



Diving Deeper into AI: Exploring the Potential and Challenges

June 2024

Erika Whitmore:

I'm really excited about today's podcast. I think we're going to dive a little bit deeper on AI and what that can mean to our audience. So I'm very excited to have Steve Shu here with us today with Superfocus.ai. And Steve, if it's okay with you, I'll let you introduce yourself.

Steve Hsu:

Sure. I'm a founder of Superfocus.ai, which is a tech company that builds AIs for enterprise. So those are AIs which use large language models. However, they are embedded in a software architecture that makes them reliable. They do not hallucinate. A little bit more about me. I'm also a professor of theoretical physics and computational science at Michigan State University, and I am a multiple deep tech founder. I've founded, or co-founded I think four or five different tech companies.

Erika Whitmore:

Wow, Steve. So I feel like your kickoff is, that's the missing link with AI. Now we're going to jump right in, but tell me more about that because I think that's one of the hugest, if not the biggest issue with AI, right? Is the hallucinations.

Steve Hsu:

Yes. So just to give you a little historical background, one of my other startups, which was founded back in 2017, had as early investors, some of the founders of OpenAI. I've known those guys for a long time and have been following the progression of GPT since all the way back to GPT one. And so I knew in advance that AI based on that architecture would have what's called a hallucination problem. And what happens is that because they train it on trillions of words of data, mostly scraped from the internet, it is good at responding in a plausible way. It understands human language and it will respond the way it has seen the data respond to similar prompts in its training, but it won't necessarily know facts, specific facts. So if you want your AI, let's suppose your AI is handling some task for you and you're a travel agency or an airline and someone says, I'd like a flight that lands direct flight, which lands in Paris before 6:00 PM, what are my options?

The model is perfectly happy to give you a response, which is consistent with all the data it has seen and it's seen many flight itineraries. However, it won't necessarily know whether that United flight to Paris is real or it only flew five years ago, or it was just a made up itinerary and it'll just give you an answer and it's perfectly happy to give an answer that's plausible, but it may not be the true factual answer. And that problem is now widely recognized as hallucination. So we knew this was coming well before chat GPT was launched. In fact,

our company was founded before chat GPT was launched. We knew this was coming. We knew GPT, this technology was getting really good, but we knew it had this hallucination problem and being experienced in enterprise software and startups, we knew that there was no way this kind of stuff was going to be deployed in the enterprise unless it was made more reliable.

And so we developed a bunch of techniques that could constrain the AI. In fact, the way that we describe it is we attach a memory to the large language model and the facts that are in the memory, the facts, the policies, the procedures that are in the memory are specified by our clients. So our clients can take the training data they use to train a human agent, travel agent or customer support agent. That same training information that they use for humans, we can embed as the memory of our AI and then our AI will only answer questions using that information.

Erika Whitmore:

That is fascinating. And I looked, and it does look like you work with several different industries. Is that correct?

Steve Hsu:

Yes. The startup that I mentioned, which some OpenAI founders are investors in is called genomic prediction. And that startup is the first to be able to genotype embryos that are produced in vitro fertilization IVF, and we use the genetic data from those embryos to be able to predict disease risks and health risks to the eventual children that are born. And so that's a very different industry. However, it is also an application of AI.

Erika Whitmore:

So the example that you gave with training the travel agents makes perfect sense to me, but in genomic prediction, how does that work? Because where I'm having trouble connecting the dots is the hallucinations to the accuracy. I'm an auditor by trade, so accuracy is super important. That makes sense in a typical business. But in a high growth startup like that, how does that work?

Steve Hsu:

If you have a webpage and all companies have webpages and you have an FAQ or something on your webpage rather than have a static FAQ, which nobody ever reads, it's very easy now to build an AI that can answer any question which is related to the FAQ information and in such a way that humans can ask the same question, same semantic content in the question in many, many different ways. The old style brittle, what we used to call flow bot AIs or chatbots can only answer the question if it's phrased in the proper way, otherwise it doesn't really understand what you're saying. Whereas the new language models, you can say it in any way and it will understand what you're saying and then it'll properly answer the question. So it's another way of communicating with customers or doing after sale support of your customers in a much more powerful, effective way. So anybody who has that kind of problem or even has say an HR problem or something like that, where your own employees are always asking questions about how many days of leave do I have left, or do I get parental leave if this happens, et cetera, et cetera. Those are simple queries that can be answered now by the AI, but in a reliable way, so according to your actual company policies rather than some thing that it saw on the internet.

