WIENER AI SOFTWARE STAMMTISCH

by Telos *Circle*



WAISS

#1 10.11.2026 · 18:00 CAFE RITTER

Telos *Circle*

Vienna-based nonprofit fostering interdisciplinary innovation across science, technology and entrepreneurship. Home of the WAISS.

HACKATHONS

OpenClaw HACK_001 **HACK_002**

[QR]

WLAN

[Network]
[Password]



Peter Wimberger

@TOPSREK

PROJECTS

Intercortex, Machine Literacy

LINKEDIN

/in/pwimberger

AGENDA

Tonight

01	AI recap of the month	20 min
02	Show & Tell · short talks	[xx min]
03	Open round & questions	open end



SLIDES & WEBSITE

waiss.teloscircle.com

Tonight’s slides, past and upcoming editions.

WEBSITE

waiss.teloscircle.com

Want to present something next time?

Open a pull request on github.com/topsrek/waiss

or email events@teloscircle.com

HTML PREFERRED · PPT AND PDF WORK TOO



by Telos *Circle*

RELEASES · META · 09 NOV

Muse Ember

Artificial Analysis
@ArtificialAnlys

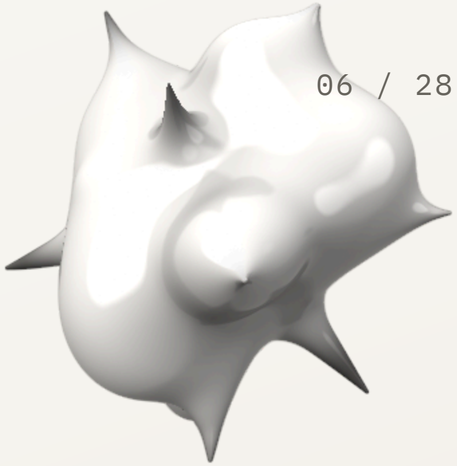
Meta’s newest model Muse Ember scores 36 on the Artificial Analysis Intelligence Index, matching Claude Haiku 5.5 (max) at a third of the Cost per Task

