

OpenAI's Sam Altman Talks ChatGPT And How Artificial General Intelligence Can 'Break Capitalism'

An edited version of Interview with CEO of of OpenAI, Sam Altman. Forbes spoke to the recently press-shy investor and entrepreneur about ChatGPT, artificial general intelligence and whether his AI tools pose a threat to Google Search.

It feels to me like we are at an inflection point with the popularity of ChatGPT, the push to monetize it and all this excitement around the partnership with Microsoft. From your standpoint, where does OpenAI feel like it is in its journey? And how would you describe the inflection point?

SAM ALTMAN: It's definitely an exciting time. But my hope is that it's still extremely early. Really this is going to be a continual exponential path of improvement of the technology and the positive impact it has on society. We could have said the same thing at the GPT-3 launch or at the DALL-E launch. We're saying it now [with ChatGPT]. I think we could say it again later. Now, we may be wrong, we may well hit a stumbling block we haven't or don't expect. But I think there's a real chance that we actually have figured out something significant here and this paradigm will take us very, very far.

Were you surprised by the response to ChatGPT?

I wanted to do it because I thought it was going to work. So, I'm surprised somewhat by the magnitude. But I was hoping and expecting people were going to really love it.

[OpenAI President] Greg Brockman told me that the team wasn't even sure it was worth launching. So not everyone felt that way.

There's a long history of the team not being as excited about trying to ship things. And we just say, "Let's just try it. Let's just try it and see what happens." This one, I pushed hard for this one. I really thought it was gonna work.

You've said in the past you think people might be surprised about how really ChatGPT came together or is run. What would you say is misunderstood?

So, one of the things is that the base model for ChatGPT had been in the API for a long time, you know, like 10 months, or whatever. [Editor's note: ChatGPT is an updated version of model GPT-3, first released as an API in 2020.] And I think one of the surprising things is, if you do a little bit of fine tuning to get [the model] to be helpful in a particular way, and figure out the right interaction paradigm, then you can get this. It's not actually fundamentally new technology that made this have a moment. It was these other things. And I think that is not well understood. Like, a lot of people still just don't believe us, and they assume this must be GPT-4.

With the froth in the entire AI ecosystem, is that a rising tide that is helpful for you? Or does it create noise that makes your job more complicated?

Both. Definitely both.

Do you think that there is a real ecosystem forming here, where companies besides OpenAI are doing important work?

Yeah, I think this is way too big for one company. And actually, I am hopeful that there is a real ecosystem here. I think that's much better. I think there should be multiple AGIs [artificial general intelligences] in the world at some point. So I really welcome that.

Do you see any parallels between where the AI market is today and, say, the emergence of cloud computing, search engines or other technologies?

Look, I think there are always parallels. And then there are always things that are a little bit idiosyncratic. And the mistake that most people make is to talk way too much about the similarities, and not about the very subtle nuances that make them different. And so it's super easy and understandable to talk about OpenAI as like, "Ah, yes, this is going to be just like the cloud computing battles. And there's going to be several of these platforms, and you'll just use one as an API." But there are a bunch of things about it that are also super different, and there are going to be very different feature choices that people make. The clouds are quite different in some ways, but you put something up, and it gets served. I think there will be much more of a spread between the various AI offerings.

People are wondering if ChatGPT replaces the traditional search engine, like Google Search. Is that something that motivates or excites you?

I mean, I don't think ChatGPT does [replace Search]. But I think someday, an AI system could. More than that, though, I think people are just totally missing the opportunity if you're focused on yesterday's news. I'm much more interested in thinking about what comes way beyond search. I don't remember what we did before web search, I'm sort of too young. I assume you are, too...

We had an Encyclopedia Britannica CD-ROM when I was a little kid.

Yeah, okay, exactly. There we go. I remember that, actually, exactly that. But, no one came along and said, "Oh, I'm going to make a slightly better version of the Encyclopedia Britannica on the CD-ROM at my elementary school." They're like, "Hey, actually we can just do this in a super different way." And the stuff that I'm excited about for these models is that it's not like, "Oh, how do you replace the experience of going on the web and typing in a search query," but, "What do we do that is totally different and way cooler?"

And that's something unlocked by AGI? Or does that happen before that?

Oh, no, I hope it happens very soon.

Do you feel that we are close to the goal of something like an AGI? And how would we know when that version of GPT, or whatever it is, is getting there?

I don't think we're super close to an AGI. But the question of how we would know is something I've been reflecting on a great deal recently. The one update I've had over the last five years, or however long I've been doing this — longer than that — is that it's not going to be such a crystal clear moment. It's going to be a much more gradual transition. It'll be what people call a "slow takeoff." And no one is going to agree on what the moment was when we had the AGI.

Do you see that being relevant to all of your interests beyond OpenAI? Do they all fit into an AGI theory, Worldcoin and these other companies?

