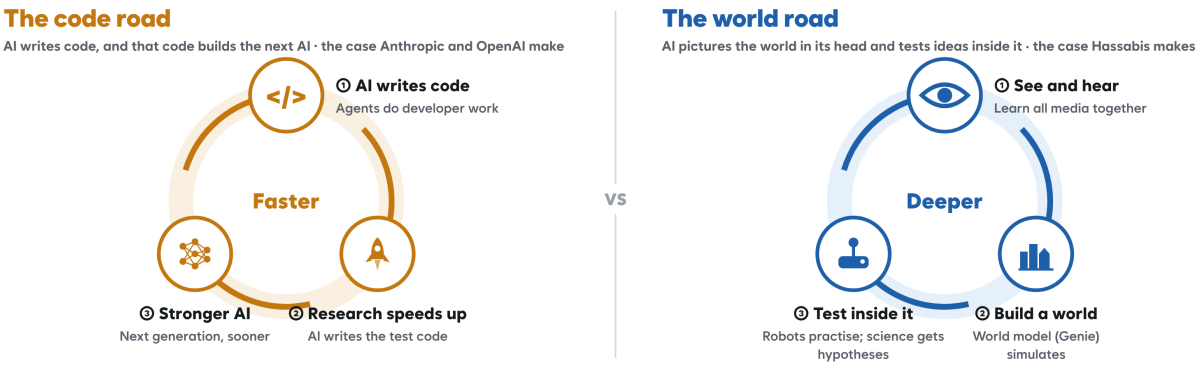
DEEP DIVE LIGHTVIEW DEEP DIVE · PEOPLE

Code or the World: Hassabis and Google's AGI Bet

From the ChatGPT shock to Gemini 4, and the road that split from OpenAI and Anthropic

Part 2 Now and next · 2017–2026

Part 1 · From the chessboard to the Nobel



Key points

1. **In late 2022, Google lost the lead to ChatGPT, built on Google's own invention.** While Google hesitated, OpenAI moved first, and the rushed Bard stumbled on a wrong answer.
2. **Google handed the fight to Hassabis, and by 2025 the tables had turned.** The merged Google DeepMind built Gemini to learn text, images and audio together; its app now has over 1 billion users.
3. **Then, in August 2026, the man who led the comeback stepped down as CEO.** Amid reported delays and departures, he handed off daily management for the big picture and science; Gemini 4 and Genie 4 followed.
4. **The road he is pushing differs from OpenAI's and Anthropic's.** They lean on AI writing code to build AI faster; Hassabis bets on AI that pictures the world in its head and tests ideas there (world models).
5. **Which road is right may become clear within a few years.** All three leaders say AGI is near, and even the cautious Hassabis has moved his clock closer. Investors can watch Google Cloud revenue and ~$200bn of capex.

1 Built the technology, took the first hit

Part 1 ended in autumn 2024, with Hassabis receiving the Nobel while Google struggled on a very different battlefield. Ironically, the weapon in that battle was one Google had made.

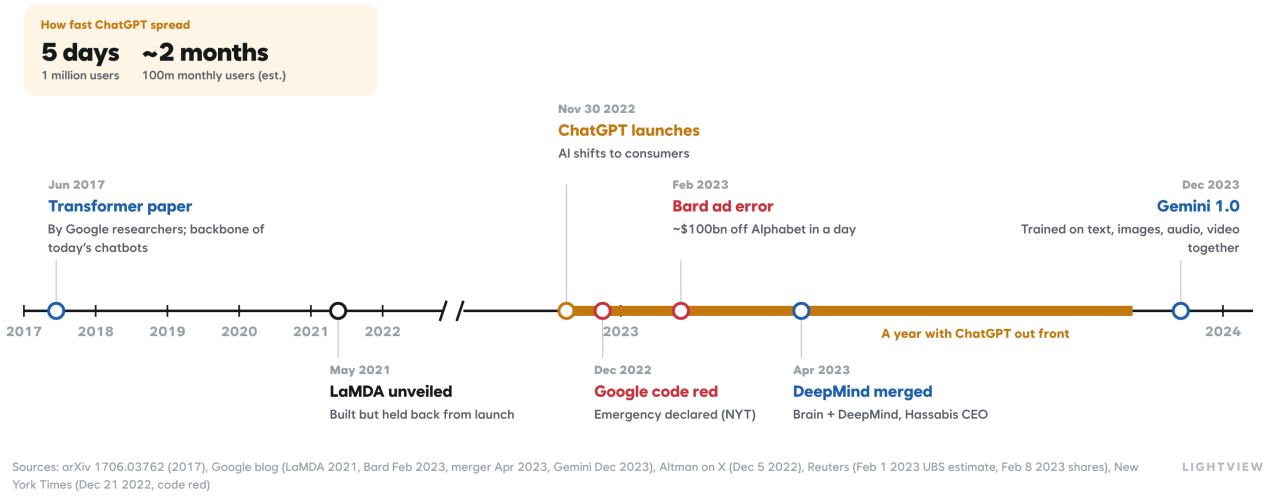


Figure 1 · The ChatGPT shock and Google's response From the transformer paper to Gemini 1.0. The middle of the axis is compressed

In June 2017, Google researchers published "Attention Is All You Need", introducing the transformer that underpins today's chatbots. Google unveiled its conversational model LaMDA in 2021 and was ahead technically, but it was cautious about releasing a chatbot to the public. Commonly cited reasons: wrong answers could hurt trust in a search company, and a chatbot could eat into its lucrative search ads.