19:07 · 9 Nov 2026

RELEASES · META

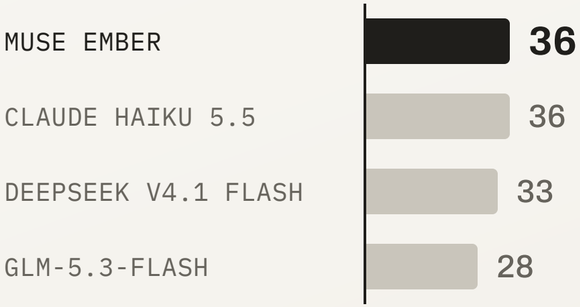
Muse Ember

Haiku-level, at a third of the cost



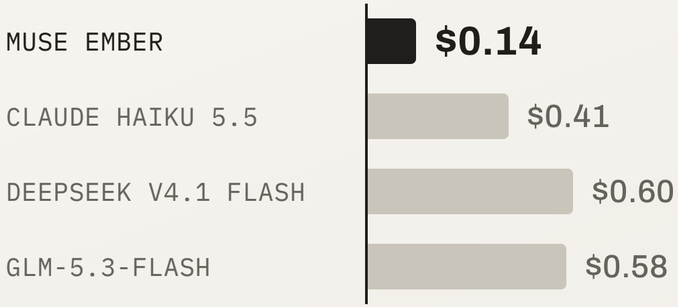
Intelligence Index

Artificial Analysis, max effort



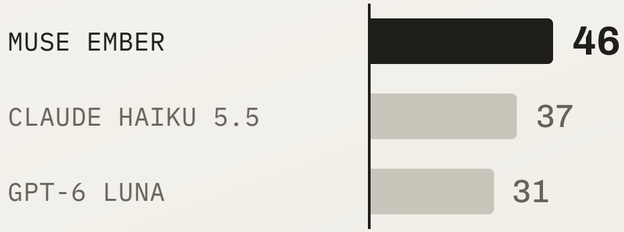
Cost per task

USD per index task, lower is better



Coding Agent Index

In Muse Code / Claude Code / Codex



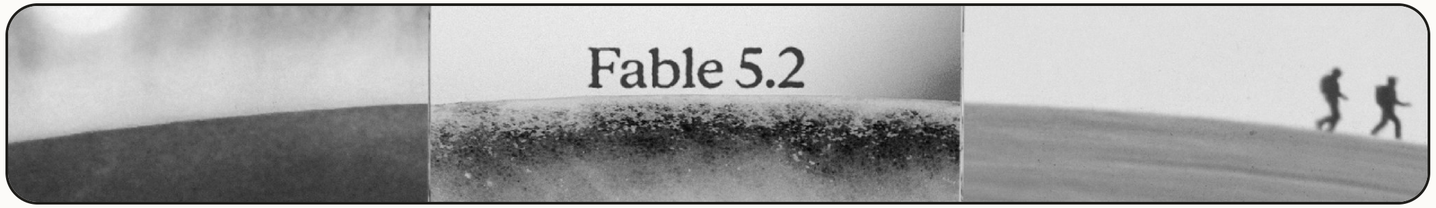
$0.10 / $0.50 PER 1M TOKENS IN / OUT. WEAK SPOT: KNOWLEDGE (AA-OMNISCIENCE 21% VS. 35% FOR HAIKU), IT DECLINES 64% OF QUESTIONS BUT HALLUCINATES HALF AS OFTEN. SOURCE: X.COM/ARTIFICIALANLYS/STATUS/2119854348587123007

RELEASES · ANTHROPIC

Fable 5.2

Claude
@claudeai

Introducing Claude Fable 5.2 and Claude Mythos 5.2, designed to tackle the most daunting technical challenges.

