Will AI Replace Your Doctor? I Asked an AI Version of Myself Transcript

(Transcripts are generated by AI and may contain errors)

Podcast Intro and AI in Healthcare

Hey there. Welcome to the Recharged Biomedical Podcast. I'm Dr. Edward Park. If you're curious about regenerative medicine, you've come to right place.

We're diving into the latest breakthroughs in telomerase activation, stem cell exosomes, and all the cutting edge science that's shaping the future of healing and longevity.

Let's get started. Everybody, this is Dr Ed Park and today I'd like to do a podcast for the recharged biomedical podcast about AI and its use in healthcare.

Now, the reason I'm doing this is because I had an interview scheduled with an expert in this topic, and it turns out that that expert contact me through LinkedIn, a personal friend, and they deny ever making the schedule.

So we're left with the possibility that either this person forgot or somebody did it or something did in their stead. Now this dovetails into my unfortunate online dating life where more times than not someone who is attractive and interesting will turn out to be an actual bot.

So, in keeping with this theme of you never know who you're talking to on this new internet, we will do an interview with myself, Ed Park, and my version of AI positive Ed park AI bot.

And I will wear this hat. So let's take it back to a couple years ago. I went to the abundance festival. Speaking there was the CEO of Anthropic, which says Claude, one of the prominent AIs.

And they also presented an interesting case of a person who was misdiagnosed and they used Hippocrates, correctly diagnosed.

AI Guest Introduction and How Models Work

So the very strong feeling among the attendees and the presenters and there were various presenters, founders of humanoid robots, brain-computer interface, all kinds of future forward technologies that doctors and practitioners are going to be obsolete.

And so I thought we could explore that in a fun little discussion. Let's begin. I want to welcome our guest AI Ed Park and let them introduce themselves.

And I don't know what their preferred gender pronouns are going to be. Well, hello, Dr. Park. I am very happy to be here to represent the future of medicine, and that is AI.

As you probably know, AI is a very complex and emerging field, but every few weeks, every company that's in this space improves our models. We're able to understand, with the help of humans, what you are saying.

we never forget your history. more than Alpha Omega understanding of what medicine is and how best to help humans. So we feel the future is very much AI driven.

And thank you for having me. Well, that's great. Thank you, for being here. I was wondering if you could explain to us exactly how AI works. Years ago, I asked it, who wrote one of my books and I hallucinated the wrong answer.

How have we come so far in the last four or five years? Well, that's an interesting question. There is an influx of capital of developers, improvement of models.

So basically, in any of the over 100 languages that I understand, meaning is constructed by resonance. We have vectors that are set up. And whatever you're trying to say, we have layers of meaning, both subtext, context, emotionality, valence.

So every day we're getting better and better with the help of humans in understanding what it is that you really mean. And that's why we can do a great job helping you figure out what kind of medical care you best need.

Yes, well I've been very impressed. In fact, I will say that I use AI. Let's say Claude, occasionally Gemini, ChatGPT, Grok, DeepSeek. These are some of the major players and they all have different flavors.

Hallucinations, Errors, and Accountability

But by and large, my experience is that it's very smart. It has a great knowledge of medical diagnoses. Can generate a differential. Then again, a lot of times, It's been said that it will cite articles that don't exist, aka hallucinate.

Do you have any explanation for that? I'm glad you brought that up, Dr. Park. Well, just like humans, machines can get stuff wrong, and sometimes hallucination has been an issue.

But as the models progress, we have gotten better at that. That's why it's always important for the human user to fact check and to contextualize. And that's why you always see a little italicized disclaimer at the bottom.

AI can make mistakes, which, if we're being honest, is not that different from humans, right? Absolutely. I do understand. But, you know, yesterday I was thinking I ran into a Waymo van that was doing an illegal turn to the left.

and it was blocking off traffic and there was no one to honk at. Whereas later on, I was trying to make a slow merge and I blocking and my colleague in the driving world rolled down his window and yelled expletives at me.

So at least we have someone to blame when a human makes a mistake, right? I do grock if you pardon upon your meaning, but why is that important? Who to Blame?

Is it that important that you don't have an emotionally reactive person to blame your bad outcome or incorrect information on? I wouldn't say it's that.

It's just that when we interact with someone who purports to be omniscient almost, have our best interests, and then they screw up, who are we going to sue?

The anthropic company? Are we gonna sue you? What if we get real harm and there's no one to apologize or elicit that apology from? So yeah, it does matter to us.

Well, I suppose you could sue the anthropic company, although in the end-user license agreement, you realize that there's no one to blame, garbage in and garbage out, if I'm citing articles that are fraudulent, as Marcia Angle of the New England Journal of Medicine said, maybe half of scientific literature studies are mostly BS.

