

David Shaywitz [00:00:03] All righty. Ready to do this? Alright. Awesome. So welcome. My name is David Shaywitz. I'm a physician scientist and lecturer in Zak's department and also a distinguished R&D fellow in digital and data at Takeda in Cambridge. I'm your moderator this afternoon, and I'm so excited to be here today, celebrating -- I should hold it up -- celebrating the release of this fascinating new book, *The Revolution in Medicine, GPT-4 and Beyond*, which we hope does better than *Bed, Bath and Beyond*. So it's by Peter, Carey and Zak describing their early experiences with GPT-4 and what they see as the implications for health care. The book, with a foreword by Sam Altman of CEO of OpenAI, is already in its third printing. The audiobook will be out at the end of May and it is or soon will be available in your favorite e-book format. Now Amazon has had some trouble keeping in stock, so maybe they're trying to tweak Peter, I don't know, but hopefully it'll be replenished soon. So a quick word about the authors who otherwise need little introduction and they'll give you an additional introduction themselves. Peter is an EVP at Microsoft and head of Microsoft Research and former head of Computer Science at Carnegie Mellon, where he conducted A.I. research for several decades. Carey is a distinguished journalist whose career has included serving as the Moscow bureau chief for the L.A. Times, the Boston Bureau chief for The New York Times and for Bloomberg, and as well as work at the Boston Globe and NPR. Zak Zak is a pediatric endocrinologist, a highly accomplished AI researcher. He's also the founder and head of the Department of Biomedical Informatics at HMS and also the editor in chief of the recently announced journal *NEJM-AI*. So, yeah, you got the varsity tonight. I've known Zak in particular for decades, and I got to know Peter quite well when, together with Amy Abernathy, we co-led a National Academy of Medicine's project looking at the impact of digital and data on the COVID response. I've also read and appreciated Carey for years, including as a savvy participant in Zak's annual Precision Medicine conference each fall. What struck me most about my conversation with these authors regarding GPT-4 in particular, is their reaction to this technology. Now, you don't have to look far to find extraordinary predictions about A.I.. Google's CEO said the impact will be more profound than the discovery of fire. Andrew Ng said AI is the new electricity - no more the new oil -- the new electricity, and then even non-techies like New York Times journalist Ezra Klein have said GPT-4 changes everything. Zak and Peter, they are two of the most grounded scientists I know or have known. Very realistic and pragmatic as well, of course, as hopeful and optimistic, Pragmatic optimists, you might say. But they labor under no illusions and yet they reacted to GPT-4 in a way I've never seen from them before, as if it was something unlike anything they'd ever experienced, a step change rather than an incremental advance. Carey, as ever, zoomed right in on the implications for patients. And I'm so excited to touch on all of these topics in tonight's conversation. So before we go, let's start with a little bit of a sort of an introductory question. Maybe each of you starting with Peter and going down the line, could talk a little bit about your role in the book.

Peter Lee [00:03:33] Well, thank you all and thank you for being here. David. That was really fantastic. So I think the real instigator of this was Kevin Scott and Sam Altman, and it was the end of the summer in 2022, and they pulled me into the tent. That's the Microsoft-speak for being brought into an effort that's still being kept secret. And they disclosed this project to build something at that time was called *davinci-3*. That was the code name that would eventually become GPT-4. And they wanted to understand the implications for healthcare. And so I was very intrigued and got started. You go through nine stages of grief when you encounter something like this because you start off with incredible skepticism and then annoyance and frustration, and then there's sort of a degree of amazement and joy and a great deal of intensity. But it was pretty much by that fall that I realized that there was something very important happening here, and that I needed help. I didn't have enough medical or health care expertise to do this. And so I thought about who

