

Unconfuse Me with Bill Gates

EPISODE 06: Sam Altman

Date aired: January 11, 2024

BILL GATES: My guest today is Sam Altman. He, of course, is the CEO of OpenAI. He's been an entrepreneur and a leader in the tech industry for a long time, including running Y Combinator, that did amazing things like funding Reddit, Dropbox, Airbnb.

A little while after I recorded this episode, I was completely taken by surprise when, at least briefly, he was let go as the CEO of OpenAI. A lot happened in the days after the firing, including a show of support from nearly all of OpenAI's employees, and Sam is back. So, before you hear the conversation that we had, let's check in with Sam and see how he's doing.

[audio – Teams call initiation]

BILL GATES: Hey, Sam.

SAM ALTMAN: Hey, Bill.

BILL GATES: How are you?

SAM ALTMAN: Oh, man. It's been so crazy. I'm all right. It's a very exciting time.

BILL GATES: How's the team doing?

SAM ALTMAN: I think, you know a lot of people have remarked on the fact that the team has never felt more productive or more optimistic or better. So, I guess that's like a silver lining of all of this.

In some sense, this was like a real moment of growing up for us, we are very motivated to become better, and sort of to become a company ready for the challenges in front of us.

BILL GATES: Fantastic.

[audio – Teams call end]

[music]

So, we won't be discussing that situation in the conversation; however, you will hear about Sam's commitment to build a safe and responsible AI. I hope you enjoy the conversation.

Welcome to *Unconfuse Me*. I'm Bill Gates.

[music fades]

BILL GATES: Today we're going to focus mostly on AI, because it's such an exciting thing, and people are also concerned. Welcome, Sam.

SAM ALTMAN: Thank you so much for having me.

BILL GATES: I was privileged to see your work as it evolved, and I was very skeptical. I didn't expect ChatGPT to get so good. It blows my mind, and we don't really understand the encoding. We know the numbers, we can watch it multiply, but the idea of where is Shakespearean encoded? Do you think we'll gain an understanding of the representation?

SAM ALTMAN: A hundred percent. Trying to do this in a human brain is very hard. You could say it's a similar problem, which is there are these neurons, they're connected. The connections are moving and we're not going to slice up your brain and watch how it's evolving, but this we can perfectly x-ray. There has been some very good work on interpretability, and I think there will be more over time. I think we will be able to understand these networks, but our current understanding is low. The little bits we do understand have, as you'd expect, been very helpful in improving these things. We're all motivated to really understand them, scientific curiosity aside, but the scale of these is so vast. We also could say, where in your brain is Shakespeare encoded, and how is that represented?

BILL GATES: We don't know.

SAM ALTMAN: We don't really know, but it somehow feels even less satisfying to say we don't know yet in these masses of numbers that we're supposed to be able to perfectly x-ray and watch and do any tests we want to on.

BILL GATES: I'm pretty sure, within the next five years, we'll understand it. In terms of both training efficiency and accuracy, that understanding would let us do far better than we're able to do today.

SAM ALTMAN: A hundred percent. You see this in a lot of the history of technology where someone makes an empirical discovery. They have no idea what's going on, but it clearly works. Then, as the scientific understanding deepens, they can make it so much better.

BILL GATES: Yes, in physics, biology, it's sometimes just messing around, and it's like, whoa – how does this actually come together?

SAM ALTMAN: In our case, the guy that built GPT-1 sort of did it off by himself and solved this, and it was somewhat impressive, but no deep understanding of how it worked or why it worked. Then we got the scaling laws. We could predict how much better it was going to be. That was why, when we told you we could do a demo, we were pretty confident it was going to work. We hadn't trained the model, but we were pretty confident. That has led us to a bunch of attempts and better and better scientific understanding of what's going on. But it really came from a place of empirical results first.

BILL GATES: When you look at the next two years, what do you think some of the key milestones will be?

SAM ALTMAN: Multimodality will definitely be important.

BILL GATES: Which means speech in, speech out?

SAM ALTMAN: Speech in, speech out. Images. Eventually video. Clearly, people really want that. We've launched images and audio, and it had a much stronger response than we expected. We'll be able to push that much further, but maybe the most important areas of progress will be around reasoning ability. Right now, GPT-4 can reason in only extremely limited ways. Also reliability. If you ask GPT-4 most questions 10,000 times, one of those 10,000 is probably pretty good, but it doesn't always know which one, and you'd like to get the best response of 10,000 each time, and so that increase in reliability will be important.

