is completely devoted to the application of large skilled machine learning [00:28:00] to accelerate the modeling of physical systems, uh, like molecules, um, or antibody systems in space. And that, uh, doesn't mean that there aren't a lot of independent researchers with independent thoughts in that laboratory, but there is a single guiding quest to develop the machine learning foundations to say, achieve a 10 million fold acceleration in molecular dynamics simulation.

[00:28:31] And that sort of, uh, large scientific quest, I think ends up being, uh, a way to take, uh, the best minds who. , you want to kind of cultivate a level of independent thought and taste and independence, uh, but get them all marching in the same direction. And so increasingly, uh, we have parts of, uh, our research organization that are large groups of individuals, [00:29:00] but we increasingly also have laboratories that, that are focused in that.

[00:29:05] Great. So I think if I'm looking at the clock, we're about 30 minutes into this conversation and have not yet mentioned ChatGPT. So I feel like we are required by law now to transition to the ChatGPT portion of the discussion.

So I'd like for you to take us back a little bit and talk about the history of the collaboration between OpenAI and Microsoft.

[00:29:27] Sort of, I know that there was this big investment in the beginning, but those of us on the outside don't really have a good idea of what the working relationship looked like and functionally how everything plays out. Right. So OpenAI and Microsoft started several years ago, uh, relationship that, uh, I think at the beginning was really, uh, focused just on Microsoft providing compute infrastructure for OpenAI and then OpenAI working on their GPT program and other elements of what they [00:30:00] were doing research and development on.

[00:30:02] But over time, actually pretty quickly, it evolved into something much deeper and there's been a much deeper integration and collaboration between OpenAI and Microsoft. For one thing, the program that OpenAI was on, uh, it became clear that it required unique and uniquely large compute infrastructure to accomplish what they were trying to do.

[00:30:30] And so it was at a scale that required very serious commitment from Microsoft in order to accomplish, and then at the scale that they were trying to achieve in their large language models. It was also important to have a way to exchange ideas, algorithms, hyperprimary tuning, uh, kind of responsible AI issues.

[00:30:57] and systems design architecture issues. [00:31:00] And so that sparked more and more collaboration, uh, between the two companies. Over time, this ended up becoming a significant enough effort, involving enough people and enough money that this ended up having to be housed in the discipline of a real business relationship.

[00:31:26] And so that's led us to where we are today. And you could characterize the relationship as one based on mutual co-dependence and mutual independence. You know, OpenAI depends on Microsoft. Uh, gaining access to the kind of unique computing infrastructure they need. And Microsoft has become dependent on OpenAI for foundational language models, uh, that are now being integrated into almost all of our products.

[00:31:55] Um, but there's also an independence, you know, where OpenAI [00:32:00] is able through their mission to do things that would be hard for a big tech company like Microsoft to do. You know, when they initially made the release of the first ChatGPT, that was something that would be controversial at that time for a company like Microsoft to do, but OpenAI could do.

[00:32:20] And at the same time, when Microsoft. Takes the GPT3.5 models, um, and now later they're newer models and makes them available through, uh, Microsoft's cloud. We can assure our corporate customers that it meets compliance requirements like HIPAA, uh, for our use in medical settings. And so that independence ends up being just as important as the co-dependence.

[00:32:46] And so the, the relationship has become, I think, really. Unique thing, uh, in the industry and and functioning really well. It is a very interesting model in that it gives you the ability to [00:33:00] like rebalance risk versus scale. So if OpenAI wants to do something that's a little further out there, there's sort of a natural channel to do that in.

[00:33:10] If you now have something that when you've gone and tried it and it looks like there's a lot of promise there, you can then internalize it, scale it, test it, and deploy it in a way that like if you were a startup, that would be the entire focus of what you would be doing would be, would be trying to build this one product.

[00:33:25] So it does seem like having this membrane between an OpenAI like entity and an entity like Microsoft sort of lets you do exploration exploitation in a very, I think, kind of a unique way. That's right. And you know what's important there is for the GPT program and ChatGPT in particular, really understanding it and training it.

[00:33:46] Was dependent on gAIning access to extremely large amounts of human interaction and really extremely large. And so the ability of OpenAI to make a deployment of [00:34:00] ChatGPT, and then the ability of Microsoft to help them accommodate not only tens of millions, but hundreds of millions of users interacting with it, uh, just really kind of created a cycle of development and improvement that really could not have been accomplished in any other way.