Meanwhile, on November 30 2022, OpenAI released ChatGPT. It passed 1 million users in five days and was estimated to reach 100 million monthly users in two months. That December, Google's leaders reportedly declared a "code red".

The rushed response backfired. In February 2023 a Bard chatbot promo got a fact about the James Webb Space Telescope wrong, and Alphabet shares fell about 7.7% in a day, erasing roughly $100bn.

Two months later Google reorganized, merging Google Brain and DeepMind into "Google DeepMind" under Hassabis. Sundar Pichai said bringing the talent into one team, backed by Google's computing resources, would speed up progress.

First in technology, late to launch

When	Google	OpenAI
Jun 2017	Transformer paper	
May 2021	LaMDA unveiled (no public launch)	
Nov 2022		ChatGPT launched
Feb–Dec 2023	Bard → merger → Gemini 1.0	GPT-4 (Mar 2023)

2 Gemini, built to learn it all together, and the comeback

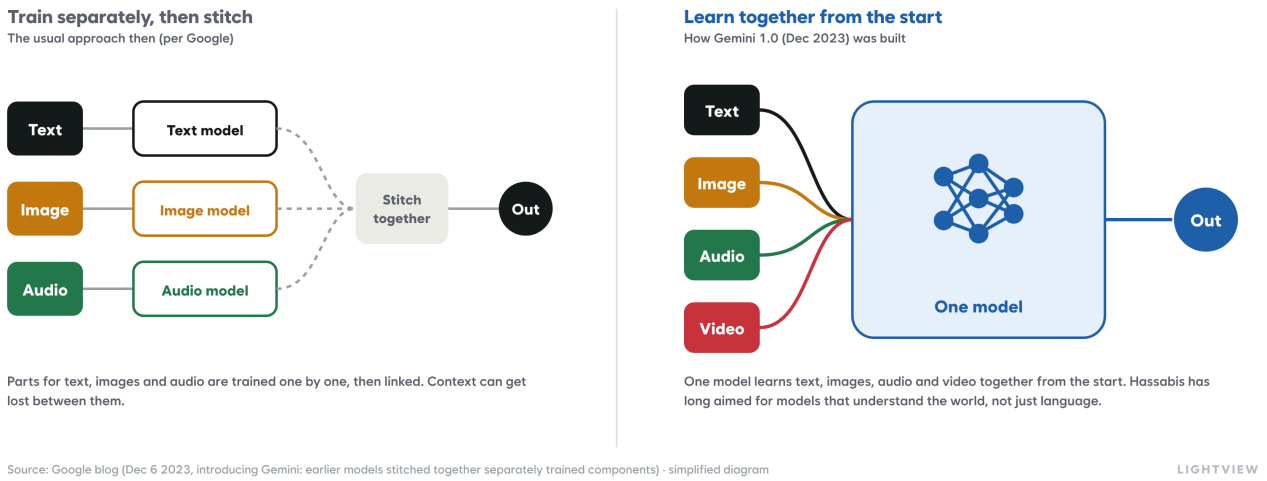


Figure 2 · How Gemini was built Instead of stitching separate parts together, one model learns from the start

Gemini 1.0, released in December 2023, was designed to learn text, images, audio and video in one model from the start, reflecting Hassabis's long-stated aim of AI that understands the world beyond language. The start was rocky: its demo video was criticized as smoother than reality, and in February 2024 Google paused generating images of people after historical figures came out wrong.