Customizability and personalization will also be very important. People want very different things out of GPT-4: different styles, different sets of assumptions. We'll make all that possible, and then also the ability to have it use your own data. The ability to know about you, your email, your calendar, how you like appointments booked, connected to other outside data sources, all of that. Those will be some of the most important areas of improvement.

BILL GATES: In the basic algorithm right now, it's just feed forward, multiply, and so to generate every new word, it's essentially doing the same thing. I'll be interested if you ever get to the point where, like in solving a complex math equation, you might have to apply transformations an arbitrary number of times, that the control logic for the reasoning may have to be quite a bit more complex than just what we do today.

SAM ALTMAN: At a minimum, it seems like we need some sort of adaptive compute. Right now, we spend the same amount of compute on each token, a dumb one, or figuring out some complicated math.

BILL GATES: Yes, when we say, "Do the Riemann hypothesis ..."

SAM ALTMAN: That deserves a lot of compute.

BILL GATES: It's the same compute as saying, "The."

SAM ALTMAN: Right, so at a minimum, we've got to get that to work. We may need much more sophisticated things beyond it.

BILL GATES: You and I were both part of a Senate Education Session, and I was pleased that about 30 senators came to that, and helping them get up to speed, since it's such a big change agent. I don't think we could ever say we did too much to draw the politicians in. And yet, when they say, "Oh, we blew it on social media, we should do better," – that is an outstanding challenge that there are very negative elements to, in terms of polarization. Even now, I'm not sure how we would deal with that.

SAM ALTMAN: I don't understand why the government was not able to be more effective around social media, but it seems worth trying to understand as a case study for what they're going to go through now with AI.

BILL GATES: It's a good case study, and when you talk about the regulation, is it clear to you what sort of regulations would be constructed?

SAM ALTMAN: I think we're starting to figure that out. It would be very easy to put way too much regulation on this space. You can look at lots of examples of where that's happened before. But also, if we are right, and we may turn out not to be, but if we are right, and this technology goes as far as we think it's going to go, it will impact society, geopolitical balance of power, so many things, that for these, still hypothetical, but future extraordinarily powerful systems – not like GPT-4, but something with 100,000 or a million times the compute power of that, we have been socialized in the idea of a global regulatory body that looks at those super-powerful systems, because they do have such global impact. One model we talk about is something like the IAEA. For nuclear energy, we decided the same thing. This needs a global agency of some sort, because of the potential for global impact. I think that could make sense. There will be a lot of shorter term issues, issues of what are these models allowed to say and not say? How do we think about copyright? Different countries are going to think about those differently and that's fine.

BILL GATES: Some people think if there are models that are so powerful, we're scared of them – the reason nuclear regulation works globally, is basically everyone, at least on the civilian side, wants to share safety practices, and it has been fantastic. When you get over into the weapons side of nuclear, you don't have that same thing. If the key is to stop the entire world from doing something dangerous, you'd almost want global government, which today for many issues, like climate, terrorism, we see that it's hard for us to cooperate. People even invoke U.S.-China competition to say why any notion of slowing down would be inappropriate. Isn't any idea of slowing down, or going slow enough to be careful, hard to enforce?

SAM ALTMAN: Yes, I think if it comes across as asking for a slowdown, that will be really hard. If it instead says, "Do what you want, but any compute cluster above a certain extremely high-power threshold" – and given the cost here, we're talking maybe five in the world, something like that – any cluster like that has to submit to the equivalent of international weapons inspectors. The model there has to be made available for safety audit, pass some tests during training, and before deployment. That feels possible to me. I wasn't that sure before, but I did a big trip around the world this year, and talked to heads of state in many of the countries that would need to participate in this, and there was almost universal support for it. That's not going to save us from everything. There are still going to be things that are going to go wrong with much smaller-scale systems, in some cases, probably pretty badly wrong. But I think that can help us with the biggest tier of risks.

BILL GATES: I do think AI, in the best case, can help us with some hard problems.

SAM ALTMAN: For sure.