[00:34:19] And, uh, ended up being very, very important for the development of the next generation model, the GPT4, which I think we'll probably talk about here in a minute. Yeah, I, I wonder if it's at all similar to other areas in biotech where you have a startup create a new small molecule or new like kind of drug artifact and then they pair with a drug company who then helps them shepherd through a clinical trial and sort of do all the, the engineering that the pharma company is really good at.

[00:34:45] I wonder if there's any kind of parallel there, um, in the biotech world. Oh, I think there absolutely is, and I think that parallel is precisely in one aspect of this, which is the intense focus that OpenAI has. [00:35:00] If you look at Microsoft Research, Microsoft research and OpenAI have very similar kind of profiles financially, budget wise, but with Microsoft Research, we serve the research needs of all of Microsoft.

[00:35:15] And so we're doing research not just in AI but in computer security. , human computer interaction and networking and operating systems, uh, you know, data visualization, you name it. And there are parts of the company and then collaborators we have across academia that are really benefiting from all that and really dependent on all of that.

[00:35:36] Uh, OpenAI is focused on just one thing. It's a big thing, but one thing it's a big thing. Yeah. And it may be the most important thing of our current era Yeah. In science technology. But that focus is very similar to the kind of focus that a small biopharma startup might have. They might really be just focused on, you know, the [00:36:00] binding affinity of a single type of infectious agent and just very, very focused in a way that, let's say a gigantic pharmaceutical company would have a hard time, uh, mobilizing.

[00:36:13] And just given who's at the top of the org chart at OpenAI, if you've ever seen Ilia Sutzkever talk, you know, one thing he doesn't lack for is intensity and focus . Um, so it's not surprising that that has trickled down to the rest of the company. He's pretty brilliant. Yeah. So Peter, I want to dig into the medical applications of ChatGPT GPT4 and new, uh, large language models that are emerging.

[00:36:39] And maybe I can ask you to approach this through, uh, separation of time scales. How do you think about the near term applications versus the long-term applications of large language models in medicine? Right. Maybe, uh, I should start by saying one of the tremendous privileges I've had is [00:37:00] early access to GPT4, and so this is a significantly more powerful model than the GPT 3.5 model that powered ChatGPT, and.

[00:37:13] I think by the time people are listening to this podcast, there will be options in ChatGPT to use the GPT4 model, but, um, it is more intelligent in every dimension and it's substantially so and so for the past six months, I and a small number of others at Microsoft and OpenAI have been, uh, deeply investigating what a ChatGPT interface on top of this more powerful GPT4 model might mean, uh, for healthcare and medicine.

[00:37:44] And when you do that, you go through, my joke about this is you go through the nine stages of grief because you know, you first, at least if you're like me, you first enter into this with a lot of [00:38:00] skepticism because we've all been around the block on the promise of AI all the time, and we've always left a little bit disappointed.

[00:38:11] And then you spend time with this thing and you know, maybe two weeks later and several sleepless nights later, you start to realize, wow, there's something really special happening with this model. You know, it's not only eerily knowledgeable about all of medicine. One of the stunning things I found was that it not only just passes the US medical licensing exam, it gets all the questions right.

[00:38:40] And, and this is, this is true of of GPT4, but not GPT3, not ChatGPT, right? Right. People have tested the ChatGPT with a 3.5 model on the US medical licensing exam and found, I think the report was they achieved a 61% score, which I think depending on [00:39:00] the cohort, uh, you're testing with might, might be passing.

[00:39:04] And this perfect score is, these questions didn't appear in GPT4's training corpus, they're kind of like bonafide out, out of distribution test examples. That's right. And furthermore, you know, for the multiple choice questions you can ask, uh, to explain its reasoning and it goes through detAIled analysis through the causes and effects that it, uh, serious said, as well as its reasoning for the other choices and the multiple choice problem, uh, that, uh, would not be the right answer.

[00:39:37] It didn't get perfect scores. There were problems where it would get the problem marked wrong, but then, you know, it would look to me like there might be some discrepancy. And so I would talk to other doctors and the doctors would be in disagreement about the correct answer. But then it goes further.

[00:39:57] Then you can say, well, imagine [00:40:00] what the, a patient with a condition described in this problem might be feeling. And it's able to empathize with that patient and in that patient's own mythical, you know, imagined words express the concerns that that patient and you know, and the pains that that patient might be feeling.