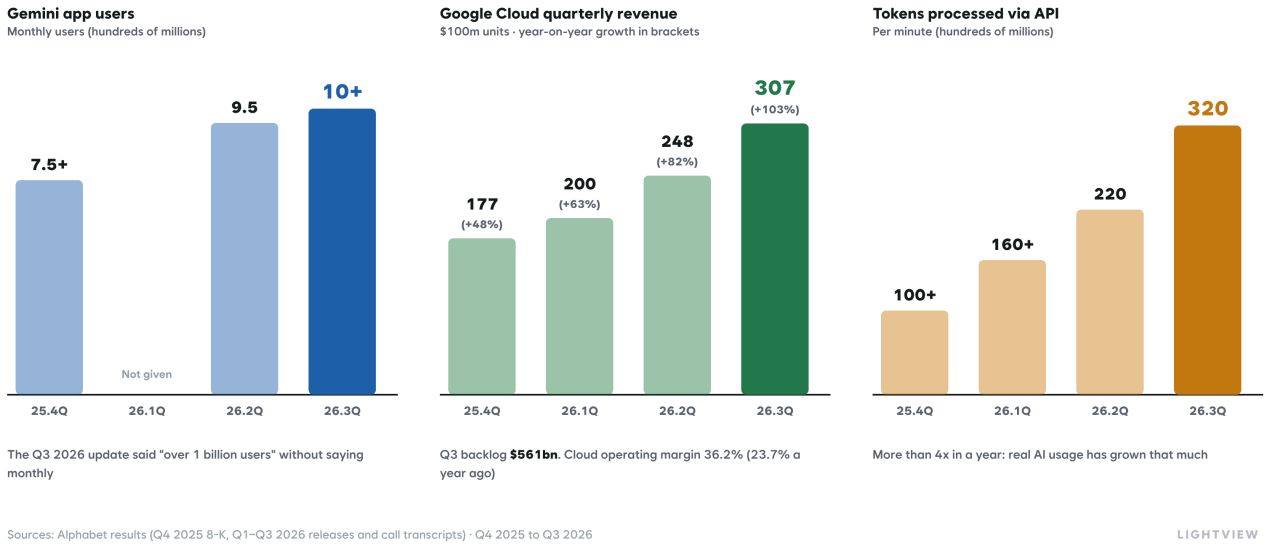


Figure 3 · The numbers came back Gemini app users, Google Cloud revenue, and tokens processed through the API

2025 was the comeback year. A string of model gains and hit features even pushed the Gemini app past ChatGPT in the App Store rankings. That December, it was OpenAI's Sam Altman who reportedly declared a "code red". In three years, the roles had flipped.

The 2025 comeback

When	What happened
Mar 2025	Gemini 2.5 Pro debuts at No. 1 on the LMArena user-vote leaderboard
Jul 2025	Gold-medal standard at the International Mathematical Olympiad, officially graded
Aug–Sep 2025	"Nano Banana" image editing takes off → Gemini app No. 1 free app in the US App Store

3 Stepping back from CEO

On August 5 2026, Hassabis stepped down as CEO of Google DeepMind. He became chair and Alphabet's Chief Scientist and continues to lead the drug-discovery company Isomorphic Labs. Day-to-day operations went to Koray Kavukcuoglu, previously CTO, as senior vice president.

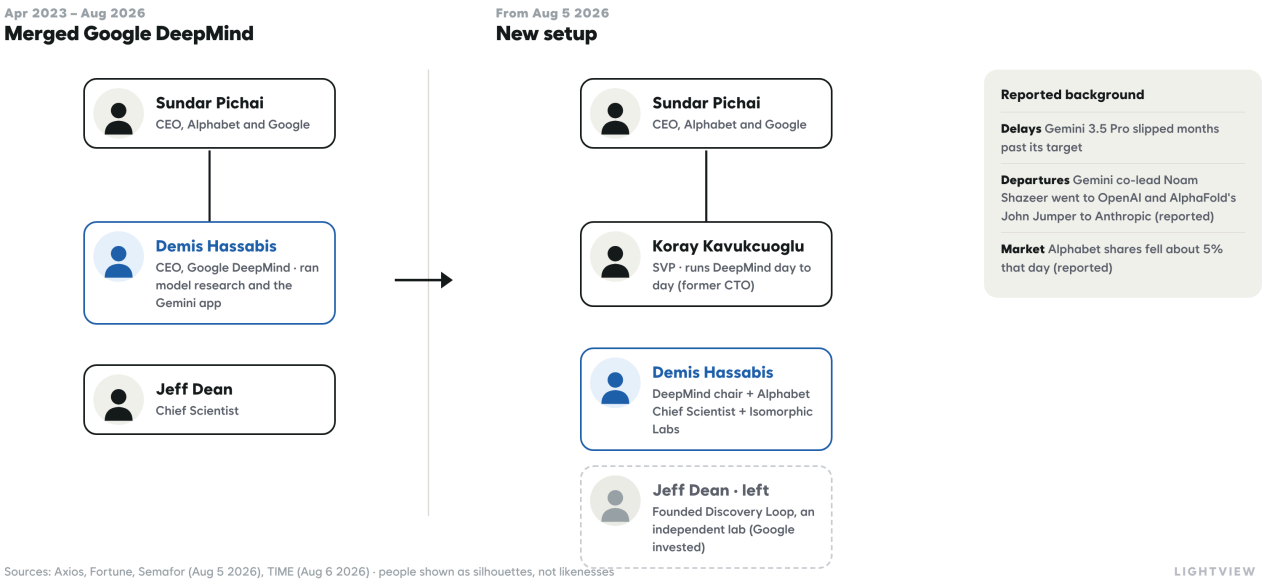


Figure 4 · The new structure Hassabis hands off daily management to become chair and Chief Scientist. Jeff Dean left the same day

The reported backdrop was rough. Gemini 3.5 Pro, aimed at a summer launch, was months late, and reports said Gemini co-lead Noam Shazeer had moved to OpenAI and John Jumper, his AlphaFold Nobel co-laureate, to Anthropic. Alphabet shares were reported to have fallen about 5% that day.

Hassabis tells it differently. In an internal memo he said he was handing over daily duties to have time and space for the big picture. Whether this was accountability for a rough stretch or a researcher returning to research is open to interpretation.

> **"I've been working towards AGI my whole life and now, like many of you, I feel it is close at hand"**
>
> Hassabis internal memo (Aug 5 2026, via Axios and TIME)

What is clear is the pace since. Within two months came Gemini 4 Argon and Carbon and the Genie 4 world model. How a Google DeepMind no longer run day to day by Hassabis turns his ideas into products is the thing to watch now.