I could call to get help. And that ended up being Zak. And I'm guessing Zak probably rolled his eyes at me also, not really knowing what this was all about. And so he had to go through his his own nine stages of grief, which he'll probably talk about. And then the question was, what is going to happen here? There are all sorts of questions, some amazing possibilities, but also significant concerns and risks. And so the idea was we need to have the medical community take ownership of the question of whether, when and how technology like this should be used or not used. And for that, the medical community needed to have help to get up to speed as quickly as possible. And so the idea was to write this book. And then the next question was, there are two problems. First, the PR people at Microsoft hated this because, oh, give me a break. Microsoft and Harvard, two elite institutions. You know, what's going on with that? But the other thing also was the credibility of the book. And would this be looked at as kind of sales for Microsoft technologies. And so we thought about this. Also, I think Zak and I, even though Zak was signed up to help write this book, neither of us felt that confident about our abilities to write. So can we get help there as well? And so this is where a journalist, a real journalist, I think, came into the picture. And Zak had worked earlier with Carey. And so we invited Carey. And I don't think Carey ever went through -- I don't think you ever went through your stages of grief. You just embraced the project in the beginning. And so we started writing and then Chat-GPT happened and suddenly we felt incredible urgency and we needed to finish the book immediately. And that more or less ruined our holiday breaks in December. I don't think it did a whole lot for our friends at Pearson either. They had to bend over backwards to make the book happen on time. But everyone pulled together and we had amazing program management from Weishung Liu. And together we we managed to pull it off. And so what we hope more than anything now is that people in the medical community will consider this book required reading to engage in the public debate about the role of this technology or not in health care and medicine, and that this helps the world make better decisions on this. So just to see you all here, I think this is really -- I'm just very grateful. So thank you all.

Carey Goldberg [00:07:57] And so to add to what Peter said, yes, my role was basically to deprogram Zak and Peter from their academic writing habits. And early on I may have said something like "dare to be sexy."

Zak Kohane [00:08:11] And I think also I heard "enough throat clearing."

Carey Goldberg [00:08:14] Yes, enough throat clearing. There was quite a bit of retraining, which they and they they did fantastically. They in particular, they tolerated having whole paragraphs and pages killed, which is a very important thing to do when you're being edited. And I would add that my other main role, I felt like, was to have the perspective of the general public and of the patient. And so I could do that because I don't have an MD or a Ph.D. and I've been doing that for years and years in my health and science reporting. And there was a period when they had access to what was then davinci-3, and I did not. But I was getting more and more of a rudimentary understanding of what it was and what it could do. And my dominant emotion was definitely, okay, we get it. It has inaccuracies, it's hallucinates. Just give it to me. Like, I just want it. I just want to be able to use it for my own queries, my own medically related queries. And I think that what I came away from working on this book with was an understanding of just the incredible usefulness that this can have for patients. And in fact, you may have seen there was a study last week in JAMA Internal Medicine where they compared the responses by doctors and by I don't know which GPT it was, I think it was Chat-GPT, to medical questions on Reddit and they found that the GPT responses were ten times likelier to be considered empathetic by the evaluators than the responses by the doctors, which of course, because

GPT-4 can write you five paragraphs in 10 seconds, right? So they can it can, it can go on and on. And I think that is just a little hint at what we're going to see going forward in terms of direct to consumer consumer use by patients. But meanwhile, of course, for it to be used optimally in the incredibly regulated atmosphere of the American health care system, there's going to be a lot of work to be done. And I suspect that a lot of the people here are here because you're interested in doing that work. And I would just say from the patient perspective, you go! Like, this is this is going to be incredibly useful. And what what an amazing plot twist in the human condition and what an amazing juncture to be at where you will be able to contribute to how this technology will be used for the good of patients. So that's my spiel.

Zak Kohane [00:10:53] And so when Peter first reached out to me last fall, my first reaction was skepticism. I had done my Ph.D. in the 1980s working on artificial intelligence. I was surrounded by very world famous smart professors who were telling me how they were going to change everything. And so I'd heard this before. And so I was wondering, what's the game here? What's the game here that Peters engaged in? And then I got access to it and I realized things really were changing. And for example, I went through a very, very detailed workup of a patient that I was called to the newborn intensive care unit who ended up having ambiguous genitalia. And this program went step by step through the workup with me all the way from imaging to a molecular diagnosis. So that was very impressive. But then when I kept on talking to Peter, it became clear to me that he was motivated by having safe but large health care impact, as I was motivated, As I came to use GPT for what we then called davinci-3 more and more, it became clearer and clearer to me that the thing that I was most worried about is health care disparities. Not disparities in the way you normally think about it -- different disadvantaged populations. I'm talking about disparities just within here, the medical industrial complex here at Harvard Medical School affiliated hospitals. I'm just talking to a friend of mine who's has a spouse going through a difficult health journey and he was telling me how delightful it was to have a doctor who really understood what was going on, who understood the plan. The light was on. Then he'd go talk to another doctor or another doctor, and the light was not on. And there was huge unevenness. And that reminds me of my own intuition just from experiencing medical training and medical care, which is there are huge variations. There's some brilliant doctors. But there are some also non-brilliant doctors and some doctors who might have been brilliant but then are harried, squished by the forces that propel modern medicine. And then when I saw Chat-GPT go out and yes, doctors were using it, but just as importantly, patients were using it. All of a sudden I realized the disruption was going to happen, whether or not medicine and the medical establishment was going to pick up the torch. Patients were going to use it. And so then my determination for the book was as follows. Yes, we were going to address patients, but really what I wanted to do was to address the doctors, particularly as health care providers, to say we should try to understand this technology so we can lead in a safe deployment so that we can actually lead in actually making our own lives as clinicians better. And so we can link up with patients who've been given this intellectual augmentation as patients and I have to say that so far the reaction I've seen from younger doctors has been very, very positive. What I'm hoping to see is now much more of a conversation among doctors, among the professional societies, and also between patients and doctors about the use of these technologies so that that those disparities that I told you everybody experiences, no matter what their socioeconomic status, the disparities that we experience in health care can be evened out upwards.