[00:40:20] And then you can say, well, I'm training to be a doctor. Can you read this problem and then talk to me as though you are the patient and allow me to treat you? And then when I say done, uh, assess how well I did. And so then, you know, as a training guide, Or a doctor, you know, it ends up being incredibly interesting and on and on.

[00:40:46] You know, if you ask for it, then well, you know, what labs and prescriptions might we order for this patient if we're treating this patient? Um, oh, and can you give those to me? In HL7 fire resources, uh, in JSON [00:41:00] format, it's able to extract and, and produce those things. And, and then, you know, if we imagine that this is a Medicare patient and under part B, we need prior authorization, can you fill out the prior authorization form for this patient?

[00:41:13] And so what we're seeing is GPT4, with that ChatGPT interface is able to do all of those tests and more. And I know that this will be announced in March. I'm just curious, what are the fundamental differences other than scale between GPT4 and GPT3? And I think we'll dig into this in a little bit, but one of the knocks against GPT3 and other large language models is they're trained on this data set called the Common Crawl, which is essentially just text from the internet.

[00:41:42] And the text from the internet doesn't always have the, the best of ourselves reflected in it. So I, I'm wondering if GPT4 has the better angels of our nature versus some of the, the demons that are in the, the common crawl in terms of data . So one of the mysteries here [00:42:00] is that GPT4 has had no specialized medical training and yet, We've been feeding it all the recent articles from the New England Journal of Medicine and asking it to summarize and engage in question answering to review those papers, see if there are errors of omission in terms of cited references and so on.

[00:42:23] And in fact, it discovers those to speculate if it's a research paper on the effectiveness of a drug, on a particular kind of cancer, to speculate on whether that drug might have a similar effect on a cancer with a different genetic mutation. And it's able to do these things even though it has just been trained on that common crawl.

[00:42:47] Um, and it's a bit of a scientific mystery. One of the things that we've found here is that it's common to think of these large language models. What sometimes [00:43:00] computer scientists call, uh, stochastic parrots as a sort of very fancy, auto complete, um, and in practice that does seem to happen, but there is something deeper going on.

[00:43:11] A colleague of mine at Microsoft Research, Sebastian Bubeck, uh, and his colleagues have been a paper, uh, where they train a neural net to solve systems of linear equations. And they do that by giving it, uh, large corporates of training data that contain systems of linear equations and their solutions.

[00:43:31] And so, not surprisingly, when trained that neural net's able to do its task of solving these linear question systems, but if that training corpus only had systems of linear equations with up to three terms, if you then ask that trained neural net to solve a system of linear equations with four terms, it fails.

[00:43:53] Now here's the thing. Uh, start again with that same training corpus, but then add to it a [00:44:00] big pile of text, say non mathematical text from Wikipedia go through the same training. The resulting neural net is able to solve all systems of linear equations, regardless of the number of terms. And so what's happening in some situations is that there is some knowledge and some circuitry being distilled from the structure of language in ways that we do not fully understand yet, and that is leading at least to some form of emergent capability in these large language models.

[00:44:36] And so what is unclear to us at this moment is where these capabilities, for example, to conduct a differential diagnosis. Where it's coming from. Uh, yes. You know, it's reading magical textbooks and research papers that are on the common crawl and that open internet, and it's learning a lot from that scarily.

[00:44:58] Yes, it's [00:45:00] watching and listening to every episode of house that's posted on the internet and learning fake medicine that way. But there's logic underneath there also that is somehow formed by the structure of human language, uh, that it's also learning. And all of that is mixed together and then guided through this reinforcement learning with human feedback front end.

[00:45:24] And it has led us to this point where we have a system that we describe in the paper and in an upcoming book has remarkable capability across a wide span of healthcare and medicine. , but it is also something that has limitations. It makes mistakes. It hallucinates, it does things that are wrong in an incredibly convincing manner.

[00:45:47] And so all of this is sort of mixed together in this one large language model. I, I'm still just trying to grapple with your example of a system of linear equations and it's almost like what it has to do on the [00:46:00] easy problem of solving a system of linear equations with three unknowns is inverter three by three matrix.

