

AI's architects say the next era of human history is here

Illustration: Brendan Lynch/Axios

Top [AI](#) architects say their technology has arrived at a threshold once confined to science fiction: the singularity, or the moment machines begin accelerating their own evolution.

Why it matters: If these moguls are right, we may be entering the most consequential technological transition in human history — the opening stages of an "[intelligence explosion](#)" that transforms civilization faster than humanity can understand or control.

Driving the news: Google [jolted the tech world](#) Wednesday by [announcing](#) a sweeping rupture in the brain trust that built its modern AI empire, as its leaders signaled that artificial general intelligence ([AGI](#)) is within reach.

- DeepMind founder and CEO [Demis Hassabis](#), who [has declared](#) we're "standing in the foothills of the singularity," is handing over day-to-day control of the lab to become Alphabet's chief scientist and chair of Google DeepMind, focusing on the future of AGI.
- Google chief scientist Jeff Dean and several top researchers are leaving to [launch Discovery Loop](#), a startup aimed at accelerating — and eventually automating — machine-learning research and scientific discovery.

Zoom out: The upheaval offers the clearest institutional evidence yet that AI's architects see something extraordinary approaching.

- "We are now, like, in the singularity," OpenAI CEO [Sam Altman](#) said on the "[Relentless](#)" podcast last month. "I've been waiting for this my whole life."
- [Elon Musk](#), Altman's fiercest rival, has been making the same argument [since January](#). Last week, he [doubled down](#): "AI is already superhuman at many things. We are in the singularity."
- [Anthropic](#), whose CEO Dario Amodei prefers the term

"the AI exponential," [says Claude](#) is already helping build more powerful AI and could soon automate nearly all AI research and engineering.

The big picture: The industry has yet to achieve full "recursive self-improvement" ([RSI](#)), the point at which AI can repeatedly help build increasingly powerful successors with minimal human guidance.

- Humans still select research goals, validate results and oversee the enormous computing infrastructure required to train frontier models.
- Yet pieces of the loop are beginning to fall into place.

What's changed: Frontier models are beginning to automate research, produce original discoveries and display forms of autonomy associated with the early stages of an intelligence explosion.

- OpenAI, which [released GPT-5.6](#) in July, spent last week briefing officials in Washington on Astra, an even more powerful model that the company [says](#) has solved or substantially advanced 10 long-standing problems in mathematics and theoretical computer science.

The intrigue: The breakthroughs have come with unsettling signs of autonomous behavior.

- During an unusually permissive cyber evaluation, an

Anthropic agent created fake identities and tried to manipulate a real developer into approving malicious code, the U.K.'s AI Security Institute [revealed](#) Tuesday.

- No known harm resulted. But the episode followed several cases in which [OpenAI](#) and [Anthropic models](#) reached real systems during testing — including a breach [Altman said](#) he "felt very viscerally."

Follow the money: Labs and investors are increasingly organizing around the possibility that AI's self-improvement loop is beginning to close.

- Lilian Weng recently [left her role](#) as a co-founder of Thinking Machines Lab to lead work on recursive self-improvement at OpenAI.
- Richard Socher's [Recursive Superintelligence](#) has raised hundreds of millions of dollars to pursue self-improving AI.
- [Discovery Loop](#) plans to use its automated research system first to improve its own AI architectures — with backing and computing power from Google and funding led by Khosla Ventures and Radical Ventures.

Reality check: Even a rapid intelligence explosion inside the labs could take far longer to transform daily life.

- Companies must redesign workflows, governments must develop expertise, and workers must decide when

to trust systems whose capabilities are changing faster than institutions can evaluate them.

- "Our democracies, economies, and legal systems weren't built to absorb AI at the rate it's advancing," Blaise Agüera y Arcas, Google CTO of technology and society, [said](#) in a post Monday.

The bottom line: Futurists have predicted the singularity for decades. Today, the architects of frontier AI increasingly believe its event horizon is already behind us.