So I think almost any company that has tasks like this that are currently handled by humans, they could instead be handled by AIs in the future.

Erika Whitmore:

That makes a lot of sense. So maybe switching gears just a little bit, so you've kind of shared the difference between the open AI or the AI that produces, at the end of the day hallucinations or inaccurate

information because of the way that it gathers information and then the difference in terms of what your company does. What do you think is possible with AI and the timeline? If you listen to a lot of different newscasts and things, obviously AI, a lot of companies are investing a lot of money in it, but who's going to be able to capitalize off of that? Is it really going to make their companies more efficient, et cetera? Those are some of the questions investors have, right? So can you talk a little bit about just what's possible and then also just from your perspective timeline?

Steve Hsu:

Yeah. This is the question that's on everyone's mind. Every CEO, every board member, every investor wants to know, and of course it's difficult to predict. Let me first say something which is not speculative at all, which is that if you take the current state-of-the-art language models like GPT-4 or some of the other ones that have been built by Google or by Facebook or by meta or by Philanthropic for example, or Mistral, these leading edge foundation models already understand human language very, very well. And so when you couple them to, in our case, an attached memory or some attached set of goal or planning capabilities, so that's another set of things you could attach to the LLM, you can get something that even though a computer scientist or philosopher might not call it a GI, artificial general intelligence, it will seem to the average person as if it is.

So in other words, you will definitely be able to have an assistant on your phone which has looked at your emails and your calendar and can answer questions and even take some simple actions on your behalf. I think that will definitely happen. There are many, many companies working on that kind of technology right now, and it'll become a reality within a few years. I think many, many human tasks are automatable, like starting with customer support. So we gave much of examples like that, but travel agents, a lot of what the travel agent does is automatable. Now, I wouldn't go so far as to say accounting could be automatable or financial analysis. I think that's currently out of reach. So currently there are no fully reliable AI systems that can then fully replace say, a consultant at McKinsey or a financial analyst at Goldman. But I think you can get...

Erika Whitmore:

To a heck of a good start though, Steve, at least that's from my experience.

Steve Hsu:

So you can definitely, for example, assist a human that's performing those tasks and make that human say 50% more efficient. I think that's for sure possible. Another area I might mention, which I think is highly automatable, is if you have ever done telemedicine where you call, you have like you say your child is sick and you call in and you're talking to say a nurse practitioner on the phone, or you might even be talking to a general practitioner on the phone, a lot of what's done, at least in terms of telemedicine is automatable. It's actually pretty similar to customer support in the sense that you can have embedded in the memory of your AI a list, a sort of checklist of things and procedures that the doctors and nurses go through and they then suggest, oh, just wait a day. If it doesn't get worse, call me. You're fine, but otherwise call me back. Things like that, arriving at that kind of conclusion, the AIs are actually going to be perfectly fine as good as humans at doing that kind of thing.

Erika Whitmore:

That makes a lot of sense. Yeah, it's interesting because my flight was, I was flying from Salt Lake City to Denver last night. My flight was significantly delayed, and I am almost certain whoever helped me change my flight was not human. And it wasn't that the task didn't get accomplished, it was that you could just kind of tell in the tone, I did it all over text anyways. As you were talking about travel agents and things like that, I am pretty sure that's how that change in flight got accomplished. It was a little rocky, but it got done. So my guess is in a couple years it'll be even smoother, and I probably won't realize it's not a human

Steve Hsu:

Being an insider in this industry. I will tell you the thing you dealt with, and this is especially likely to me, given that it was over text, could have been AI, but it wouldn't have incorporated the modern power of large language models. That was probably a rather relatively brittle system, which is designed so that if you have a mass cancellation that it then allows quick rebooking by the people that were on the canceled flight. That kind of thing is a little more structured if then kind of statements than a full blown LLM AI, which the LLM AI could have sort of chatted with you and you could have said something like, well, it's actually better for me to stay an extra day here in Salt Lake City. I don't have to work tomorrow and I can blah, blah, blah. And the model then would know, it might just come back with, okay, how about if I schedule for another flight tomorrow then and the more brittle system today wouldn't be able to do that. But we build systems that can do what I just described.