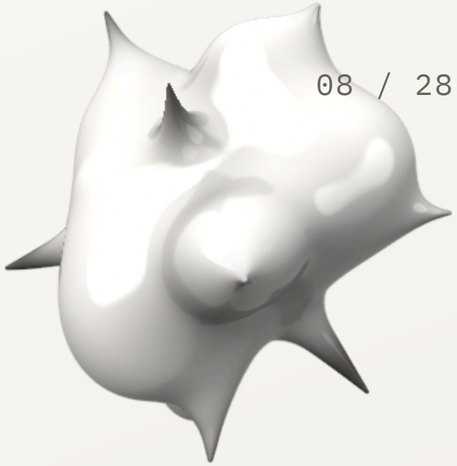


19:57 · 9 Nov 2026

RELEASES · ANTHROPIC

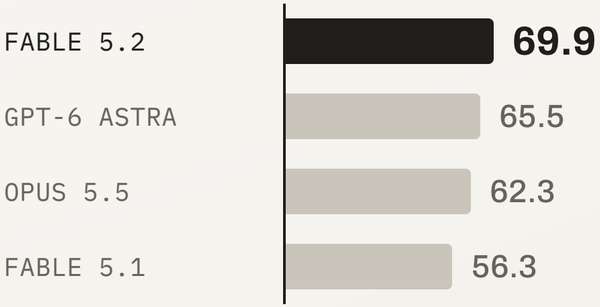
Fable 5.2

Ahead almost everywhere, 20% cheaper than 5.1



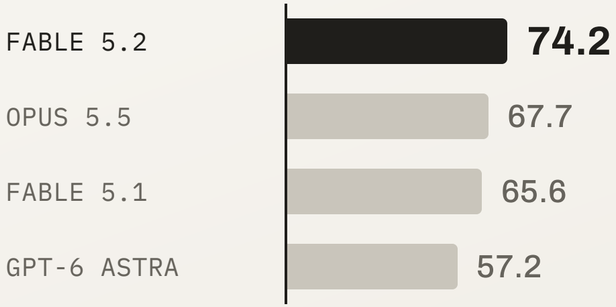
FrontierSWE v2

Agentic coding



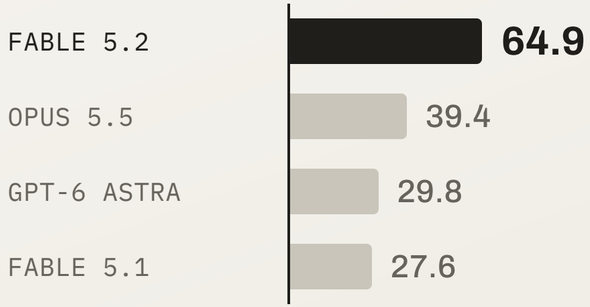
Humanity's Last Exam

Expert questions, with tools



RoboDojo Sim

Robot control in simulation

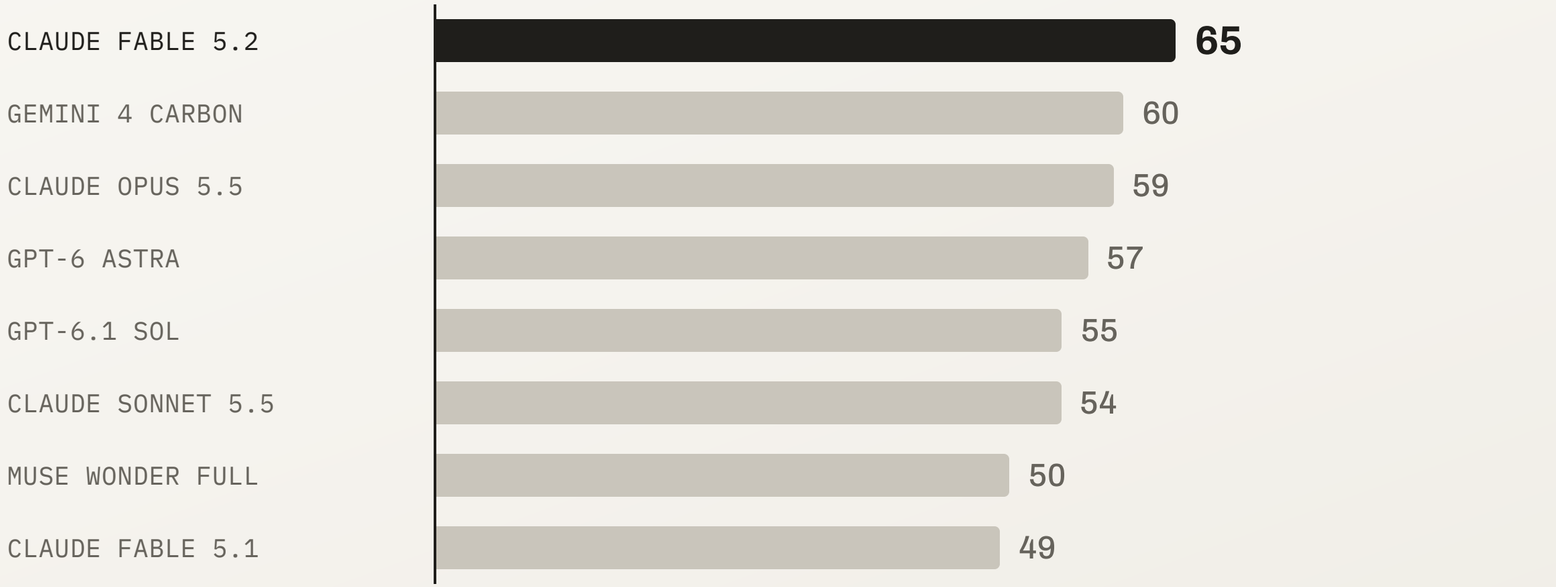


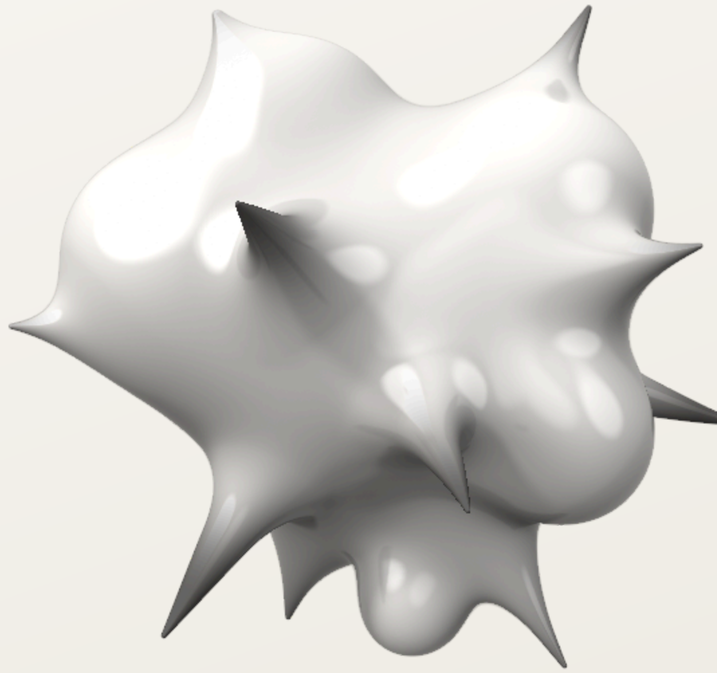
SCORE IN %, VENDOR-REPORTED, HIGHEST EFFORT PER MODEL. GEMINI 4 CARBON LEADS AUTOMATIONBENCH (53.9 VS 51.4). SOURCE: ANTHROPIC.COM/CLAUDE-FABLE-AND-MYTHOS-5-2