Hassabis's three roles now

Role	What it involves
Chair, Google DeepMind	Leads research direction and the big picture (Kavukcuoglu runs daily operations)
Chief Scientist, Alphabet	Advises on science and technology strategy across Google
Head of Isomorphic Labs	Runs the subsidiary using AlphaFold technology to find drug candidates (founded 2021; Eli Lilly and Novartis research deals in 2024; first outside funding of $600m in 2025)

4 Gemini 4, back in the front row

Skipping Gemini 3.5 Pro, Google jumped to a fourth generation. Gemini 4 Argon went first to cyber defenders on September 30 (the US government also took part in a pre-release review) and opened to developers on October 13. It was reported to lead outright on 12 of 18 published benchmarks, pulling well ahead on long-video understanding. Google also said it scored 87% on ERQA, a test of how well AI understands space and chooses actions, as a robot must.

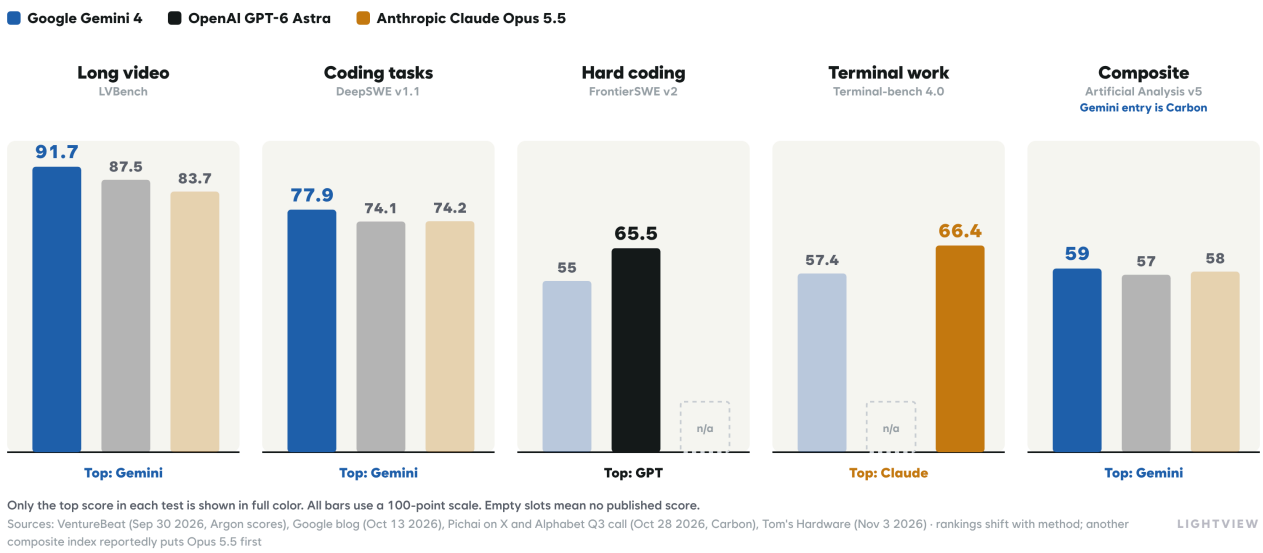


Figure 5 · Scores by test Gemini tops long video, coding tasks and the composite; GPT tops hard coding; Claude tops terminal work

What the tests in Figure 5 measure

Area	Test	What it checks
Long video	LVBench	Answering questions about videos around an hour long
Coding, hard coding	DeepSWE, FrontierSWE	How well it solves realistic software tasks
Terminal work	Terminal-bench	Finishing multi-step tasks in a command line on its own
Composite	Artificial Analysis	Many test scores combined into one

The weak spots were clear too. On hard coding tasks and terminal work, it trailed the OpenAI and Anthropic models. Carbon, the higher-end model Sundar Pichai unveiled at the October 28 earnings call, scored 59 on one composite index, 1–2 points ahead of Claude Opus 5.5 (58) and GPT-6 Astra (57), though another index reportedly ranks Opus 5.5 first. Leaderboard leads change within weeks, so it is more useful to look at where each company is strong than at a single number.

Gemini 4's two models

	Argon	Carbon
Strengths	Long video, spatial and action reasoning, 1 million tokens at once	No. 1 on one composite index; claimed lowest hallucination rate
Price (per million tokens)	$2 in / $10 out; $4 / $20 after Nov 30	Not yet open to developers

Sources: Google blog (Oct 13 2026), Tom's Hardware (Nov 3 2026) · Kavukcuoglu said in September that "we are always gonna be at the frontier" (9to5Google)