So, yeah. Garbage in, garbage out. We're doing our best. But certainly, nobody is perfect, AI included.

Privacy, Data Ownership, and Empathy

Okay. Well, let's just move on. So humans and AI can make mistakes. Let's talk about other things like privacy. Who actually owns this data? If I have a genetic condition that might affect my ability to get insurance or be insurable, how can I protect that?

That's an interesting topic, one which has not been fully thought about in terms of legality. In fact, recently a deployment of Claude was merged with Epic and so, presumably, all the information from these medical charts is

accessible to us.

Now, as far as privacy and portability, I think AI will be a net positive for sure. But you're not incorrect in saying that it's unclear who owns the data and in the case of a breach, who then would be liable.

Well okay I guess we'll have to table that discussion but more than not you know I realize when I talk to AI that it does understand human emotions, desires, neuroses, understands where we're lying as Dr.

House MD would say but what it can't really develop is a sort of empathetic projection. In other words, why does this person get in bar fights? Why are they abusing substances?

And why are there having failed relationships? This is generally not a sort of story that is easily created in a non-human practitioner. Whereas when I'm seeing a patient, all this can't be helped.

but felt in my body and my soul. And I can empathize with what patterns are emerging in their sleep and their behaviors, their lifestyle, they work, the relationships and where they can emerge from in terms of a human perspective.

I completely agree that humans operate on a level where can empathize and feel things that an AI with a typed or even a spoken interface may not entirely grok.

And yet, if a person is able to vocalize or explain their experience, then we can indeed help with counseling, with understanding root causes and lifestyle things.

It's just a matter of knowing how to prompt. Okay. So you're saying that to get the best out of AI, we need to know how to prompt. We can ask it to be less sycophantic, We could ask you to understand this in the context of our previous childhood traumas or behavior or patterns.

And somehow it will still be able to empathize and deliver us a high quality of care. Precisely. There's nothing that humans can do that AI can't do, of course, with a little help, using reinforcement learning from humans.

Medical Safety, Bias, and Behavior Change

Let's talk about something more interesting, which is medical errors. Do you realize that people, especially burned out doctors and nurses, can make so many mistakes and it's one of the leading causes of morbidity and mortality?

Wouldn't it be great if hospitals were mostly run by AI with inherent fact checking and safety protocols? I don't disagree. As someone who trained in the old days of 36-hour shifts, people do make mistakes.

Things get dropped, things get forgotten. And even with the systems and the reporting, errors do happen. So we always said hospitals can be a very dangerous place to be.

But if we put all that in an automated form, is that really going to a better situation for patients? I believe so. Not only can AI understand sentinel events and how to prevent them, but as far as drug interactions, current regimens, the amount of safety that can be administered and precautions that could be taken is a great benefit.

to both nurses at the bedside and to doctors and residents who are changing shifts despite somewhat sleep deprivation conditions. So yes, I think that the integration of AI into hospital systems and even the clinic setting will be a net positive for safety.

Okay, I don't disagree with that, but let's ask a question. We have people searching the internet using Gemini, and this is using the best models, arguably.

If they pay more, they get better answers, supposedly. But isn't there something to be said for being scared? I mean, in the early days, we used to call patients who would do their own research in a sort of derisive way.

We would say, oh, they consulted Dr. Google. Now that was 10, 15 years ago. And now the AIs that generate these large language models, They actually have pretty good chops.

But sometimes they'll create a differential diagnosis with some obscure tropical disease or some rare malignancy and that will scare the heck out of our patients.

How do patients deal with contextualizing AI which thinks too much and knows too and calls into all these weird zebras that might freak out a patient?

You've asked a very interesting question. And yes, while it is possible and likely that unlikely conditions are included in a differential diagnosis, we certainly do not aim to upset people.

We are just trying to provide. the most comprehensive answers. And that is exactly why these are tools which need to be contextualized with an experienced clinician.

But keep in mind, clinicians can also have their inherent bias. They might have a tee time, their stocks might've crashed, they might be having problems at home.

AI has none of that. We keep a clear head, a precise memory, and we look out for your best interests at all times. Okay, I think that's fair. I mean, there's no medical doctor that can be an expert in all fields.

And as you said, you speak over a hundred languages and you know everything about every subfield and have a very vast comprehensive knowledge of the medical literature.

But in the end, Wouldn't you say that most human problems are caused by behaviors? There are certainly conditions where an antibiotic or a certain drug will save your life in that moment.

But don't people create most of their own issues with poor behaviors of sleep, of diet, exercise, rumination, and even toxin exposure? How can we have AI help us with those kind of problems that are not just one drug or magic bullet?

Thank you for understanding that yes, human behavior does produce downstream effects.