BILL GATES: Including polarization because potentially that breaks democracy and that would be a super-bad thing. Right now, we're looking at a lot of productivity improvement from AI, which is overwhelmingly a very good thing. Which areas are you most excited about?

SAM ALTMAN: First of all, I always think it's worth remembering that we're on this long, continuous curve. Right now, we have AI systems that can do tasks. They certainly can't do jobs, but they can do tasks, and there's productivity gain there. Eventually, they will be able to do more things that we think of like a job today, and we will, of course, find new jobs and better jobs. I totally

believe that if you give people way more powerful tools, it's not just that they can work a little faster, they can do qualitatively different things. Right now, maybe we can speed up a programmer 3x. That's about what we see, and that's one of the categories that we're most excited about it. It's working super-well. But if you make a programmer three times more effective, it's not just that they can do three times more stuff, it's that they can – at that higher level of abstraction, using more of their brainpower – they can now think of totally different things. It's like going from punch cards to higher level languages didn't just let us program a little faster, it let us do these qualitatively new things. We're really seeing that.

As we look at these next steps of things that can do a more complete task, you can imagine a little agent that you can say, "Go write this whole program for me, I'll ask you a few questions along the way, but it won't just be writing a few functions at a time." That'll enable a bunch of new stuff. And then again, it'll do even more complex stuff. Someday, maybe there's an AI where you can say, "Go start and run this company for me." And then someday, there's maybe an AI where you can say, "Go discover new physics." The stuff that we're seeing now is very exciting and wonderful, but I think it's worth always putting it in context of this technology that, at least for the next five or ten years, will be on a very steep improvement curve. These are the stupidest the models will ever be.

Coding is probably the single area from a productivity gain we're most excited about today. It's massively deployed and at scaled usage at this point. Healthcare and education are two things that are coming up that curve that we're very excited about too.

BILL GATES: The thing that is a little daunting is, unlike previous technology improvements, this one could improve very rapidly, and there's kind of no upper bound. The idea that it achieves human levels on a lot of areas of work, even if it's not doing unique science, it can do support calls and sales calls. I guess you and I do have some concern, along with this good thing, that it'll force us to adapt faster than we've had to ever before.

SAM ALTMAN: That's the scary part. It's not that we have to adapt. It's not that humanity is not super-adaptable. We've been through these massive technological shifts, and a massive percentage of the jobs that people do can change over a couple of generations, and over a couple of generations, we seem to absorb that just fine. We've seen that with the great technological revolutions of the past. Each technological revolution has gotten faster, and this will be the fastest by far. That's the part that I find potentially a little scary, is the speed with which society is going to have to adapt, and that the labor market will change.

BILL GATES: One aspect of AI is robotics, or blue-collar jobs, when you get hands and feet that are at human-level capability. The incredible ChatGPT breakthrough has kind of gotten us focused on the white-collar thing, which is super appropriate, but I do worry that people are losing the focus on the blue-collar piece. So how do you see robotics?

SAM ALTMAN: Super-excited for that. We started robots too early, so we had to put that project on hold. It was hard for the wrong reasons. It wasn't helping us make progress with the difficult parts of the ML research. We were dealing with bad simulators and breaking tendons and things like that. We also realized more and more over time that we first needed intelligence and cognition, and then we could figure out how to adapt it to physicality. It was easier to start with that with the way we built these language models. But we have always planned to come back to it.

We've started investing a little bit in robotics companies. On the physical hardware side, there's finally, for the first time that I've ever seen, really exciting new platforms being built there. At some point, we will be able to use our models, as you were saying, with their language understanding and future video understanding, to say, "All right, let's do amazing things with a robot."

BILL GATES: If the hardware guys who've done a good job on legs actually get the arms, hands, fingers piece, and then we couple it, and it's not ridiculously expensive, that could change the job market for a lot of the blue-collar type work, pretty rapidly.

SAM ALTMAN: Yes. Certainly, the prediction, the consensus prediction, if we rewind seven or ten years, was that the impact was going to be blue-collar work first, white-collar work second, creativity maybe never, but certainly last, because that was magic and human.

Obviously, it's gone exactly the other direction. I think there are a lot of interesting takeaways about why that happened. Creative work, the hallucinations of the GPT models is a feature, not a bug. It lets you discover some new things. Whereas if you're having a robot move heavy machinery around, you'd better be really precise with that. I think this is just a case of you've got to follow where technology goes. You have preconceptions, but sometimes the science doesn't want to go that way.