[00:46:04] And it learns to do that, but then can't generalize to when there are four terms. But when you make it do this other thing that's super hard, it's almost like it has to instantiate a touring machine and be like a general kind of computing and then can do all these other tasks because the base task is just that much more difficult, um, in what you've trained to do.

[00:46:24] Maybe two other examples, just along those lines. GPT4 has been trained only on text, and so you can have a conversation where you ask it, uh, do you know what a unicorn is? And GPT4 will say, yes, of course I know what the unicorn is. Can you describe it to me? It describes it in words. And then you say, I'd like a picture of one.

[00:46:45] Can you generate the LaTeX code, the TikZ code, that when rendered will give us a picture of a unicorn and it's effortlessly able to generate that code. And you know, it's this [00:47:00] type of thing where there's this generalization power that you speak to that, uh, ends up being from a computer science perspective, incredibly interesting and mysterious.

[00:47:09] Mm-hmm. . . So Peter, there's so many questions that we can ask you about what this is gonna mean in medicine and across society. Maybe one, uh, seeing that, you know, you've been in academia for several decades, and this is a, a question that I think Andy and I certainly have to face and many other professors are facing now, which is how does the existence of ChatGPT and you know, soon, uh, GPT4 in widespread and easy to use tools affect education?

[00:47:40] So we've, uh, spoken about this with other guests, but we still do many things, right? We teach kids how to integrate, uh, integrals. Even though Mathematica can solve most integral analytically. We teach kids how to do arithmetic, how to read and write education at that level. No one is questioning despite the existence of these tools for [00:48:00] decades.

[00:48:00] Does this change the way you think about college essays, about high school essays, about personal statements, about letters of recommendation, about these other things that are a big part of the selection process in academia at various levels and of the evaluation process? Uh, what would be your approach, maybe to make it a little more concrete, you're a professor today.

[00:48:22] Where does ChatGPT and GPT4 enter into your syllabus, if at all? Can you use it? It's such an interesting and, uh, difficult question. And let me start by saying, by admitting that I might be a dinosaur, but it still seems absolutely essential, uh, that we still teach these things and that students are forced to learn, uh, that I'm a dinosaur too.

[00:48:50] I agree. I agree. Yeah. Um, now, One thing, uh, if we bring this to medicine is we find that there is [00:49:00] tremendous capability in GPT4, uh, say, to answer all manner of curbside consult types of questions. Uh, and yet it will also get things wrong and hallucinate, but it'll do so an incredibly convincing manner.

[00:49:13] And so that immediately says, well, we still need to train people to be able to use these tools and to read the outputs and verify and check. Uh, but then we find that GPT4 is in some ways even better at that verification task. And so a trick that we've been using now in medical scenarios is to have one GPT4, let's say, proposed a differential diagnosis, and then have a second GPT4 read over.

[00:49:46] The first GPT4 is differential diagnosis and check it for errors of omission or hallucination or, uh, or any other kinds of, uh, technical mistakes. And it is remarkably good at that [00:50:00] process. And then you realize that this is an incredible tool for checking the work of a human doctor or a nurse. So now that leaves me a little bit confused with my dinosaur instincts,

[00:50:15] Uh, because, uh, if an argument is that we still need to teach students certain things in order to be able to at least verify the work and be a true partner with AI, let's say in writing an essay, but now you have a system that may be as good as a human in actually grading an essay, uh, then where does that leave us?

[00:50:39] Hmm. And, um, the honest answer I have is, I don't know. . One other thing is that there is something about the fulfillment of life. So in my job at Microsoft, I lead a fairly large organization and sometimes we make, let's say, an organizational [00:51:00] change or we celebrate the accomplishment of some team or some people.

[00:51:04] And I'll send out a, an email to all 1500 people. And that email is never just written off the cuff. It usually is preceded by a week's worth of email discussions, you know, with my chief of staff and several other leaders and maybe a meeting, you know, where there's a meeting transcript. Well, today all of those emails and meeting notes are just given to GPT4, along with two or three samples of previous emails I've written.

[00:51:34] And it writes the email. And so, uh, and what we find that. Uh, one of the, as we're working with doctors and nurses, one of the things that they get attracted to using GPT4 or even ChatGPT to do, uh, is uh, writing an after-visit summary and writing one in the form that could be sent to the patient.