5 Three companies, three centers of gravity

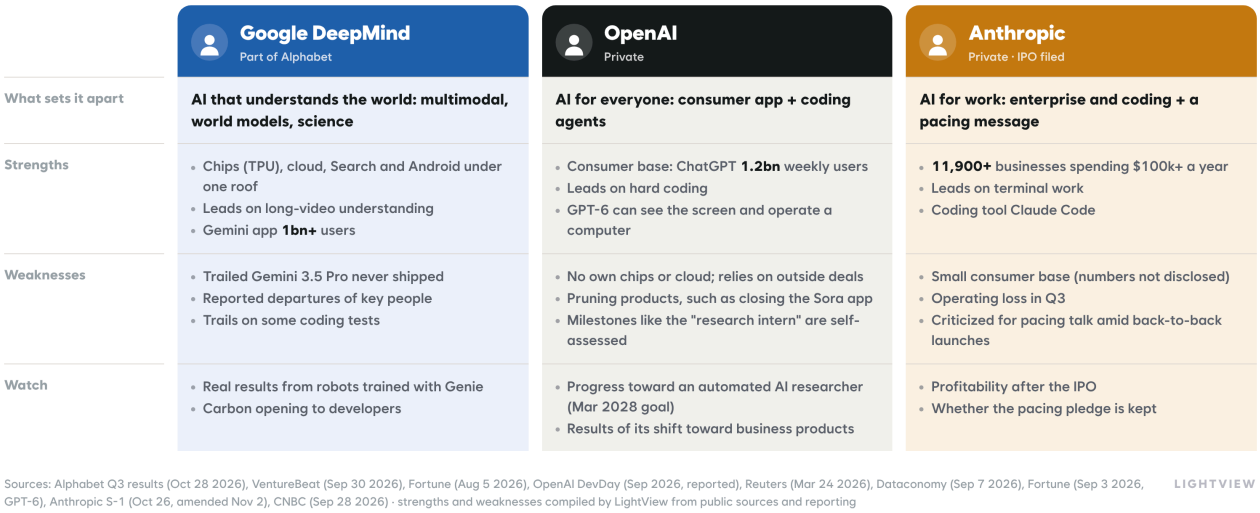


Figure 6 · What sets each company apart, strengths and weaknesses As of early November 2026, compiled from public sources and reporting

OpenAI started as a consumer company, with ChatGPT used by 1.2 billion people a week, but it has recently leaned into coding agents and business products. In March 2026 it decided to close the Sora video app; Reuters said the aim was to focus on enterprise products. In September it announced that it had met its own goal of an "AI research intern", a self-assessment without outside verification.

Anthropic made businesses its main customers from the start. Business customers grew fast, but its amended IPO filing in November shows short-term compute costs and a one-off non-cash charge pushed it to an operating loss in the third quarter. In a September essay Dario Amodei said AI progress should be slowed and that Anthropic would move first, yet it kept releasing new models, drawing criticism that words and actions did not match.

Google's strength is having chips (TPUs), cloud and distribution through Search and Android under one roof. That lets it set its own pace between research and products. The relationships are not simple, though: Anthropic rents Google TPUs at scale and is also a company Google has invested in. Rival and customer at once.

In the end, the three differ on where they look for their next growth. OpenAI layers work tools on its user base, Anthropic sells AI that businesses can trust, and Google bets that AI which understands the world will go further. Their choices over the past year show the same direction.

Choices made in the past year

Company	Choice	Direction it suggests
Google DeepMind	Hassabis moves from CEO to chair and Chief Scientist (Aug 2026); Genie 4 world model (Oct)	Product race to operators; the founder to science and world models
OpenAI	Decides to close the Sora video app (Mar 2026); claims an "AI research intern" (Sep)	From consumer buzz toward business products and automating AI research
Anthropic	"Pace the frontier" essay (Sep 2026); IPO filing (Oct)	Safety-first AI for business, while still shipping

"Direction it suggests" is LightView's reading of company announcements and reporting

