

URL: <https://stvp.stanford.edu/clips/responsible-ai-deployment>

Sam Altman, CEO of OpenAI, explains that he is increasingly worried about how to responsibly deploy more capable AI models. He predicts that more advanced models will require more granular deployment, as well as stakeholders working together to determine the rules of AI.



Transcript

Student I sort of wanted to talk to you about 00:00:04,620 responsible deployment of AGI.. So, as you guys continually inch closer to that, how do you plan to deploy that responsibly, at Open AI, you know, to prevent, stifling human innovation and continue to spur that? - So I'm actually not worried at all about 00:00:21,600 stifling of human innovation.. I really deeply believe that people will just surprise us on the upside with better tools.. I think all of history suggests that if you give people more leverage, they do more amazing things.. And that's kind of like we all get to benefit from that.. That's just kind of great.. I am though increasingly worried about how we're gonna do this all responsibly.. I think as the models get more capable, we have a higher and higher bar.. We do a lot of things like red teaming and external audits, and I think those are all really good.. But I think as the models get more capable, we'll have to deploy even more iteratively, have an even tighter feedback loop on looking at how they're used and where they work and where they don't work..

And this world that we used to do where we can release a major model update every couple of years, we probably have to find ways to like increase the granularity on that and deploy more iterative than we have in the past.. And it's not super obvious to us yet how to do that, but I think that'll be key to responsible deployment and also the way we kind of have all of the stakeholders negotiate what the rules of AI need to be.. That's gonna get more complex over time...