BILL GATES: So what application on your phone do you use the most?

SAM ALTMAN: Slack.

BILL GATES: Really?

SAM ALTMAN: Yes. I wish I could say ChatGPT.

BILL GATES: [*laughs*] Even more than e-mail?

SAM ALTMAN: Way more than e-mail. The only thing that I was thinking possibly was iMessages, but yes, more than that.

BILL GATES: Inside OpenAI, there's a lot of coordination going on.

SAM ALTMAN: Yes. What about you?

BILL GATES: It's Outlook. I'm this old-style e-mail guy, either that or the browser, because, of course, a lot of my news is coming through the browser.

SAM ALTMAN: I didn't quite count the browser as an app. It's possible I use it more, but I still would bet Slack. I'm on Slack all day.

BILL GATES: Incredible.

BILL GATES: Well, we've got a turntable here. I asked Sam, like I have for other guests, to bring one of his favorite records. So, what have we got?

SAM ALTMAN: I brought *The New Four Seasons - Vivaldi Recomposed* by Max Richter. I like music with no words for working. That had the old comfort of Vivaldi and pieces I knew really well, but enough new notes that it was a totally different experience. There are pieces of music that you form these strong emotional attachments to, because you listened to them a lot in a key period of your life. This was something that I listened to a lot while we were starting OpenAI.

I think it's very beautiful music. It's soaring and optimistic, and just perfect for me for working. I thought the new version is just super great.

BILL GATES: Is it performed by an orchestra?

SAM ALTMAN: It is. The Chineke! Orchestra.

BILL GATES: Fantastic.

SAM ALTMAN: Should I play it?

BILL GATES: Yes, let's.

[music – “The New Four Seasons – Vivaldi Recomposed: Spring 1” by Max Richter]

SAM ALTMAN: This is the intro to the sound we're going for.

[music]

BILL GATES: Do you wear headphones?

SAM ALTMAN: I do.

BILL GATES: Do your colleagues give you a hard time about listening to classical music?

SAM ALTMAN: I don't think they know what I listen to, because I do wear headphones. It's very hard for me to work in silence. I can do it, but it's not my natural state.

BILL GATES: It's fascinating. Songs with words, I agree, I would find that distracting, but this is more of a mood type thing.

SAM ALTMAN: Yes, and I have it quiet. I can't listen to loud music either, but it's just somehow always what I've done.

BILL GATES: It's fantastic. Thanks for bringing it.

[music fades]

BILL GATES: Now, with AI, to me, if you do get to the incredible capability, AGI, AGI+, there are three things I worry about. One is that a bad guy is in control of the system. If we have good guys who have equally powerful systems that hopefully minimizes that problem. There's the chance

of the system taking control. For some reasons, I'm less concerned about that. I'm glad other people are. The one that sort of befuddles me is human purpose. I get a lot of excitement that, hey, I'm good at working on malaria, and malaria eradication, and getting smart people and applying resources to that. When the machine says to me, "Bill, go play pickleball, I've got malaria eradication. You're just a slow thinker," then it is a philosophically confusing thing. How do you organize society? Yes, we're going to improve education, but education to do what, if you get to this extreme, which we still have a big uncertainty. For the first time, the chance that might come in the next 20 years is not zero.

SAM ALTMAN: There's a lot of psychologically difficult parts of working on the technology, but this is for me, the most difficult, because I also get a lot of satisfaction from that.

BILL GATES: You have real value added.

SAM ALTMAN: In some real sense, this might be the last hard thing I ever do.

BILL GATES: Our minds are so organized around scarcity; scarcity of teachers and doctors and good ideas that, partly, I do wonder if a generation that grows up without that scarcity will find the philosophical notion of how to organize society and what to do. Maybe they'll come up with a solution. I'm afraid my mind is so shaped around scarcity, I even have a hard time thinking of it.