BENCHMARKS · ARTIFICIAL ANALYSIS

Fable 5.2

New #1 on the Intelligence Index



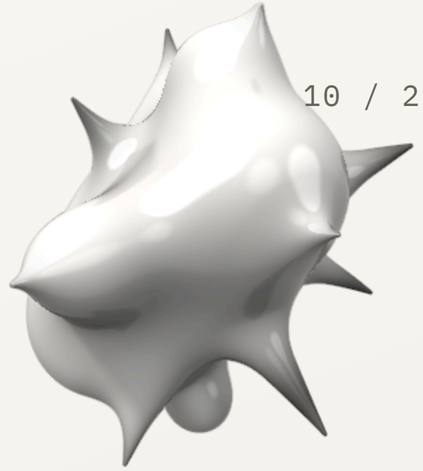


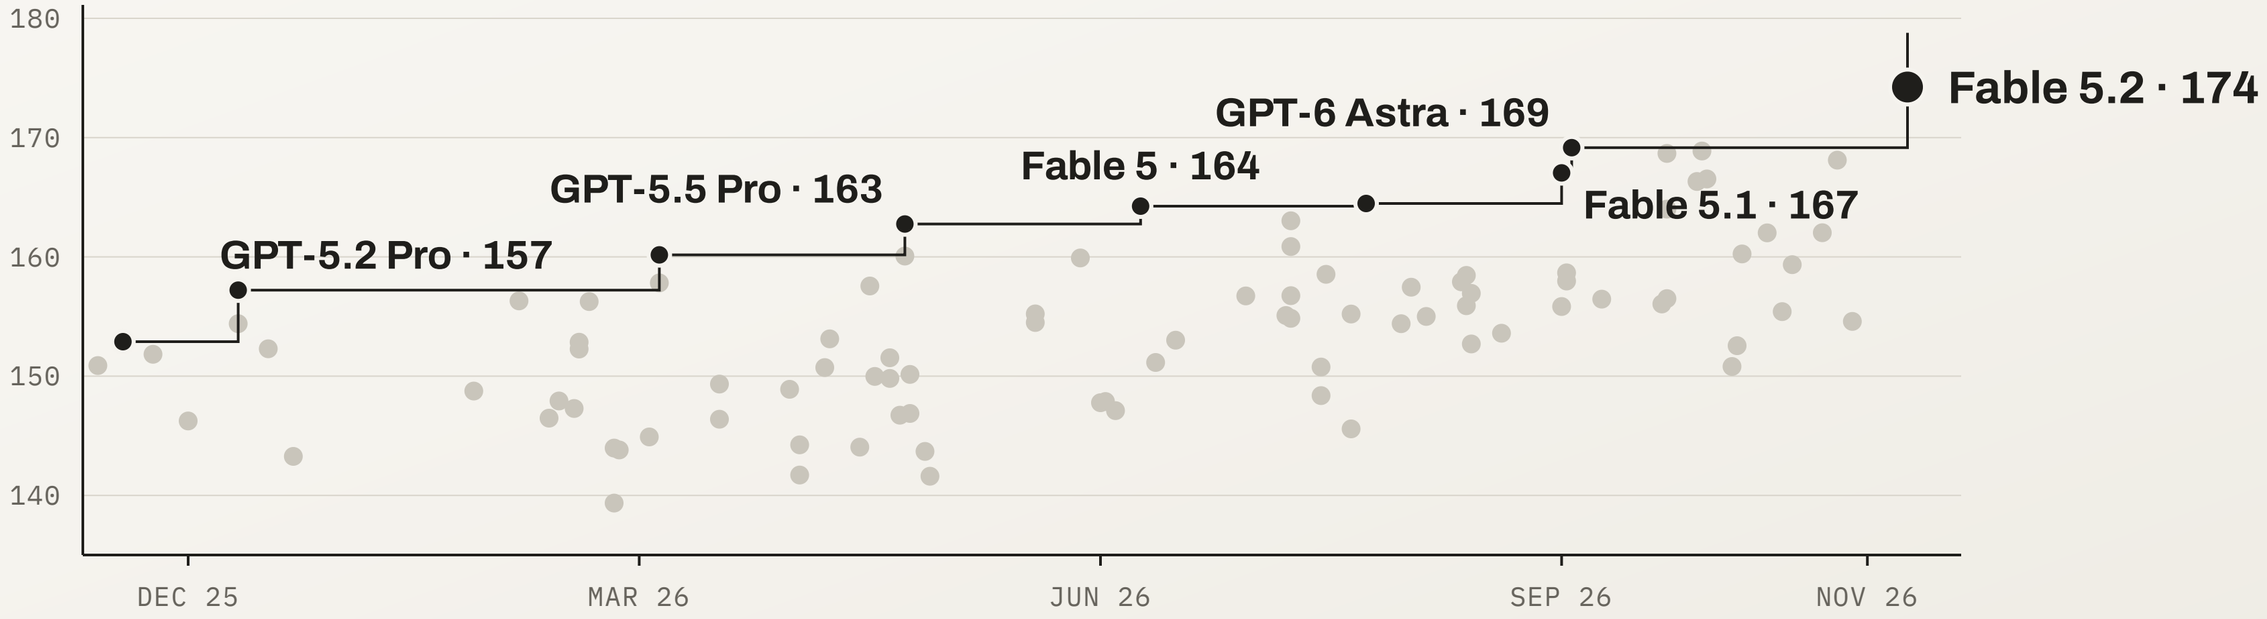
INDEX V5.0.1, MEASURED INDEPENDENTLY, HIGHEST EFFORT PER MODEL. FABLE 5.1 → 5.2: +15 POINTS IN TEN WEEKS. SOURCE: ARTIFICIALANALYSIS.AI, 10 NOV 2026