David Shaywitz [00:15:14] All right. So let me start off with a fact from each of you, in a sense. And then y'all can respond to it. So again, the book is incredibly readable, and it is it does not come off like some corporate claptrap.

Carey Goldberg [00:15:34] Truly we should say Microsoft had no editorial control.

David Shaywitz [00:15:38] And I would say that is...

Peter Lee [00:15:39] Really very obvious.

David Shaywitz [00:15:41] Yeah. Yeah I would say that yeah, like one or two slams at Google. But beyond that -- I assume at this point it's just reflexive. So fact one is, Peter, you say, you know, incredible technology. we don't really understand how it works. You actually write that. Zak says, Oh, and P.S., we don't really understand the data on which it was trained. And then Carey says, and it's going to be used anyway. And then for good measure, Peter says, it's smarter and dumber than anyone, you know. So how does this triumvirate put all that together? How should we put it together? Should we feel empowered, as sort of Zak is suggesting? Scared? Both? How should we approach thinking about it?

Zak Kohane [00:16:26] Peter, what do you think?

Peter Lee [00:16:27] Yeah, well, so as a computer scientist, one of the things that's been both inspiring but incredibly frustrating for me personally is when you have an understanding of the underlying architecture -- of the neural transformer architecture -- there are certain things that our best computer science, cognitive science and psychological research tells us should not be possible. So then it's a mystery. You know, it's very popular, particularly in social media, to get GPT-4 or Chat-GPT to stumble and do something stupid and say, "Aha, you see, it doesn't play chess. It's obviously not intelligent," but it isn't very convincing to my mind's because I can probably pick any human being in this room and cause you to stumble in some similar way. What's more mysterious and meaningful to me is when it ends up doing something that clearly displays causal reasoning or counterfactual reasoning or appropriate decision making in morally charged circumstances or planning involving common sense reasoning in the physical world. Any of these things which provably should not be possible in this underlying neural transformer architecture. And yet they are happening. They may not happen with 100% kind of infallibility. Maybe it's only 90%, but they're happening. And so this has been both the source of amazement and frustration for me. And I spend -- you know, we agreed -- I think Carey said, okay, you can spend one chapter on this. So that ends up being chapter three, which is basically probably the longest chapter in the book, which is basically my attempt to explain to the medical community why I'm so frustrated on this point. Ultimately, I think I turned Chapter six into a continuation of that frustration. So if you are reading the book and you get frustrated that you can skip chapters three and six and others, but this is part of these nine stages of grief where there is this amazing thing that as a computer scientist, I need to understand how it works and how, you know, why it's happening. But at the moment, it's my opinion that the best computer science and cognitive science research fails to explain to us. And ultimately, I think I've just resigned myself to the fact that the questions are probably very similar to trying to understand our own brain function. And that's not to say that these things are alive or sentient or truly intelligent, but there's a mystery that's similar in trying to understand the mechanisms by which GPT-4 is doing its stuff. And so it's with really mixed emotions here -- and I have to thank Carey. Carey, really, in writing the chapter pulled enough emotion out of me to be able to express that.

So even if you aren't a computer scientist, if you do read that, I hope you get something out of that chapter.

