

Toward Humanist Superintelligence

Nov 7, 2025 | MUSTAFA SULEYMAN

REDMOND – For decades, the Turing test was AI researchers’ North Star. Today, it’s been quietly surpassed. With reasoning models and agentic capabilities emerging, and with the pace of AI infrastructure build increasing, we have crossed an inflection point on the journey to superintelligence: the point at which AI exceeds human-level performance at all tasks. Indeed, the most consequential question for our time is not whether AI will surpass us, because in some ways it already has (try beating an AI at general knowledge), in many other ways, it will, and in some ways we will always be unique. The real question, then, is whether we can shape AI to advance human flourishing rather than undermine it. That is the most important challenge of our time.

To be sure, everyone is primed by now to roll their eyes at AI hype. I get it. But the stakes could not be higher. Science and technology have always been humanity’s greatest engine of progress. Over the last 250 years, that engine has doubled life expectancy, lifted billions of people out of poverty, and given us antibiotics, electricity, and instant global communication. AI is the next chapter in this story. It represents our best shot at accelerating scientific discovery, economic growth, and human well-being. Whenever you hear about AI, this potential is worth keeping in mind.

But harnessing AI’s potential will work out right only if we *build* AI the right way. The costs of getting it wrong are immense. No one yet has reassuring answers about how we contain or align these systems. We are caught at an odd moment, faced with history’s most powerful technologies and unsure how they can be controlled or whether they will remain beneficial.

Make your inbox smarter.
Select Newsletters

I think we can cut through the noise and understand it like this: AI, like all technology, can be judged by a simple test. Does it improve human life? Is it clearly working in service of people?

As we embark on the next phase of AI, the answer to these questions lies in what I call Humanist Superintelligence (HSI): advanced AI designed to remain controllable, aligned, and firmly in service to humanity. This project is explicitly about avoiding, at all costs, an unbounded entity with total autonomy.

Instead, we must focus on domain-specific superintelligence. Rather than simply making a system that can endlessly improve and run away with itself for whatever purpose it might eventually arrive at, the core purpose is to deliver practical real-world benefits to billions of people. It must forever remain *unequivocally subordinate* to humanity. This is the vision of our Superintelligence Team at Microsoft, where our core mission is to keep humanity secure and firmly in control.

Why humanism? Because history has demonstrated the humanist tradition’s enduring power to preserve human dignity. AI built in that spirit can unlock extraordinary benefits while avoiding

catastrophic risks. We need a vision of AI that supports humanity, amplifies creativity, and protects our fragile environment – not one that sidelines us.

The prize for humanity is enormous: a world of rapid advances in living standards and science, and a time of new art forms, culture, and growth. It is a truly inspiring mission that has motivated me for decades. We should celebrate and accelerate technology as the greatest engine of progress that humanity has ever known. That's why we need much, much more of it.

HSI offers a safer path forward. Remaining grounded in domain-specific breakthroughs with profound societal impact is an example of this. Imagine AI companions that ease the mental load of daily life, enhance productivity, and transform education through adaptive, individualized learning. Imagine medical superintelligence delivering accurate, affordable expert-level diagnostics that could revolutionize global health care, capabilities already previewed by our health team at Microsoft AI. And consider the potential for AI-driven advances in clean energy that will enable abundant, low-cost power generation, storage, and carbon removal to meet soaring demand while protecting the planet.

With HSI, these are not speculative dreams. They are achievable goals that can benefit people around the world, providing concrete improvements to everyday life.

To state the obvious, humans matter more than tech or AI. Superintelligence could be the best invention ever, but only if it sticks to this maxim. That means ensuring accountability and transparency, and a willingness to make safety a top priority. Our goal is not to build a superintelligence at any cost, but to follow a careful path toward one that is contained, value-aligned, and always focused on human well-being.

Everybody needs to ask themselves this: What kind of AI do we actually want? The answer will shape the future of civilization. For me, that answer is Humanist Superintelligence.

MUSTAFA SULEYMAN

Mustafa Suleyman is CEO of Microsoft AI and the author of *The Coming Wave: Technology, Power, and the Twenty-First Century's Greatest Dilemma* (Crown, 2023). He previously Co-Founded Inflection AI and DeepMind.

<https://prosyn.org/eE7EpJ>

Support Project Syndicate

Subscribe Upgrade Donate

Make your inbox smarter.
Select Newsletters

[Content](#)

[PS Quarterly](#)

[Sections](#)

[Topics](#)

[Archive](#)

[Columnists](#)

[Videos](#)

[Podcasts](#)

[Subscriber Exclusive](#)

[OnPoint](#)

[Longer Reads](#)

[Insider Interviews](#)

[Big Picture](#)

[Global Bookmark](#)

[Rewards](#)

[Project Syndicate](#)

[About Us](#)

[Contact](#)

[Careers](#)

[FAQ](#)

[Submission policy](#)

[Newsletters](#)

[Member Publications](#)

[Network](#)

[Syndication](#)

[PS Magazines](#)

[PS Events](#)

[Institutions](#)

[Corporate](#)

© Project Syndicate - 2025 [Privacy Policy](#) [Terms & Conditions](#) [Cookie Policy](#)