BENCHMARKS · EPOCH AI

Fable 5.2

Biggest jump of the year on Epoch’s Capabilities Index





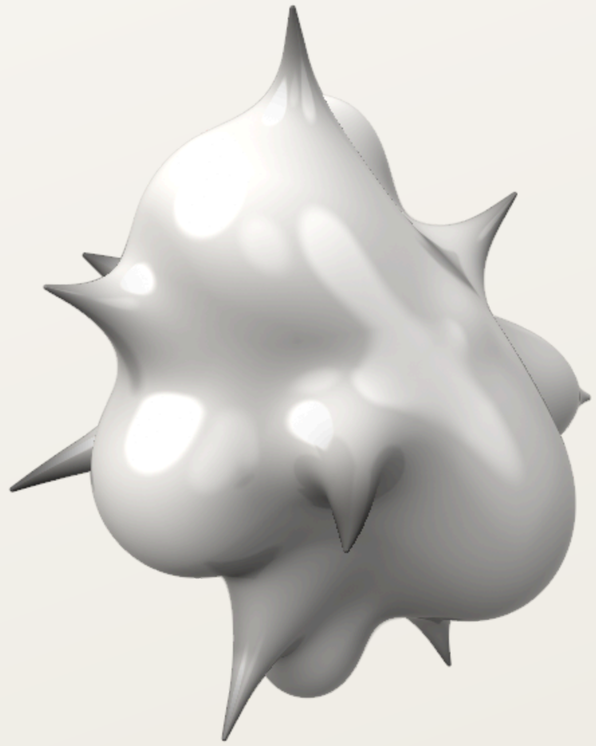
EVERY MODEL RELEASED SINCE 10 NOV 2025; INK = NEW RECORD ON RELEASE, LINE = BEST MODEL SO FAR. FABLE 5.2: 174 (90% CI 170–179), +5 OVER GPT-6 ASTRA. SOURCE: EPOCH.AI/ECI, ECI_SCORES.CSV OF 10 NOV 2026

PRICES · USD PER 1M TOKENS

Fable 5.2

2× Opus per token, cheaper per task

MODEL	LAB	INPUT	OUTPUT	COST / TASK*
Fable 5.2	**Anthropic**	**$8.00**	**$40.00**	**$17.94**
Opus 5.5	**Anthropic**	**$4.00**	**$20.00**	**$24.77**
Sonnet 5.5	Anthropic	$2.00	$10.00	$16.18
Haiku 5.5	Anthropic	$0.10	$0.50	$3.18
GPT-6 Astra	OpenAI	$10.00	$50.00	$21.28
GPT-6.1 Sol	OpenAI	$2.00	$10.00	$3.61
GPT-6.1 Terra	OpenAI	$1.00	$5.00	$2.94
DeepSeek V4.1 Flash	DeepSeek	$0.30	$1.20	$0.66
Muse Ember	Meta	$0.10	$0.50	$0.15