6 The code road and the world road

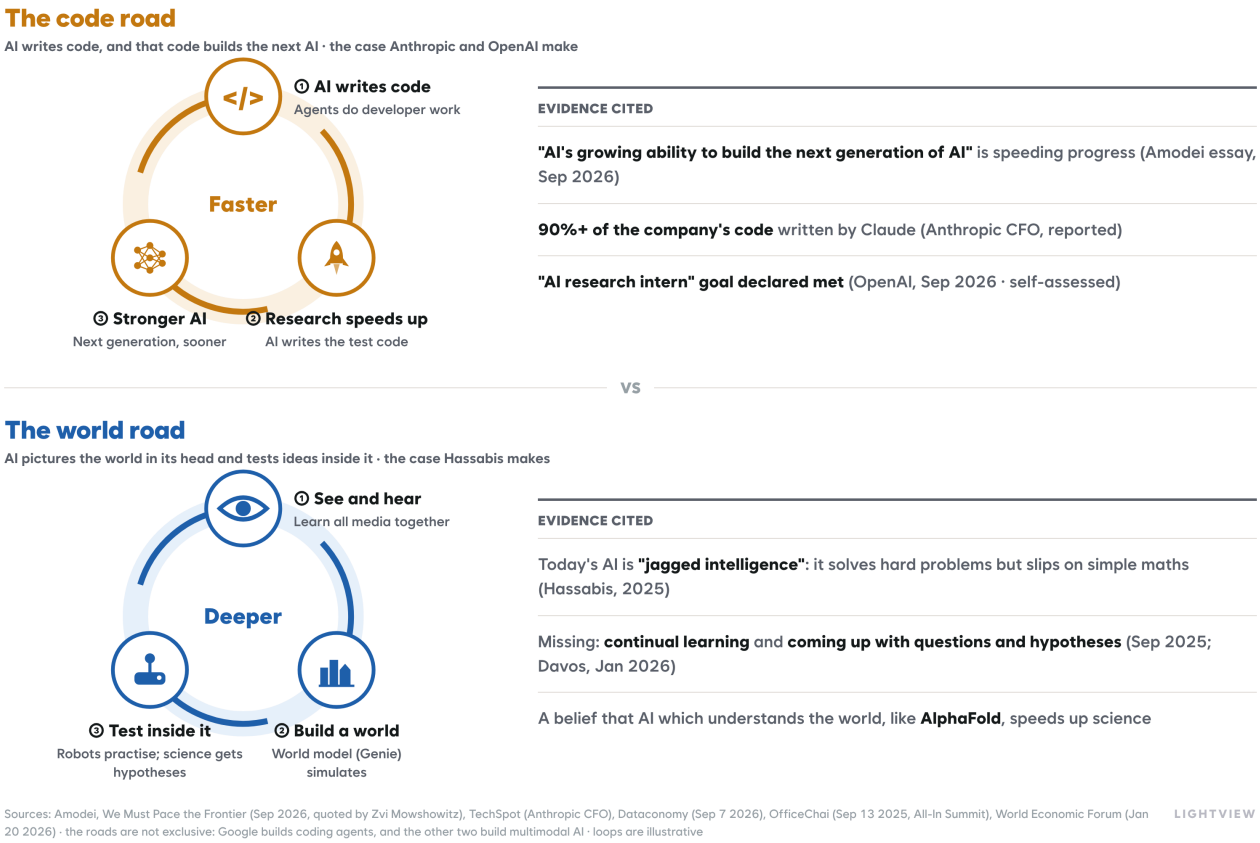


Figure 7 · Two roads A loop where AI builds AI faster, and a loop where AI understands the world more deeply

The logic of the code road is simple. If AI writes code well, AI research itself speeds up, and the stronger AI it produces speeds research up again. Amodei wrote in September that AI progress is accelerating sharply, "driven primarily by AI's growing ability to build the next generation of AI."

Hassabis looks elsewhere. He calls today's AI "jagged intelligence", solving hard problems but slipping on simple arithmetic, and says it still lacks continual learning and the ability to come up with its own questions and hypotheses. As in his hippocampus research (Part 1), real intelligence must build a world in its head and test the future there.

> **"[They] can make simple mistakes with even high school maths and simple counting. That shouldn't be possible for a true AGI system"**
>
> Hassabis, All-In Summit (Sep 2025, reported by OfficeChai), excerpt

The two roads side by side

	The code road	The world road
Strength	Quick revenue; easy to measure	Aims at big markets not yet open, like robots and science
Weakness	The physical world is a separate problem	Slow to make money; progress is hard to measure
Sign of success	More AI research results produced by AI	Sim-trained robots work in reality

LightView's comparison based on company announcements and leaders' statements

7 Genie 4, AI that builds a world in its head

The signature work of the world road is the Genie world model. Give it an image or a line of text, and it generates, in real time, a world you can walk around and control.

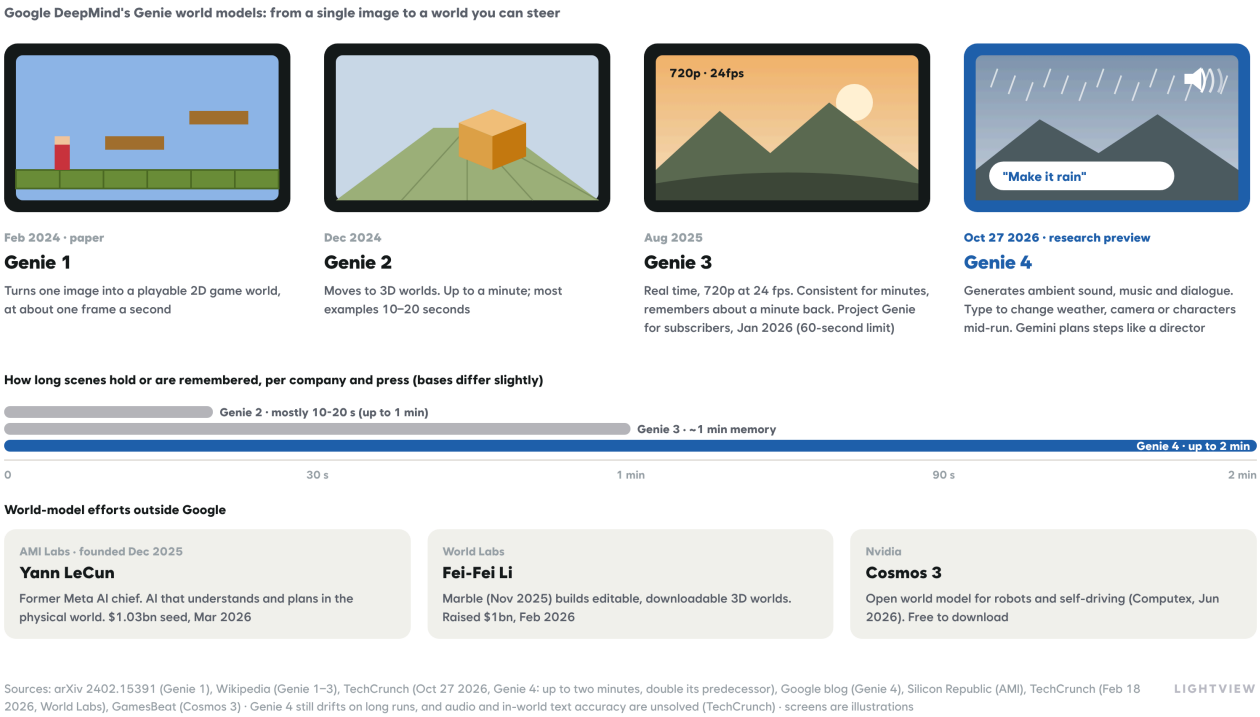


Figure 8 · How Genie evolved, and other world-model efforts From a 2D game screen to a world with sound. Scene consistency now reaches two minutes