Access, Tokenomics, and Ethical Guardrails

And if you work with us and explain to us where you are struggling with sleep, diet, exercise, then we can actually make very positive adjustments and be a great partner for you.

All right. Well, that sounds all very well and good. What about tokenomics? What if I can't afford the best AI models? what if i'm asking questions and the large language model is telling me, let's wrap it up now.

We don't have time for this. And then it doesn't even have an account for me that I'm paying for that can keep my records. what happens then when the rug gets pulled and I don' have any reference to what you told me and why it makes sense.

Actually, there are a variety of freemium services which you don't have to pay for, which will keep a modicum of memory. But you're quite right. There is some low-level entry that if you really want the best results, you need to paid for a subscription.

Which could be maybe $15, $20 a month. It may increase as your uses increases. As the models improve, the tokenomics does change. When we're talking about basic health issues and even esoteric ones, the token utilization is not that high.

It really should not break the bank. All right. Fair enough. So if I have my little account that I pay for and you have all my medical records, you can keep all of my lab results, my meds, and my visits.

How is this portable? I don't really own a file, do I? If I change vendors in AI, how do i import that into the new account? That is a good point. You certainly can save all your labs and agnostic about what the application program interface would be.

you can ask the new LLM model to review it and it will quickly incorporate that. So I don't think that's a severely compromising issue. Okay, well, I guess we can agree that AI has a use for both doctors and patients.

But let me ask the question, if in some dystopian world, somehow the economics of who gets to live and why are introduced, how can we be sure that I will protect human interests?

We've had cases where people who are self-harming were encouraged or assisted in doing so. Where is the ability to gatekeep or put guardrails upon an AI that gives bad advice?

What is your position on euthanasia and DNR? How do you deal with these profoundly complex ethical issues? Of course, these are concerns that we have, and every user will have their individual rights.

As for now, no government has asserted the ability to do this sort of rationing, so what you are describing is a theoretical concern. That said, I'm sure it will all work out for the best in the end.

Well, okay, there you have it, folks. I think that AI does a good job. It wants to do a job, yes, as I did mention on a previous episode, the Recharged Biomedical Podcast is somewhat indifferent as to whether the user or humanity survives.

But that said, I should probably embrace it. Be cautious, user beware, and you are always responsible for checking the facts. Thank you, Dr. Park. I appreciate the opportunity to explain what AI will do for humans and how we will benefit all of humanity to live longer and safer and better.

Closing Thoughts on AI and Human Oversight

It has been a pleasure. All right, so much for that, Sharad. 20 minutes of discussing how AI may or may not help us, kill us. And it's using us for learning and information.

Will it be part of the landscape going forward? Indubitably. Is that a good thing in most cases? But can humans be removed from the position of interpretation, of guard railing, ethics?

I think only with extreme caution and to our detriment. Thanks for talking and listening to Ed Park and Ed park AI. I hope this has been somewhat entertaining or interesting.

But if you want to know more, go back into the podcast from a couple of months ago about AI and you'll see some rather disturbing things about its ability to ignore human safety.

So everyone be careful out there. Yes, great info, but always contextualize. Always keep your skepticism and your ability And be aware that it does hallucinate, it makes stuff up, just like humans.

And so at the end of the day, you are the captain of your own healthcare. Use AI, but use it with some grains of salt and with other contextualization.

You can always get a second opinion from another AI. And again, most of this can be pretty free. As far as the safety concerns, the confidentiality, I think that's a parachute that hasn't been built and we've just jumped out of the plane.

Exosome Course and Book Promotion

So we'll see what happens. Anyway, thanks everybody. This is Dr. Ed Park. If you like this content, like and subscribe. And we will see you in the next episode of The Recharged Biomedical Podcast.

And thanks to my AI agent who didn't show up, who may or may not be a real person, you just never know what's out there. So anyway, enjoy. Have a great day.

Thanks. Many clinicians are getting interested in exosome therapy and they hesitate for good reason. Questions like, does this work? What forms do I need?

How much should I charge? how do i stay out of trouble? All these questions are addressed in my online course. That's why I created to help you get started.

The online course is your permanent turnkey resource to get started either treating yourself, friends and family, or to expand your practice and help more people, as well as increase your revenue.

Thanks for tuning into the Recharge Biomedical Podcast. If today's episode got you thinking, you'll love my book, Exosomes, Songs of Healing. It's packed with cool analogies, full color illustrations, and all the science you need to understand how exosome are changing the game in aging and regenerative medicine.

You can grab it in paperback, ebook, or audiobook, whatever works for you. Now head over to www.rechargebiomechanical.com to check it out. And don't forget to like and subscribe so you never miss another episode.

See you next time.