Yeah, it is. That is, at least, the framework in which I think. [AGI] is the thrust that drives all my actions. Some are more direct than others, but

many that don't seem direct, still are. And then there is also the goal of getting to a world of abundance. I think energy is really important, for example, but energy is also really important to create AGI.

Greg [Brockman] has said that while OpenAI is research driven, it's not anti-capitalist. How are you navigating the wire act between being for-profit with investors who want a return and the broader goal of OpenAI?

I think capitalism is awesome. I love capitalism. Of all of the bad systems the world has, it's the best one — or the least bad one we found so far. I hope we find a way better one. And I think that if AGI really truly fully happens, I can imagine all these ways that it breaks capitalism.

We've tried to design a structure that is, as far as I know, unlike any other corporate structure out there, because we actually believe in what we're doing. If we just thought this was going to be another tech company, I'd say, "Great, I know this playbook because I've been doing it my whole career, so let's make a really big company." But if we really, truly get AGI and it breaks, we'll need something different [in company structure]. So I'm very excited for our team and our investors to do super well, but I don't think any one company should own the AI universe out there. How the profits of AGI are shared, how access to is shared and how governance is distributed, those are three questions that are going to require new thinking.

Greg walked me through the idea of a third-party API future alongside first-party products —enterprise tools, perhaps. As you productize, how do you maintain an ethos of OpenAI staying open?

I think the most important way we do that is by putting out open tools like ChatGPT. Google does not put these things out for public use. Other research labs don't do it for other reasons; there are some people who fear it's unsafe. But I really believe we need society to get a feel for this, to wrestle with it, to see the benefits, to understand the downsides. So I think the most important thing we do is to put these things out there so the world can start to understand what's coming. Of all the things I'm proud of OpenAI for, one of the biggest is that we have been able to push the Overton Window [Editor's note: a model for understanding what policies are politically acceptable to the public at a given time] on AGI in a way that I think is healthy and important — even if it's sometimes uncomfortable.

Beyond that, we want to offer increasingly powerful APIs as we are able to make them safer. We will continue to open source things like we open-sourced CLIP [Editor's note: a visual neural network released in 2021]. Open source in really what led to the image generation boom. More

recently, we open sourced Whisper and Triton [automatic speech recognition and a programming language]. So I believe it's a multi-pronged strategy of getting stuff out into the world, while balancing the risks and benefits of each particular thing.

What would you say to people who might be concerned that you're hitching your wagon to [CEO] Satya [Nadella] and Microsoft?

I would say we have carefully constructed any deals we've done with them to make sure we can still fulfill our mission. And also, Satya and Microsoft are awesome. I think they are, by far, the tech company that is most aligned with our values. And every time we've gone to them and said, "Hey, we need to do this weird thing that you're probably going to hate, because it's very different than what a standard deal would do, like capping your return or having these safety override provisions," they have said, "That's awesome."

So you feel like the business pressures or realities of the for-profit side of OpenAI will not conflict with the overall mission of the company?

Not at all. You could reference me with anyone. I'm sort of well known for not putting up with anything I don't want to put up with. I wouldn't do a deal if I thought that.

You guys are not monks in hair shirts saying, "We don't want to make a profit off of this." At the same time, it feels like you're not motivated by wealth creation, either.

I think it is a balance for sure. We want to make people very successful, making a great return [on their equity], that's great, as long as it's at a normal, reasonable level. If the full AGI thing breaks, we want something different for that paradigm. And we want the ability to bake in now how we're going to share this with society. I think we've done it in a nice way that balances it.

What has been the coolest thing you've seen someone do with GPT so far? And what's the thing that scares you most?

It's really hard to pick one coolest thing. It has been remarkable to see the diversity of things people have done. I could tell you the things that I have found the most personal utility in. Summarization has been absolutely huge for me, much more than I thought it would be. The fact that I can just have full articles or long email threads summarized has been way more useful than I would have thought. Also, the ability to ask esoteric programming

questions or help debug code in a way that feels like I've got a super brilliant programmer that I can talk to.

As far as a scary thing? I definitely have been watching with great concern the revenge porn generation that's been happening with the open source image generators. I think that's causing huge and predictable harm.

Do you think the companies who are behind these tools have a responsibility to ensure that kind of thing doesn't happen? Or is this just an unavoidable side of human nature?

I think it's both. There's this question of like, where do you want to regulate it? In some sense, it'd be great if we could just point to those companies and say, "Hey, you can't do these things." But I think people are going to open source models regardless, and it's mostly going to be great, but there will be some bad things that happen. Companies that are building on top of them, companies that have the last relationship with the end user, are going to have to have some responsibility, too. And so, I think it's going to be joint responsibility and accountability there.

(IICA duly acknowledge the authorship/ownership of the content and republishing the same with the sole objective of education.)

(Views/opinion mentioned in the document does not necessarily represent the IICA's stand.)