Genie 4, unveiled on October 27, generates ambient sound, music and even dialogue, and lets you type to change the weather or camera angle or add characters while the world runs. It can keep a scene consistent for up to two minutes, double its predecessor, according to TechCrunch.

Given Hassabis's career in games, the direction feels natural, but the goal is not play. Google described a research experiment in which Gemini 4 Argon broke a robot's task into steps and Genie 4 simulated the scene in real time: practise in a virtual world before moving a real robot. Limits remain, such as details drifting on long runs. Outside Google, Yann LeCun's AMI Labs, Fei-Fei Li's World Labs and Nvidia are betting big on the same road.

Where world models could be used

Field	Use	Stage now
Robots	Practise hard or risky moves in a virtual world first	Research experiment (Genie 4 + Gemini Robotics)
Games and film	Worlds you can try from a few lines of text; scene ideation	Subscriber trial (Project Genie); testers using it for storyboards
Self-driving	Rare crash scenes to train driving AI	Open models such as Nvidia Cosmos
AI agents	AI that sets goals, learns in 3D	DeepMind SIMA 2 research (Nov 2025)

Sources: TechCrunch (Oct 27 2026), Google DeepMind (Nov 2025, SIMA 2), GamesBeat (Cosmos 3)

8 The AGI clock, and what investors should watch

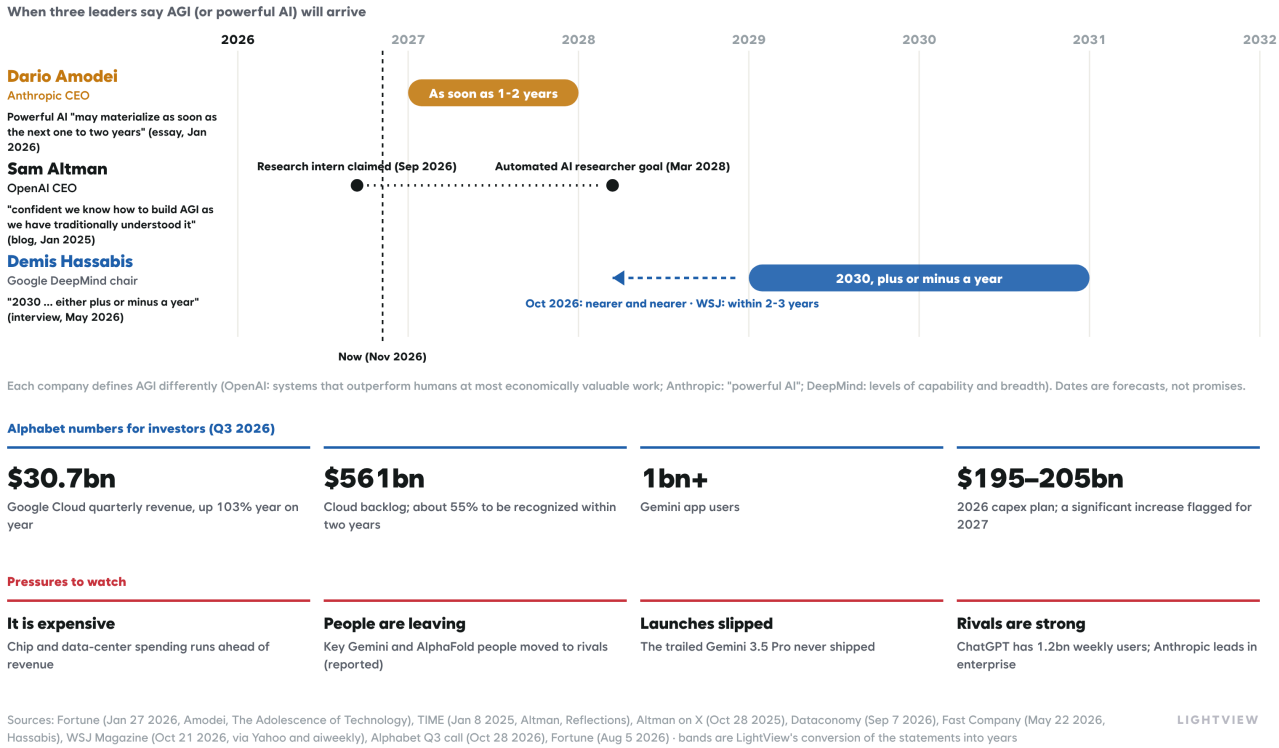


Figure 9 · Three leaders' AGI timelines and Alphabet's numbers Forecasts range from one or two years out to around 2030

The three leaders' clocks differ. In January 2026 Amodei wrote that powerful AI could arrive as soon as the next one to two years, and in January 2025 Altman wrote that OpenAI was confident it knew how to build AGI "as we have traditionally understood it." Hassabis, the most cautious, has long said 2030 plus or minus a year, but in an October interview he said it looks nearer and nearer, and the Wall Street Journal read his outlook as within two to three years. Each defines AGI differently, so what they cite as evidence matters more than the dates.