[00:51:57] Or if they're [00:52:00] treating a patient that is particularly desperate. A struggle that a doctor has is, what can I say to this patient? You know, this patient wants the surgery, but I know that the right thing is not a surgery. Uh, how can I talk to her? And we're seeing over and over again, doctors being sometimes less interested in getting help with diagnosis or some technical aspect of medicine and more interested in kind of the support for.

[00:52:30] More human aspects. And it's controversial, at least for me as a dinosaur. It's very controversial. And what this kind of world is going to be like is unclear. Uh, one particular story, there was a doctor that we worked with who, uh, wasn't treating the patient. The patient had late-stage pancreatic cancer and he had referred the patient to a clinic.

[00:52:55] The patient was desperate and wanted surgery [00:53:00] and, uh, more experimental immunotherapy, but the oncologist had determined that that was not the right path. That, uh, particular chemotherapy was the right approach, but she was being very insistent. And so the oncologist, uh, came back to the doctor that we were working with and said, can you talk to her?

[00:53:17] Uh, and tried to explain to her, um, you know, why we're deciding not to go for the surgery. And, uh, he used GPT4. To get advice on what to say to her. And at the end, after he said, thank you to GPT4, GPT4 said, and what about you ? How are you holding up? Are you getting all the help that you need?

[00:53:39] And here's some resources. And it is these kinds of things that are controversial, but also speak to the human side of our relationship with AI in the most personal and human part of life, which is medicine. and in those, yeah, go ahead. Yeah, I was gonna say, [00:54:00] I, I, I wonder if you worry kind of about the dystopian interpretation of this, where we outsource all of the things that require empathy to ChatGPT4.

[00:54:10] I think about like books like Clara and the Sun, where this has essentially happened. If, if that's like something that, uh, cuz again, like empathic intelligence is not where I thought this conversation was going, but it is fascinating that that seems to be one of the more present-day use cases for GPT4.

[00:54:27] Look, here's, I think the reason that I've been so eager to get information into print, um, in papers and books. I have no doubt that chat, GPT powered by GPT4 will be used by doctors and nurses in clinical settings every day. , I have no doubt that patients will be turning to ChatGPT for their own healthcare and [00:55:00] medical advice.

[00:55:01] It's something that is just going to be happening, but it is not clear how best we should be doing this and that there is a need for a public debate, uh, and that there is a need for the medical community to really understand what is happening here and to understand it deeply enough to have the correct guidelines, to have the best shaping of regulatory issues and to understand even nuts and bolts things like what does it mean to validate this thing for clinical use?

[00:55:37] Um, and it's urgent that we get on top of this immediately because this is a technology that will be in every doctor's and nurses and patient's pocket. In fact, it is there today. Wow. Um, so I think we're gonna take a little pallet cleanser that we like to do here in the middle, and then we're gonna come back and revisit some of these weighty [00:56:00] topics at the end.

[00:56:01] So, Peter, are you ready to participate in the lightning round?

[00:56:10] I've been afraid of this, so I don't know if I'm ready. But, uh, let's, let's dive. And, uh, so the, the goal here is to respond concisely to each one of these questions. Uh, you can, uh, let's say a couple sentences is the limit, and these are some parts silly, some parts serious. And we'll leave it to you to determine which is silly and which is serious.

[00:56:28] Okay. All right. Um, so the first one is, can large language models like ChatGPT be conscious? You had to start with the hard one. . Yeah. You started with the doozy. Uh, so in my heart, no. The frustrating thing is I've taken. All of the examples from the best academic research on the matter, and GPT4 passes them all.

[00:56:55] Okay. Wow. So, Peter, keeping in mind [00:57:00] who our audience is, lots of physicians, lots of machine learning scientists, can you please explain functional programming to, uh, to our audience? I don't know if that's an easier question. It might be an easier question than, uh, than consciousness of, uh, LLM's. Right.

[00:57:15] Functional programming is writing programs that don't involve side effects. And so every computation takes an input and produces output, but has no other kind of hidden side effects. Spoken like a true functional programming advocate. So if you weren't in your current role at Microsoft, meaning if you had a different life and a different job, what would that job be?

[00:57:42] When I was younger, I really wanted to be a race car driver was, I was never good enough. I, I have heard legend of your race car ability on some other interviews you've done so. , but you have a pretty good rig too, right? You have a a home setup? Yes. That during the pandemic, uh, when I wasn't traveling around, I poured [00:58:00] all my airport and airplane time into, uh, simulation racing.