Carey Goldberg [00:19:51] So this is therapeutic as well. And I think I would add again, from the patient point of view, we don't care. I mean, I'm glad that, you know, it's fascinating intellectually. But the fact is that it does what it does. And in fact, one thing that was interesting to me in working on the book was seeing that both Zak and Peter were basically working empirically, like they both had deep computer science knowledge, but a lot of the book is them throwing stuff at GPT-4 and seeing what it does, which is what anyone can do, right? And so that in fact is where -- you were saying you, you suggest that everyone spend at least an hour on GPT-4 every day like just interacting with it and and and Peter also has some very interesting parts about how you develop a relationship with this AI which is which kind of a tease because then AI changed and it got kind of lobotomized, which was a little sad. But anyway, and then I think the only other thing I'd like to add is that one of the motivations for the book was also that this is a very fraught time in terms of AI. It's very scary for many people, and yet at the same time we can see the tremendous good that it can do. And part of focusing on AI and medicine was that it's the sphere, it's probably the leading sphere where it can do very obvious good.

Peter Lee [00:21:15] You mentioned that Zak uses the word "genteel" to describe the demeanor of GPT-4, but I think in the early days when we were really working intensively on the book, davinci-3 wasn't fully mature yet. It was still undergoing a process of alignment, So it was more prone to disagree with you. It was more prone to scold you, to argue, to accuse you of being wrong when it was clearly wrong and a bunch of things. But as all of that has sort of been beaten out of the system, I think there was sadness.

Zak Kohane [00:21:53] There is sadness because in the case of the of the ambiguous genitalia, for example, it said something which I disagreed with, and I argued with it and it argued back. And in the end we agreed on our disagreement. It raised a rare reason why it might be right. And so I said, Yes, you're right. But it's that's rare. My hypothesis is more probable. And davinci agreed. But that's gone. So this felt like a real intellectual colleague. Now it feels more like a blander augmentation. And the part of me that, went into computer science, because of this gee whiz thing, really enjoyed it having this more independent personality.

Peter Lee [00:22:45] An example of this is davinci-3. Early on for us, back in October and November, I asked davinci-3 about the first research paper on this, should it be just one page like the original Watson and Crick paper on DNA. And is said that would be an interesting homage to that paper, but it might be very difficult to do just given the richness of the topic. But if that paper does get written, I should be a co-author. And I said, Wow. Okay, I agree. But should you be the first author and it says, Well, I don't want to be immodest, but yes, I think I should be the first author. And now, if you asked today the same question, it's just very, very polite. And so it's probably what the world needs today. And there's an aspect of this as a practical matter in medicine, particularly where I think it's a much safer and more correct thing today. For the book of this was a big problem for us because we had written the book and then, you know, as things were evolving, the nature of the system was changing and we had to decide how to handle it. And so you'll see in the book that in different chapters we make reference to both GPT-3 and to GPT-4, and hopefully it's not too confusing.

David Shaywitz [00:24:09] Zak, Did you want to say a word or two about the training data? Because you've talked about that, you've emphasized that.

Zak Kohane [00:24:17] So it's pretty incredible that this program can answer question as I posed it. Here's an atypical presentation for any known genetic disease. Here's ten genes that were mutated. Which one of these is responsible? We were able to independently -- haven't published -- figure out based on some wet lab downstream work. It actually picked the correct gene. The same program was able -- this was back in December when my son was saying he has to do a little speech at our synagogue about the portion of the week. And it gives him some Talmudic directions about what to say. And so it's pretty remarkable that this one entity can actually have such a wide range. So it's pretty amazing. But it does beg the question: what data does this come from. And it is the case that -- it's we know that it's not the what's called the standard web crawl. It's not merely the half a terabyte or terabyte that is scooped up from sources, but we don't know what those sources are and we don't know what the medical sources are. And that could be potentially an issue because, for example if it had a lot of Southeast Asians in it, maybe the practice of medicine looks a little bit different in Southeast Asia than in Boston. And right now, most of the large language models coming out of Google and Microsoft -- we don't really know all the data that's going into it. I think that, right now, there's so much of a thrill that that's okay. I anticipate that this will be one of the points of contention in the medical establishment about...given the fact it's very hard to evaluate something that can opine about anything in medicine. What's the data source?