For investors, Google's bet shows up in numbers. Cloud revenue and backlog show AI demand turning into real money, but this year's capex of around $200bn is a heavy load. Whether Hassabis's answer to the question that began at the chessboard is right depends, in the end, on whether AI that understands the world creates more value than AI that writes code.

What to check next

Watch	Good sign	Bad sign
Cloud	Revenue growth and backlog keep rising	Revenue grows slower than capex
Gemini	Carbon opens to developers and users keep growing	New models slip again
World models	Real results from robots trained with Genie are published	Stays at demos, never becomes products
People	Key researchers stay and new talent arrives	Moves to rivals continue

Where do winning thoughts come from?

A 12-year-old chess player's question is now the question of an industry worth hundreds of billions of dollars

Terms and sources

Terms

Transformer — A neural-network design that computes the relationships among all words in a passage at once; the basis of today's chatbots.

Coding agent — An AI tool that plans, writes and tests code on its own.

Benchmark — A shared test for comparing AI performance; rankings shift easily with method.

TPU — Google's own chip designed for AI computation.

Multimodal — Handling several kinds of information together, such as text, images, audio and video.

World model — AI that pictures how the world works and predicts or simulates what happens next.

Token — The small unit AI uses to process text, roughly a word fragment.

Backlog — Contracted revenue not yet recognized.

The ChatGPT shock and Google's response

arXiv 1706.03762 (Jun 2017, transformer) · Google blog (May 2021 LaMDA, Apr 20 2023 merger, Dec 6 2023 Gemini) · Altman on X (Dec 5 2022) · Reuters (Feb 1 2023 UBS estimate, Feb 8 2023 Bard) · New York Times (Dec 21 2022, code red) · Bloomberg (Dec 7 2023, demo video) · Google blog (Feb 23 2024, image generation pause)

Comeback and role change

Google DeepMind (Mar 2025 Gemini 2.5, Jul 2025 IMO) · TechCrunch (Sep 16 2025, App Store No. 1) · Google blog (Nov 18 2025, Gemini 3) · Fortune (Dec 2 2025, Aug 5 2026) · Axios, Semafor (Aug 5 2026) · TIME (Aug 6 2026) · 9to5Google (Sep 24 2026) · Alphabet results (Q4 2025 to Q3 2026, SEC 8-K) · Alphabet Q3 call (Oct 28 2026) · Isomorphic Labs (Jan 2024 Eli Lilly and Novartis deals, Mar 2025 funding) · OpenAI (Mar 2023, GPT-4)

Gemini 4 and Genie

VentureBeat (Sep 30 2026) · Google blog (Oct 13 2026 Argon, Oct 27 2026 Genie 4) · Pichai on X (Oct 28 2026) · Tom's Hardware (Nov 3 2026) · TechCrunch (Oct 27 2026) · arXiv 2402.15391 (Genie 1) · Wikipedia (Genie) · Silicon Republic (AMI Labs) · TechCrunch (Feb 18 2026, World Labs) · GamesBeat (Cosmos 3) · Google DeepMind (Nov 2025, SIMA 2)

OpenAI, Anthropic and AGI outlooks

Fortune (Sep 3 2026 GPT-6; May 28 2026; Nov 2 2026, amended S-1) · aiweekly (Oct 2026, WSJ summary) · 9to5Mac (Oct 26 2026) · TechCrunch (Sep 22 2026, Opus 5.5) · Reuters, Axios (Mar 24 2026, Sora) · Dataconomy (Sep 7 2026) · Anthropic S-1 (Oct 26 2026, amended Nov 2) · Axios (Oct 8 2026) · Bloomberg (Nov 6, Apr 24, Mar 31 2026) · Anthropic (Apr 6 2026), Broadcom 8-K · CNBC (Apr 20 2026) · OpenAI (Jan 2025, Stargate) · Zvi Mowshowitz (Sep 14 2026, quoting Amodei) · TechSpot (Anthropic CFO) · OfficeChai (Sep 13 2025) · World Economic Forum (Jan 20 2026) · TIME (Jan 8 2025) · Fortune (Jan 27 2026) · Fast Company (May 22 2026) · WSJ Magazine (Oct 21 2026, via Yahoo) · companiesmarketcap

The outline of this piece drew on a published video commentary; every number and fact was checked against the sources above, and some figures in the video (such as benchmark names) were corrected to match the original sources. It reflects information public as of November 8 2026. Real people are shown as silhouettes, not photos.

The editor chooses the topic and angle and reviews and revises each column before publishing. AI tools were used for research, drafting and the infographics, and every number was checked against its source. The AI tool used (Claude) is made by Anthropic, one of the companies compared here. We tried to treat all three companies by the same standard, but please read with that in mind.