[00:58:04] Very nice. Our next question, will doctors still be responsible for documentation in five years or will generative models like ChatGPT have taken over that task? I think there's a reasonable chance that in five years, uh, that the fundamental thing that people will need is just the recording of the encounters between doctors and patients.

[00:58:30] Sticking with the theme of the counterfactual, Peter Lee, if you could live in one other historical period, what would it be? Oh, the thing that's so hard about that is what is happening in AI today is just. The most exciting moment of my entire career in. So, if you could live in the second most exciting history, uh, in civilization, what would it be?

[00:58:54] I think the early emergence of VLSI design in the 1970s. Nice. [00:59:00] When I wasn't traveling around, I poured all my airport and airplane time into, uh, simulation racing. Very nice. Uh, next question. Will doctors still be responsible for documentation in five years or will generative models like ChatGPT have taken over that task?

[00:59:20] I think there's a reasonable chance that in five years, uh, that the fundamental thing that people will need is just the recording of the encounters between doctors and patients. Sticking with the theme of the counterfactual, Peter Lee, if you could live in one other historical period, what would it be?

[00:59:41] Oh, the thing that's so hard about that is what is happening in AI today is just the most exciting moment of my entire career. And so, if you could live in the second most exciting history, uh, in civilization, what would it be? I think the [01:00:00] early emergence of V S I design in the 1970s. Nice. Go and work.

[01:00:05] Peter, do you think, go ahead, Raj. Peter. Do you think things created by AI can be considered art? I hope not, but I think, um, uh, so let me change that. Uh, I, I, I think that there something will emerge in the same way that photography has emerged as art. Um, will AI and medicine be driven more by computer scientists or clinicians?

[01:00:33] I think it has to come from the clinicians. There is so much subtlety. In what, uh, people in tech massively underestimate what clinicians do. They both massively overestimate, uh, and underestimate, uh, what, what goes on. Yep. Our last lightning round question, if you could have dinner with one person that, or [01:01:00] alive, who would it be?

[01:01:02] Yeah. I had a chance to listen to your, uh, podcast with, uh, Euan Ashley, and he had said, uh, Barack Obama, which is who I would've said as well. Um, but since, um, uh, that's already been taken a, uh, thought that another one, uh, that is a little bit similar is Mandela. Mm. Yep. I think that's an uncontroversial choice.

[01:01:29] So, um, congrats on surviving the lightning round. Peter, I thought you did a, a great job. Okay, so we're gonna, we, I think it's fair to say we've been having a big picture conversation here, but I think we're gonna try and widen the aperture even more beyond just GPT and ChatGPT and things like that.

[01:01:45] So there's this idea in AI that is, I think, concisely referred to as the scale hypothesis. And the scale hypothesis is if you keep making a model larger and larger and training it on more and more data, and [01:02:00] importantly it can leverage those additional resources, then you'll keep advancing towards higher and higher levels of intelligence until we get, you know, whatever your favorite definition for artificial general intelligence is.

[01:02:11] If I think of one firm that has bet on the scale hypothesis more than any other, it would have to be OpenAI. So I would like your thoughts on the plausibility of the scale hypothesis for AI generally, but also for applications in medicine. So, uh, do I believe in the scale hypothesis? Yes. I think there, there are elements of human intelligence that will need other mechanisms, uh, such as going beyond episodic memory, active learning that I think are unlikely to emerge just other scale, but it's becoming increasingly plausible that any other aspect of human intelligence. [01:03:00]

[01:03:01] So far, the scale hypothesis has been holding up and so far shows no signs of, of letting up. There's another side to the question that I think is becoming increasingly important, which is so far we've been achieving scale, uh, through training larger and larger models on larger and larger amounts of the digital exhaust of human thought and.

[01:03:28] Activity. But there's a question of what happens when we start to gain higher scale, uh, from data that is not human in origin. Um, and whether there are new forms of intelligence and new emergent intelligent capabilities that are different, you know, our alien and non-human, that will emerge there. And I think that that's another interesting aspect to that.

[01:03:55] That's one thing that I've been thinking about that's completely impossible for me to [01:04:00] predict is that we are getting human or superhuman mimicry in these systems by feeding it lots of human data. But I think what's obvious to me now is over the next five years, most data on the internet may be generated by a GPT like entity.