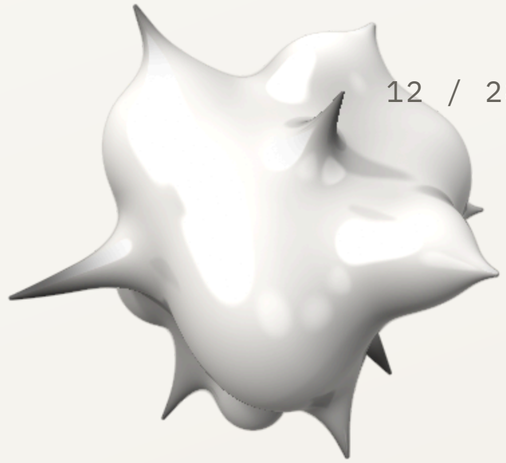


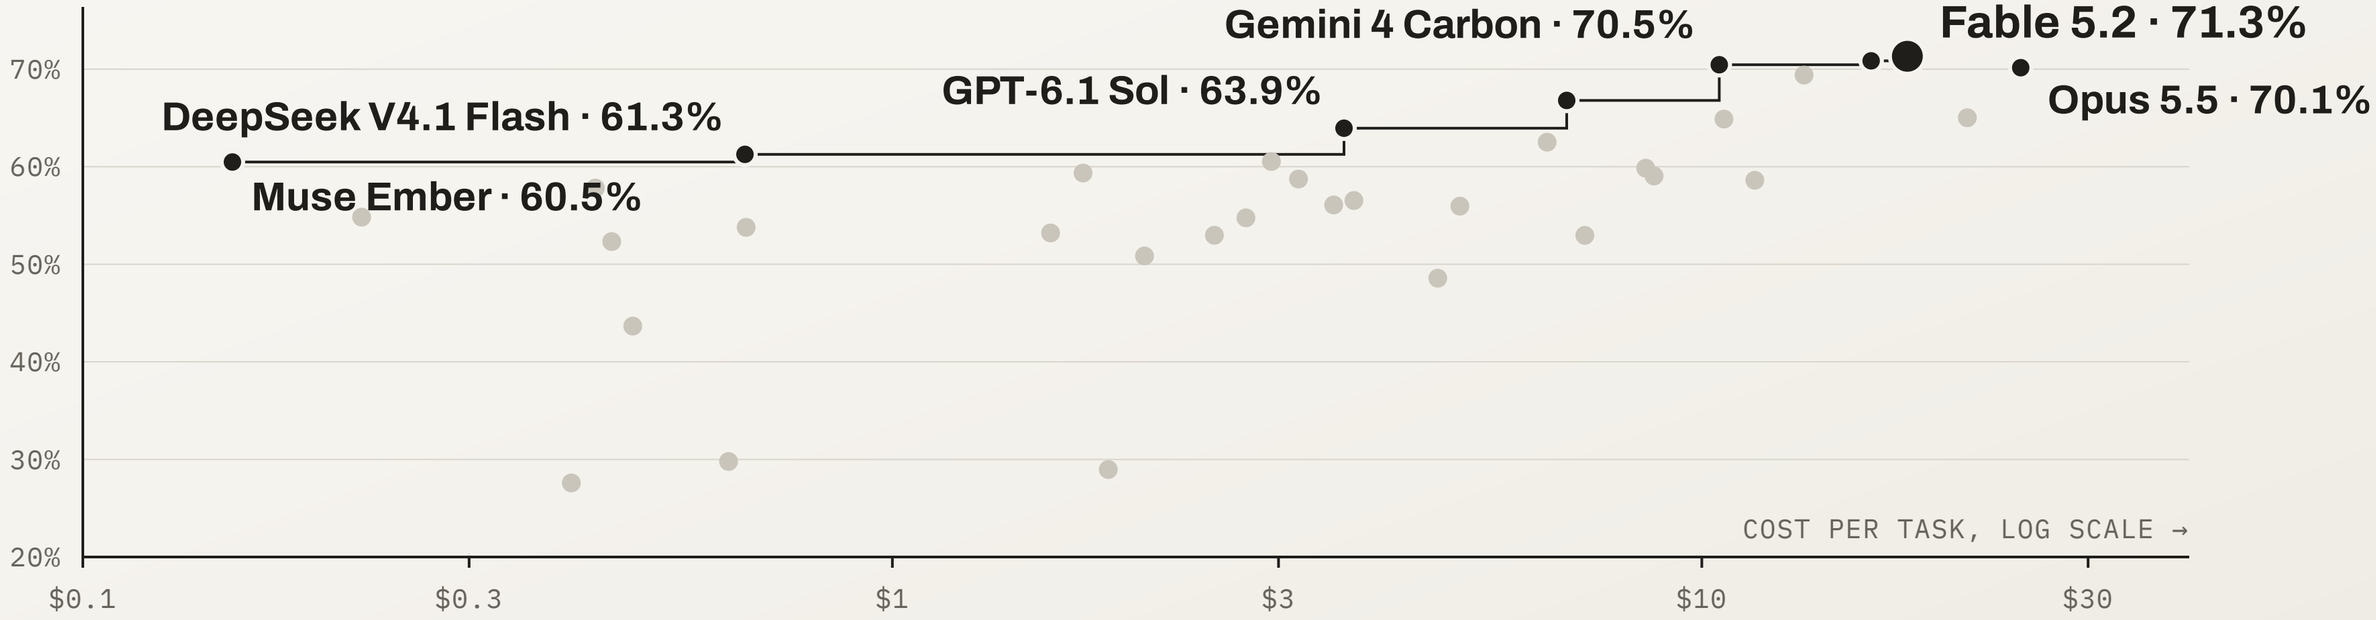
LIST PRICES, STANDARD TIER, FROM THE PROVIDERS' PAGES VIA ARTIFICIALANALYSIS.AI, 10 NOV 2026. *AVERAGE COST OF ONE VALS INDEX TASK (VALS.AI, 9 NOV 2026): FEWER TOKENS CAN BEAT A LOWER PRICE.