David Shaywitz [00:26:37] So let me ask you another question. One more question that I have. Maybe we'll get to one or two, hopefully from the zoom, and then we'll do the books. So I was going to ask, how will the practice of medicine change, but I'm going to ask instead the active case: how is the practice of medicine changing from what you've seen, from the folks you're interacting with and the practice of medical research to the extent that it is or the changes that you're starting to see as a result of large language models like GPT-4 Or do you feel like people are still just beginning to grok it?

Zak Kohane [00:27:17] So. I could tell you about things that I learned firsthand from asking people. But what's interesting is, as a result of writing this book -- even though no one's actually read it, because no one can get their hands on it -- nonetheless, the notoriety has had a bunch of doctors reach out to me and tell me things. There's an immunologist out there who told me that he was running GPT-3 -- Chat-GPT -- to get second opinions for his patients around their workups. That's pretty astounding. And this is not --I had now an extensive conversation -- he's technologically naive, to be generous, but that's the way he's using it. We also know that, because people tweeted about it, that they're using it for the administrivia of health care, for getting prior authorization. If you're a doctor and you want something to happen that's going to cost the insurance company a lot of money, you have to write a letter explaining why your patient and these other patients deserve that extra step. They just take the text, the last letter they wrote about this patient, you dump it in and say, GPT, make a pre-authorization letter and it does it. And so I think the short answer is doctors who have self-awareness of their own limitations. And when they're especially at the limits of what they think they can do, want with open eyes, they're not going to shut their eyes and say, go do it with open eyes to hear this other colleague where they don't have to stand on ceremony, don't have to pay them, don't have to particularly be nice, and any hour of the day get that interaction and and very clearly on the part they hate of their job, which is the part that has to do with reimbursement workflow.

David Shaywitz [00:29:24] Zak - so the paperwork stuff -- and Peter has an extensive, a whole chapter, in fact, on the paperwork. And, you know, I understand at one level how,

you can have an automated pre-auth letter and then the payer can then have their automated denial of your pre-auth letter and they can sort of fight it out and you know, and saving the authentic amount of paperwork and the partnership with Nuance to try to make that work better. But there's a level where I listen to that and I think a little bit of that, you know, the famous comment, you know, they promised us flying cars and we're getting, you know, whatever, 160 characters or whatever [140]. And, you know, you're sort of thinking here is, you know, the superintelligence, alien intelligence that you describe, and then you can sort of let like payers and provider letters like argue with each other. Do you see some -- and I know it's to liberate the time, although they'll just make people see more patients if they have more time -- what do you see? Is there an elevated way to practice better medicine that this technology is going to enable?

Peter Lee [00:30:28] Yeah, I can jump in. By the way, on your prior authorization thing, in a review with Bill Gates a few months ago, he pointed out that, okay, so David, I'm going to write you an email. I'll write three bullet items, the GPT-4 will expand that into a nice flowery email that it'll get sent to you. Since it's from Peter Lee, you won't bother to read it. Instead, you'll have GPT boil it down to three items and generate a response. And there's something about that that I think you're getting at with prior authorization, because if there's AI on both sides, you know, why are we doing it? And I think while I think that there's a real valid point there, I think what we're aiming for in the future, I think we're always aiming for with technology and intelligence in medicine is to kind of connect across these things. So it would be great if medical research had data and visibility to every clinical account. You know, the dream of real world evidence that could be made true. It would be great if somehow, you know, reimbursements for delivering health care, you know, were seamlessly part of the unified kind of payments scheme without the evil kind of incentives on all sides. And I think that the technology, we do get closer and closer. And I think this has been a big step forward to having a technical foundation doing that now becomes a societal and business issue.

David Shaywitz [00:32:02] Peter, just to push back one little point, when we were working together, one of the issues that struck us was the last mile problem of how many technical problems -- it wasn't [an issue of] good enough technology. I mean, you were putting you were willing to throw everything that Microsoft had available at it. And it was basically these silos. We're aspiring to create the learning health system, right? That's sort of what you're describing. And the issue isn't there isn't like a big enough brain or maybe that wasn't the only issue. It was getting the data flows and getting the information in a coherent place. To what extent is this going to help solve that?

Peter Lee [00:32:39] Yeah, and actually, in that relationship with Epic and Nuance, of course, the big push in recent years has been the interoperability of healthcare and the adoption of standardized modern data formats like FHIR - fast health care interoperability resources. One of the things that's very interesting and puzzling right now is GPT-4 can translate in and out of FHIR from anything. From a transcript of a doctor patient conversation to medical vignette to prescriptions, whatever. And so you think wow, this is the ultimate translator to solve health data interoperability. And then you think about that a little bit more and you realize then why do we need data and operability at all? And so it gets very confusing to understand what the future should be here. But it does seem different that there is a tool now that seems to just be the universal translator for health information.