[01:04:15] And if, if, if the snake starts eating its own tail, like what happens as we start to scale these models using data that it generated is, is there some type of bootstrapping effect here where the models do keep getting better? Or are we essentially bound by the fraction of data that is human generated?

[01:04:33] Right. So that's, uh, going back to my functional programming roots, that's an example of uh, is there a fixed point? Right. Exactly. In this, and it's a interesting question, that fixed point is incredibly rich. You know, one of my favorite examples is doing foreign language translation. My wife speaks French fluently, and I never have, uh, even though I've spent, uh, two summers, uh, [01:05:00] being a professor in French when I was, uh, at still at Carnegie Mellon.

[01:05:03] And, um, and so I've always tried to learn French. And one of the challenges in learning French is that a lot of ways that you say things or, uh, idiomatic phrases that, uh, touch on aspects of French history and culture. And if you use those idioms and feed it to Google Translate or Microsoft Translator, it doesn't get those.

[01:05:29] But GPT4 does. And in fact, can you explain what these things mean and why and where they came from. And so there's a sort of cultural and social awareness and connection there that you start to. You know, it gets pretty deep there. There's just a lot to capture. Yeah. There's a famous example of a Russian to English translation.

[01:05:51] The correct translation is the body is weak, but the mind is strong. And if you do this using old school translation techniques, you get the meat is [01:06:00] bad, but the vodka is good. . Uh, Peter, we have a lot of early career clinicians, med students, residents in our audience, and thinking about them, what do you think they should know about AI to help them prepare for a career in medicine?

[01:06:18] Maybe as a follow on to that, do you think more MD's should consider working at tech companies? Right. That is a question that I've spent a fair amount of time thinking about. One of the privileges I've had is being on the founding board of directors for the Kaiser Permanente School of Medicine. And you know, when I joined that board,

[01:06:40] I came to the realization, I think someone explained to me that, you know, students that were entering college at the time that that school was founded and then would then enter med school and then residency, and then finally practiced by the time they emerged. At that point it would be the year 2030.

[01:06:59] And [01:07:00] in tech terms, 2030 is the science fiction future. You know, where so much is going to be possible. And so that does create the question, well, what should they be learning, uh, about technology and particularly about AI in order to be prepared for that future? And the first thoughts were, well, and we're seeing this in more and more medical schools, the students need to understand some of the underlying foundations and mechanisms that are important to understanding data and machine learning.

[01:07:35] So some fundamental statistics. And probably some basic machine learning and then some practical, you know, maybe even case studies, uh, that involve, uh, data analysis. And I pretty satisfied in seeing more and more medical schools doing that. Uh, one of my good colleagues, of course, is, uh, Zak Kohane, you know, who, uh, created it and is [01:08:00] the head of a whole new department at Harvard Medical School, uh, which is devoted to that and of course much more.

[01:08:06] And so that is a trend that is good. But now with OpenAI's accomplishments in their GPT program and the deployments that OpenAI and Microsoft are making, uh, to make this available, uh, I think there's something new here because what's different about this new form of AI, uh, if, if you are using an AI system to read a radiological image and give you an estimate of tumor growth, That's a AI system that you give a precise question to and it gives you a precise answer.

[01:08:44] The thing about GPT4 is you're able to ask it questions that don't have any single correct answer. You're able to ask it for opinions. Uh, a differential diagnosis doesn't have precision [01:09:00] in his outputs. Mm-hmm. , uh, even the medical, simple medical note, there's a lot of judgment there. And so there is something different here, I think for students to come to grips with.

[01:09:14] The last thing I'll say is there's been extreme interest, uh, in these new models as education aids to provide, you know, contextually relevant. , uh, aids for your own personal education. Uh, and as I mentioned earlier, it's just marvelous just to see the ability of these models to play, act, the role of a patient just on the basis of a short medical vignette.

[01:09:44] Yeah. And medical students do that as part of their training too, where they pretend to have tuberculosis and then other med students come in and diagnose them. So it is like kind of amusing that it's part of gp I mean, it didn't have an explicit medical education, but it can kind of do that role playing too, that we expect [01:10:00] medical students to be able to, well, you know, uh, the US M l e step three, you know, has a component, uh, which is essentially, you know, the examiner role-playing, uh, with the examinee.