Erika Whitmore:

So I love that Steve, because that's a tangible example. So do you have another example like that of here's how things kind of are today, and yes, they use some type of AI, but this is what I envision in the future in another either a service industry or something that would impact quite a few people?

Steve Hsu:

Yeah, so I'll give you an example. So we work with a very big consumer electronics manufacturer that makes smart TVs and we've built an AI that can troubleshoot problems people have with a TV across hundreds of models that they've sold over the last 20 years or so, the actual knowledge base that their human agents have to use or navigate to troubleshoot all these problems. These problems could be things like it's saying it's not connecting to Wi-Fi or my Netflix app keeps crashing, or there's weird static on the screen. There's so many different problems that you can have with your TV and then you might call them to get assistance. We now have an AI that can troubleshoot all of those problems and is much better than even the 90th percentile human agent that's working for them. And by the way, I didn't mention this. The other problem that we've solved at the startup is how to do all this inference with the language model and the non-hallucination in the memory and stuff, how to do all of that inference in under two to three seconds so that if you're using the voice mode and you're talking to the AI, when you finish what you have to say, you want there to be less than a two or three second gap before the AI starts talking back to you.

Otherwise it sounds weird. And we've now been able to do that. We've now compressed everything. We've sped everything up to the point where you can have a kind of natural conversation with AI. And so with this consumer products company, we've more or less demonstrated to them now that the AI is superior to their average call center worker. And so all of those jobs, there are 90% of those jobs can be more or less automated now.

Erika Whitmore:

Which makes sense because at least the research that I've done, and again, just what's out there, call centers is obviously a clear target. What do you think are the other bigger areas of human work that'll be impacted? And I don't know, at least from my perspective, this has happened over decades and hundreds of years where things evolve and we always figure it out. I don't think it's a bad thing. It is the next evolution, right?

Steve Hsu:

Absolutely. To give you an example, there's a big shortage of doctors and nurses and in the developing world, there are countries where the number of doctors or nurses per thousand people in the population is 10 times or even a hundred times lower than say in the United States. And so to have an AI that nevertheless, even in these developing countries now, a majority of people are a pretty good fraction of people have either a phone or a smartphone for them to be able to call an AI and with a medical issue like,

oh, my father has blood in his stool, what should I do? And then the AI starts having a conversation with them and then 20 minutes later, the AI either says, wait 24 hours and see if it goes away. If not, you might consider traveling to this clinic, but the nearest clinic might be six hours travel away for these people. So there are enormous benefits that are going to come from building AIs that are able to do the things we've already discussed today.

Erika Whitmore:

Yeah, absolutely. I think the last question that I've got, it's another big and broad one, Steve, so hopefully that works. Just your views on being a startup in the current environment. There's kind of a lot going on. The economy seems to be doing well, but there's pockets that are maybe struggling and doing some layoffs. What is your view on being a startup in today's environment?

Steve Hsu:

So our perspective as an AI startup is very different from other companies because we're in the hot space. It's very unusual because so many CEOs of even public companies and directors of public companies are trying to seek answers to the exact questions that you asked me during the course of this interview so that we are able to get meetings as one of the companies that's actually innovating technologically in this space. It's easy for us to get meetings with really top-level people at these companies because they themselves want to figure out what is going on. So we're in a unique situation. We're somewhat insulated from what's happening in the general economy because every company has a budget to deploy or at least investigate AI if they think it can really return value to the company that's protected, they might be laying people off, but at the same time, they're trying to figure out how to use AI for higher efficiencies.