SAM ALTMAN: That's what I tell myself too, and it's what I truly believe, that although we are giving something up here, in some sense, we are going to have things that are smarter than us. If we can get into this world of post-scarcity, we will find new things to do. They will feel very different. Maybe instead of solving malaria, you're deciding which galaxy you like, and what you're going to do with it. I'm confident we're never going to run out of problems, and we're never going to run out of different ways to find fulfilment and do things for each other and understand how we play our human games for other humans in this way that's going to remain really important. It is going to be different for sure, but I think the only way out is through. We have to go do this thing. It's going to happen. This is now an unstoppable technological course. The value is too great. And I'm pretty confident, very confident, we'll make it work, but it does feel like it's going to be so different.

BILL GATES: The way to apply this to certain current problems, like getting kids a tutor and helping to motivate them, or discover drugs for Alzheimer's, I think it's pretty clear how to do that. Whether AI can help us go to war less, be less polarized; you'd think as you drive intelligence, and not being polarized kind of is common sense, and not having war is common sense, but I do think a lot of people would be skeptical. I'd love to have people working on the hardest human problems, like whether we get along with each other. I think that would be extremely positive, if we thought the AI could contribute to humans getting along with each other.

SAM ALTMAN: I believe that it will surprise us on the upside there. The technology will surprise us with how much it can do. We've got to find out and see, but I'm very optimistic. I agree with you, what a contribution that would be.

BILL GATES: In terms of equity, technology is often expensive, like a PC or Internet connection, and it takes time to come down in cost. I guess the costs of running these AI systems, it looks pretty good that the cost per evaluation is going to come down a lot?

SAM ALTMAN: It's come down an enormous amount already. GPT-3, which is the model we've had out the longest and the most time to optimize, in the three and a little bit years that it has been out, we've been able to bring the cost down by a factor of 40. For three years' time, that's a pretty good start. For 3.5, we've brought it down, I would bet, close to 10 at this point. Four is newer, so we haven't had as much time to bring the cost down there, but we will continue to bring the cost down. I think we are on the steepest curve of cost reduction ever of any technology I know, way better than Moore's Law. It's not only that we figured out how to make the models more efficient, but also, as we understand the research better, we can get more knowledge, we can get more ability into a smaller model. I think we are going to drive the cost of intelligence down to so close to zero that it will be this before-and-after transformation for society.

Right now, my basic model of the world is cost of intelligence, cost of energy. *[Bill laughs]* Those are the two biggest inputs to quality of life, particularly for poor people, but overall. If you can drive both of those way down at the same time, the amount of stuff you can have, the amount of improvement you can deliver for people, it's quite enormous. We are on a curve, at least for intelligence, we will really, really deliver on that promise. Even at the current cost, which again, this is the highest it will ever be and much more than we want, for 20 bucks a month, you get a lot of GPT-4 access, and way more than 20 bucks' worth of value. We've come down pretty far.

BILL GATES: What about the competition? Is that kind of a fun thing that many people are working on this all at once?

SAM ALTMAN: It's both annoying and motivating and fun. *[Bill laughs]* I'm sure you've felt similarly. It does push us to be better and do things faster. We are very confident in our approach. We have a lot of people that I think are skating to where the puck was, and we're going to where the puck is going. It feels all right.

BILL GATES: I think people would be surprised at how small OpenAI is. How many employees do you have?

SAM ALTMAN: About 500, so we're a little bigger than before.

BILL GATES: But that's tiny. *[laughs]* By Google, Microsoft, Apple standards –

SAM ALTMAN: It's tiny. We have to not only run the research lab, but now we have to run a real business and two products.

BILL GATES: The scaling of all your capacities, including talking to everybody in the world, and listening to all those constituencies, that's got to be fascinating for you right now.

SAM ALTMAN: It's very fascinating.

BILL GATES: Is it mostly a young company?

SAM ALTMAN: It's an older company than average.

BILL GATES: Okay.

SAM ALTMAN: It's not a bunch of 24-year-old programmers.

BILL GATES: It's true, my perspective is warped, because I'm in my 60s. I see you, and you're younger, but you're right. You have a lot in their 40s.

SAM ALTMAN: Thirties, 40s, 50s.

BILL GATES: It's not the early Apple, Microsoft, which we were really kids.

SAM ALTMAN: It's not, and I've reflected on that. I think companies have gotten older in general, and I don't know quite what to make of that. I think it's somehow a bad sign for society, but I tracked this at YC. The best founders have trended older over time.

BILL GATES: That's fascinating.

SAM ALTMAN: Then in our case, it's a little bit older than the average, even still.