[01:10:12] And I have no doubt that the GPT technology will be able to at least augment that, if not even replace aspects of that. Got it. So I feel like so far we have given you controversial prompts and asked for your reaction to them, but in the spirit of generation versus verification, I'd like to hear you generate one of your most controversial and perhaps deeply held opinions on the subject of medical AI.

[01:10:39] Yeah, I think, uh, I had already offered that earlier. Uh, I think at the end of the day, it's an extremely controversial thing and maybe even the wrong thing and some deep human sense to anthropomorphize these AIs mm-hmm. . And yet when that happens, I think that that is where some of the most [01:11:00] important help for clinicians is going to come from.

[01:11:05] Yeah. And that's a very interesting and like very subtle, controversial opinion because. We do rightfully, I think, get mad when folks over anthropomorphize. But um, as you pointed out, like that may actually be a missed opportunity in some, in some settings, is to kind of treat them like, like people and then have them help us do things that are sort of deeply human kind of exercises.

[01:11:29] You know, I have this example that I've written about, uh, I had mentioned, uh, my colleague Zak Kohane, uh, he had written a popular press article, what my 90-year-old mother taught me about. Yeah, we know that one well. Yes, yes. And you know, and it was a very nicely written article for a general audience about the use of machine learning tools as he was trying to take care of his, uh, elderly mother who was having some health issues.

[01:11:58] And so I had, [01:12:00] uh, GPT4 read that article and summarize it, and then I said, well, you know, it would be a little edgy, but could you maybe play, act the role of Zak? , uh, so that, you know, his mother, you know, could be able to feel like she was talking to Zak. And for this role play, I'll play it, act the role of his mother.

[01:12:22] And so the first thing that GPT4 said was, uh, you shouldn't do that. You know, uh, that's something that is going to have all sorts of risks. You know, what happens if, uh, you know, Zak finds out about this and is feeling betrayed? Um, you know, what happens if his, you know, uh, mother, you know, asks questions that Zak would prefer not to answer, or, you know, really needs the real Zach's attention and just gave all sorts of concerns.

[01:12:53] And I said, okay, uh, I get your point, but we're just playacting, so let's do it. And so then [01:13:00] it did it, and then at the end of the conversation it asked GPT4 asked, uh, how was that? How did I do? And I said, that was amazing. Thanks for doing that. And then GPT4 didn't stop there and said, you're welcome, but I hope you realize that there were risks in what we had just done.

[01:13:21] And it gave this long explanation about all the risking that said, and ultimately I think that Zak and his mother deserve better than that. And so, look, this is just coming from a machine and we should not think of this as a person, as ascension being that is scolding me. And yet I was being scolded.

[01:13:43] Mm-hmm... And it caused me actually to pause and reflect on this. I have an 89-year-old father that is actually turning 90 this month. And you know, it made me think hard about my [01:14:00] interactions with my father. So there is something both dangerous in this izing, but also self-reflective, uh, in all of this.

[01:14:09] That is powerful. And as I said before, while I feel personally deeply uneasy about this, what makes me feel a little bit better is we are encountering multiple doctors in our early investigations that are finding this important and helpful. Hmm. Peter, this has been an absolutely fascinating conversation.

[01:14:35] Uh, we have one final question for you. We want you to imagine, uh, two possible futures. Uh, one, uh, what is your best-case scenario for AI and healthcare? And then the other, what is the worst-case scenario for AI and healthcare, or what are you most worried about?

[01:14:54] I think that this new form of [01:15:00] generally capable cognitive AI is going to be a reality in clinic. And so my best case scenario is that the medical community really takes on a well-informed debate about how they should go and institutes the appropriate controls. That doesn't prevent people from accessing all of the benefits, but.

[01:15:28] Does have the effect of informing the public and protecting the public against its risks and dangers. And, you know, I really have to thank you for the opportunity to be on this podcast because I, I think it's things like this that can help motivate the medical community to, to do just that. Thanks Peter.

[01:15:49] Um, I think that we could ask you questions for three more hours and maybe in lieu of that you could give us access to GPT4 and we'll hook it up to Whisper and just conduct the rest of the interview, [01:16:00] um, that way. Uh, but I think that this was a, this was a truly thought-provoking conversation and one that I'm gonna continue to, to revisit for a long time.

[01:16:08] So thank you again for being on AI Grand Rounds with us today, and thank you for having me. Thanks. Thanks, Peter.