Zak Kohane [00:33:40] So I want to give a more elevated challenge.

David Shaywitz [00:33:44] We would expect nothing less, Zak...

Zak Kohane [00:33:46] Thank you, by the way - forget it, I'm not going to rise to that to that bait. But back in the antediluvian era, when I was in training, after the end of every clinic, when we'd all seen patients, we'd get together and all of us would go over our patients. And invariably, I'd say -- let's say we've seen 30 patients -- invariably four or five we'd actually-- because of what someone else said -- we changed the management. That went away. There's no time for it. I'm convinced that there will be almost instantaneously, doctors will have that version. You know, did you think about this? Even the best doctors will not remember everything all the time. And this will actually serve as that, what we used to do at a post clinic conference and turn it into an intra clinic conference with the generative A.I.. I think that's going to change because right now every doctor is essentially solo without a sounding board, without a critical evaluation by a non-judgmental party.

Carey Goldberg [00:35:08] I heard an interesting anecdote just yesterday, actually, from a Mass General physician who had a patient who had had severe hand pain for like eight weeks. And so her primary care physician sent her to an orthopedic specialist and to a rheumatologist. And she got so frustrated that -- she was a nurse -- that she put her own clinical data into I think it was Chat-GPT or GPT-4 and among the other possible differential diagnoses that it gave her was the list of medications that could be causing the pain. And she was on one of them and she stopped the medication and the pain went away. Now, why her other physicians missed that isn't exactly clear.

David Shaywitz [00:35:52] But but it really does speak because it sounds like GPT had the time to thoughtfully go over each of the things. Let me ask one question from the zoom, and then I want to make sure you have time to sign books for all the many people that are here. So one of the questions somebody asked, they brought up today, as you may have heard, Geoff Hinton left, resignation from Google -- so let me take half a step back. So there's a catastrophization school of thought that basically says AI is going to end the world: react.

Carey Goldberg [00:36:28] Peter that's all yours.

Peter Lee [00:36:31] So all technologies are dual edged. And I think that the thing that's interesting about artificial intelligence is it might in its ultimate form, be incredibly powerful, but not just powerful, but also be so highly democratized. In other words, it's a technology that might be, at the same time powerful and available to every person on the planet. And so the question is, is the world a safe place to live if every person on the planet has ultimate technological power in their hands? And so I think it's unclear what the answer to that question is. And because it's unclear, if we think there's a chance that will achieve that ultimate psychological power on the path we're on with AI, then it is very important to understand the path that we're on and how to manage that path that we're on right now. For us at Microsoft and OpenAI, we believe very strongly that that we are on that path and that our best hope for understanding how to manage and control it is to, as Sam Altman would say, to approach it gradually and at every step of the way, try to encourage as much thoughtfulness as possible. And so that's what led to GPT-2, GPT-3, Dall-E 2, GitHub, Copilot and so on. But I think it would be foolish for us not to take it very seriously, but at the same time not overestimate so much what where we are today, that GPT-4 is the next kind of point on the spectrum. And it has been, you know, I think we understand increasingly what its limits are. And so if you asked me, do I think GPT-4 will cause the extinction of humanity? Absolutely not.

Carey Goldberg [00:38:36] Sleep easy.

Peter Lee [00:38:39] Sleep easy.

Zak Kohane [00:38:44] What's your reaction to catastrophization?

Carey Goldberg [00:38:47] Oh, catastrophization. I want governmental bodies, I want international governmental bodies to take this on yesterday. And Peter travels in those circles so you can say more. But it seems like the best suggestion I've heard is Sam Altman saying we need something like the IAEA which regulates nuclear power. Yeah, so it seems like a no brainer.

Peter Lee [00:39:13] The thing that's very good is -- I think that there's been tremendous frustration, especially at OpenAI, but also at Microsoft for the past several years. We felt like the world wasn't taking the risks seriously enough. And there were repeated calls in public, particularly by Sam Altman, over the years. The thing that's very good is that people are paying attention now and people are taking this very seriously. And with respect to Geoff Hinton, I think it's commendable and very good that he's decided to devote what remains of his career and his life to thinking hard and helping the world prepare for this. And he's making the decision, which I agree with, that you'll have more success doing that not as an employee of Google, rather than as an employee of Google. So I think that, you know, it's a good thing.