BILL GATES: You got to learn a lot by your role at Y Combinator, helping these companies. I guess that was good training for what you're doing now. *[laughs]*

SAM ALTMAN: That was super helpful.

BILL GATES: Including seeing mistakes.

SAM ALTMAN: Totally. OpenAI did a lot of things that are very against the standard YC advice. We took four and a half years to launch our first product. We started the company without any idea of what a product would be. We were not talking to users. I still don't recommend that for most companies, but having learned the rules and seen them at YC made me feel like I understood when and how and why we could break them. We really did things that were just so different than any other company I've seen.

BILL GATES: The key was the talent that you assembled, and letting them be focused on the big, big problem, not some near-term revenue thing.

SAM ALTMAN: I think Silicon Valley investors would not have supported us at the level we needed, because we had to spend so much capital on the research before getting to the product. We just said, "Eventually the model will be good enough that we know it's going to be valuable to people." But we were very grateful for the partnership with Microsoft, because this kind of way-ahead-of-revenue investing is not something that the venture capital industry is good at.

BILL GATES: No, and the capital costs were reasonably significant, almost at the edge of what venture would ever be comfortable with.

SAM ALTMAN: Maybe past.

BILL GATES: Maybe past. I give Satya incredible credit for thinking through 'how do you take this brilliant AI organization, and couple it into the large software company?' It has been very, very synergistic.

SAM ALTMAN: It's been wonderful, yes. You really touched on it, though, and this was something I learned from Y Combinator. We said, we are going to get the best people in the world at this. We are going to make sure that we're all aligned at where we're going and this AGI mission. But beyond that, we're going to let people do their thing. We're going to realize it's going to go through some twists and turns and take a while.

We had a theory that turned out to be roughly right, but a lot of the tactics along the way turned out to be super wrong. We just tried to follow the science.

BILL GATES: I remember going and seeing the demonstration and thinking, okay, what's the path to revenue on that one? What is that like? In these frenzied times, you're still holding on to an incredible team.

SAM ALTMAN: Yes. Great people really want to work with great colleagues.

BILL GATES: That's an attractive force.

SAM ALTMAN: There's a deep center of gravity there. Also, it sounds so cliché, and every company says it, but people feel the mission so deeply. Everyone wants to be in the room for the creation of AGI.

BILL GATES: It must be exciting. I can see the energy when you come up and blow me away again with the demos; I'm seeing new people, new ideas. You're continuing to move at a really incredible speed.

SAM ALTMAN: What's the piece of advice you give most often?

BILL GATES: There are so many different forms of talent. Early in my career, I thought, just pure IQ, like engineering IQ, and of course, you can apply that to financial and sales. That turned out to be so wrong. Building teams where you have the right mix of skills is so important. Getting people to think, for their problem, how do they build that team that has all the different skills, that's probably the one that I think is the most helpful. Yes, telling kids, you know, math, science is cool, if you like it, but it's that talent mix that really surprised me.

What about you? What advice do you give?

SAM ALTMAN: It's something about how most people are mis-calibrated on risk. They're afraid to leave the soft, cushy job behind to go do the thing they really want to do, when, in fact, if they don't do that, they look back at their lives like, "Man, I never went to go start this company I wanted to start, or I never tried to go be an AI researcher." I think that's sort of much riskier.

Related to that, being clear about what you want to do, and asking people for what you want goes a surprisingly long way. A lot of people get trapped in spending their time in not the way they want to do. Probably the most frequent advice I give is to try to fix that some way or other.

BILL GATES: If you can get people into a job where they feel they have a purpose, it's more fun. Sometimes that's how they can have gigantic impact.

SAM ALTMAN: That's for sure.

BILL GATES: Thanks for coming. It was a fantastic conversation. In the years ahead, I'm sure we'll get to talk a lot more, as we try to shape AI in the best way possible.

SAM ALTMAN: Thanks a lot for having me. I really enjoyed it.

[music]

BILL GATES: *Unconfuse Me* is a production of the Gates Notes. Special thanks to my guest today, Sam Altman.

BILL GATES: Remind me what your first computer was?

SAM ALTMAN: A Mac LC2.

BILL GATES: Nice choice.

SAM ALTMAN: It was a good one. I still have it; it still works.