And also of all the areas, sectors that venture capital is flowing into, AI is definitely the hottest area. So we're not really affected by what's happening in the general economy, at least at Superfocus. In some of the other companies that I've founded, you might feel it, we might be more sensitive to what's happening in the general economy. I think one thing I would just say to all Americans just broadening out the discussion a little bit is whether you feel the economy's doing well or it's struggling along, this may depend on where you are, whether you're a working-class person or a professional, et cetera. Regardless of how you feel about it, the economy itself, you should be reminded that this country's borrowing one and a half trillion dollars every year, and our kids have to pay for that money. So the Fed, if you want to think of it this way, is literally printing money for the level of prosperity that we currently have, whether you think that's good or bad, whether you think the current level of economic health is good or bad, just know that we're borrowing unprecedented amounts of money from the future from our kids that they have to pay off.

We're borrowing that money to get what we're getting right now. And that's something that being an older, I'm not a young kid anymore. I remember times when there were huge debates in Congress over whether we should run these deficits, which were much smaller than what we currently run. So whatever we have, we're paying for it in a kind of underhanded way by borrowing from the future. And I think every American should think about that.

Erika Whitmore:

No, that makes sense. That makes a lot of sense. Coming back to all the meetings that you're having with these different CEOs, regardless of the size of company and companies having budget for AI, what are some of the aha moments that some of these high level leaders have had in your meetings in terms of, oh, I can see how we can gain efficiencies there, or I can see how I can make my employees' lives better. You know what I mean? What are some of the aha moments that they've had in those meetings?

Steve Hsu:

Yeah, so great question. So number one, if you're following AI and you're interested in what is really happening, if you listen to public earnings calls with big company CEOs, those quarterly earnings calls you'll often have. For example, recently you had the CEO of Salesforce.com, which is a pretty big,

well-known company. Someone asks him like, oh, all these AI initiatives you guys are rolling out, because on the PR side, you'll see Salesforce saying like, oh, we rolled out this AI called Einstein, and it's going to be great people using Salesforce will be able to get help from Einstein all the time or whatever. But then on the earnings call where the CEO, he's legally liable for anything he says on that earnings call, some investor will dial in and say, well, what do you anticipate as revenues in the coming quarter from AI initiatives? And then you'll see the CEO back off and he'll say something like, well, there's great potential here, but as far as the next quarter or two, it's de minimis. We don't expect any impact material impact on our bottom line from this AI stuff. So you'll actually see a gap between the PR stuff, the media, the pitching of new super AI capabilities. And when he's forced to be honest, actually says about the revenue impact of what they're doing. And we would say from the Superfocus perspective, this is because they know internally the AIs that a lot of times they'll just grab some LLM from OpenAI or some other company and just try to throw it in there. Maybe they'll train it fine, tune it on some of their use cases and things like this, but generally it won't really be fully reliable. And they know that, and they're very worried about that. So there is a huge gap between practical utility of these AIs, what is actually possible unless you deploy a technology like ours and what most companies have. So that's the state of play right now. But then when you go into a meeting and you can actually show them that you've built an AI that is reliable, that doesn't hallucinate and it knows literally everything in the HR manual, for example, at Company X, then they have an aha moment where they realize, well, this is how powerful this stuff is really going to be at the end.

Erika Whitmore:

Yeah, absolutely. No, I think that's pretty common that there's a lot of companies doing a lot of things, but in terms of how it's going to impact the bottom line is yet to be determined. And it does feel like a lot of companies are going in a lot of different directions because there's a lot of good things happening. And then to your earlier point, there's maybe some not bad things happening, but maybe not as robust as the users are hoping. And I would say even at KPMG, we've got AI embedded for all of our practices, so advisory, tax and audit. And it's really interesting to see how it evolves literally over a month. So it does get better and better. But you still have the caveat, you signed a waiver that says, I understand that I'm responsible for everything that comes out of here because I need to make sure it's right.

Steve Hsu:

What you just said is exactly how I would describe the situation. Upper management is trying to roll out AI tools to make each and every KPMG person more effective and efficient. However, it's still the case that a human really has to review everything that comes out of the AI could save you time sometimes maybe it costs you time. You have to track down some really subtle error that the AI made, but that's sort of the state of play. There are very limited set of situations where the AI is on its own making important decisions or supplying information that's not going to be further reviewed by another employee of the company. So when a customer support AI is talking to the end customer, there is no human in the middle there. And so those are the use cases where, okay, the AI is fully doing this job and those are very rare right now. There are very limited situations where that has materialized, but we're sure that is possible. It's just now that technology is getting rolled out.