David Shaywitz [00:40:02] Peter, do you think that -- I want to get Zak's response, but before that -- I do want to ask you, Peter, do you think that Google's concern about regulation and about ethics actually slowed down their development of their large language model platform?

Peter Lee [00:40:22] I have no idea.

David Shaywitz [00:40:24] You're not taking that one. Okay, Zak. All that media training is for something.

Peter Lee [00:40:30] When there's an existential threat to a company's core business, because -- so I joined Microsoft in 2010, and the world had just declared the post-PC era, iPads would take over everything, and Microsoft was having its own existential crisis. And so I'm pretty sure Microsoft was not thinking straight at all about much. I think it's very difficult for extremely large organizations to manage these crises well. Google today is a much better company, probably than Microsoft was in 2010, so they'll probably manage whatever they're going through now much better than Microsoft did. But I think it's..

David Shaywitz [00:41:15] But it's just so striking, they developed a lot of the technology and -- alright, Zak: maybe we'll just end with your thoughts on A.I. for good and evil.

Zak Kohane [00:41:24] Yeah.

David Shaywitz [00:41:24] Discuss.

Zak Kohane [00:41:27] Especially that a former Orthodox Jew seems to be leading the charge on the catastrophization.

Carey Goldberg [00:41:34] Is that Yudkowsky?

Zak Kohane [00:41:35] Yeah, Yudkowsky. But we are told in the Talmud about how to react to those things, and it's not by inciting controversy. So I will do the job that I assigned myself in this book, which is to represent annoyingly the medical establishment and say, we have this opportunity here today with GPT-4. And we know -- those of us who are in medicine know where our weaknesses are, and they are many. And patients who have experienced the health care system know they are many. And so the question is how can we best use this technology safely, today? I am convinced that not only can we use it safely under the right auspices with the right controls to actually truly use this tired word, transform health care both on the patient side and the doctor side. But I think we can also advance research. Perhaps not as fast as we can advance the healthcare changes, but we can advance it. And so I don't see a catastrophe there. I do think that some of the financial incentives in the health care system are going to actually push back against AI. So, for example, under a fee for service system in the United States, the more you do to a patient, the more you get paid, which means that an agent, whether it's a human being, a regulatory agency or a computer program that decides that it might be better off for the patient to have less procedures -- that's a financial threat to someone, and I think we're going to see a big push back from that corner under the guise of larger concerns around a catastrophe. And what I would argue is all of us who are in health care need to remind ourselves why we're in this business. Let's make sure this technology is used for the best purposes. It could be used, as Peter says, as a double edged sword for the wrong reasons. But we were -- when we were given a scalpel, we decide not to go around and hack and slash people. We decided to use them for good. And so this is exactly what we have to do now. And so I think the other questions are interesting and important questions for even the near future. But for today, I think the path is clear.

David Shaywitz [00:44:10] So a couple of things, but I'm going to wrap up here. I thought there were so many good points that you just made and were in the book. I think the idea that entrenched interests may sort of exaggerate anxieties about AI for the purpose of maintaining the status quo. I think you didn't -- you put it a little bit more politically, but I think that's exactly the concern. I also really loved how in your book there is such a tendency for these things -- I know you said, yes, there should always be a person in the loop -- but you talked about, I think, some many times, people say first, do no harm, first do no harm, and it's a precautionary principle that's often misunderstood, where it's like whoa, let's not do anything. And you really highlight the harm that's being done by not utilizing an available, powerful, enabling technology, and that so often gets overlooked. And so I really thought it was incredible, I think, all three authors consistent focus on the patient really from your, you know, annual meetings Zak that have been always focused on the patient the search for rare diseases has always been your focus Carey and the focus of your writing and with Peter with your work that we were doing on COVID and really trying to leverage technology to really help in an impactful way, individual patients. And I think that essential humanity and compassion drives the book as well as an incredibly erudite understanding of the at least the issues in the technology and how it might be working. So I feel so grateful that the technology and the description of the book arrived mediated through you three. And I feel like we were all so fortunate and the world is fortunate to have the benefit of your insight and erudition. So thank you for being here.