Erika Whitmore:

In terms of back to the timeline, you said earlier that in terms of call centers and things like that where maybe the tasks are a little bit less complicated, that'll be a reality either now or very soon, and then the rest of it will follow you think a few years after.

Steve Hsu:

I think that it's unknown. I mean, to fully replace a tax auditor with an AI, I think it's unknown how long that is going to take. I wouldn't be surprised at all if it takes another five years or more for that, something like that to happen. But for these easier jobs, less cognitively demanding, less complex set of inputs and outputs, it's already possible. It just hasn't been done at scale yet. It's just starting to be done at scale now.

Erika Whitmore:

So my other favorite question, this is a random one, Steve, that I ask anybody who's on to talk about AI is when will we be ordering things from a fast food or maybe even a fast casual restaurant with no human? I mean, I know we can do that now, but where they're literally not a human except maybe they're back cooking.

Steve Hsu:

Yeah, we can do that now. It just has not been deployed at scale. So it is possible. So for example, if we got an engagement with McDonald's and they wanted an AI that lives on the McDonald's app on your phone and you can just talk to it and get your order in with no human participation, we can definitely build that. Now, part of the issue with this particular scenario that you raised is that the gooey UIs are already really good. In other words, is it really better for me to be talking to an AI than to just navigate the app where I can just click extra large fries milkshake. I'm pretty good at navigating that gooey and just touching some buttons to get what I want. And is it actually better to talk to an AI? So in some of these cases, the interface engineering is already good enough that the gain from having the AI come along isn't necessarily that great. And I would say restaurant ordering is kind of borderline in that case. Something like someone calling in to say, my tv, this is not working, and there's a 10 step troubleshooting process that the agent has to walk you through to get your TV back, working online for that, you're not going to do that through some gooey UI. You need the actual, either you need a human or you need an AI to do it.

Erika Whitmore:

And that is a very common problem. So that in and of itself will be revolutionary. Well, Steve, any parting thoughts for our listeners?

Steve Hsu:

I think you've asked all the right questions, and I think we are in for an unprecedented pace of change. I think usually what happens is you initially have this super hyped stuff, and we've been through that in the last year and a half overhyped expectations about AI, and now people are realizing, okay, this is kind of more realistically where it is. It can help me, it can assist me, but it can't fully replace me at this point in time. But in the background, these busy beaver technologists have been continuing to push things forward and you will start to see job categories where the human is fully eliminated in favor of an AI.

Erika Whitmore:

Yeah, absolutely. Well, Steve, I can't thank you enough for joining me today. I think this is going to be probably one of our more popular podcasts. So thank you so much for taking the time out of your day, and we hope you have a good one and we'd love to have you back soon.

Steve Hsu:

It's my pleasure. And for your listeners who obviously like podcasts, I'll just mention that I have a podcast called Manifold, M-A-N-I-F-O-L-D. If you google my name, you'll find it right away, manifold. And in that podcast I interview AI researchers, technologists, business people, startup entrepreneurs. It probably would be something that your audience is interested in.

Erika Whitmore:

Absolutely. I am going to go listen to some myself. So awesome, Steve. Well, thank you so much. Have a great day, and we will talk to you soon.

Steve Hsu:

Thank you. It's been a pleasure.

Erika Whitmore:

Thanks, Steve.

Contact



Erika Whitmore
T: 303-382-7717
E: elommen@kpmg.com

About the Privately Speaking podcast series

Brought to you by our Private Enterprise practice, Privately Speaking features conversations with business leaders to explore the opportunities and challenges of your business journey.

About KPMG Private Enterprise

The world has changed. Thriving in today's dynamic marketplace requires deep insight into the new market fundamentals. And it takes a clear understanding of the rapidly evolving regulatory and tax landscape, investor trends, supply and demand disruptions, and changing customer expectations. KPMG Private Enterprise advisers understand the trends influencing the new reality and are hard at work helping private companies uncover and capture new opportunities in their markets. Learn more at kpmg.us/industries/private-enterprise.

The views and opinions expressed herein are those of the interviewees and do not necessarily represent the views and opinions of KPMG LLP.

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us:



kpmg.com

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2024 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization. USCS017510-2A