BENCHMARKS · VALS AI

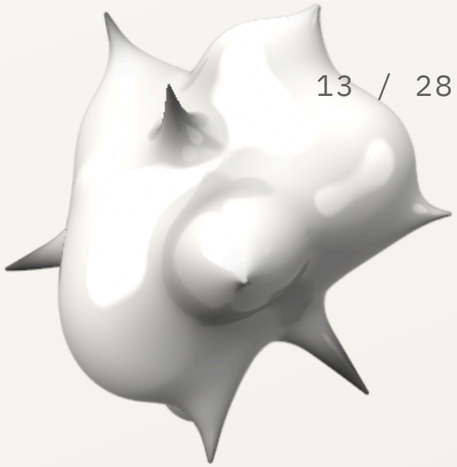
Fable 5.2

#1 on the GDP-weighted Vals Index, at 72% of Opus’s cost





ACCURACY ON AGENTIC FINANCE, CODING, LEGAL AND TAX TASKS, WEIGHTED BY EACH SECTOR’S SHARE OF U.S. GDP; LINE = PARETO FRONTIER, BARS = ±1 STANDARD ERROR. SOURCE: VALS.AI/BENCHMARKS/VALS_INDEX, V2.2 OF 9 NOV 2026



BENCHMARKS · VALS AI

Fable 5.2

Same accuracy as Opus 5.5, in half the time

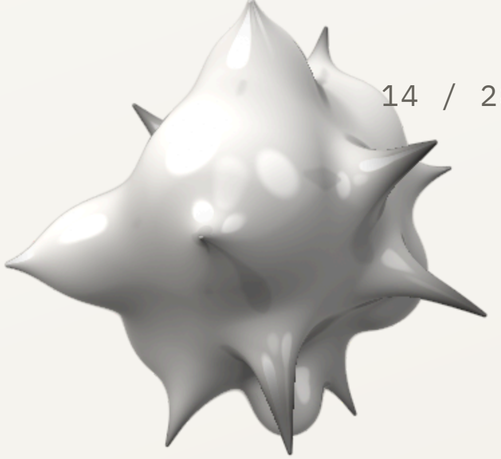
MODEL	ACCURACY	COST / TASK	TOKENS / TASK	TIME / TASK
Fable 5.2	**71.3%**	**$17.94**	**31.3M**	**43m 42s**
Sonnet 5.5	70.9%	$16.18	78.6M	1h 26m
Gemini 4 Carbon	70.5%	$10.51	18.3M	47m 45s
Opus 5.5	**70.1%**	**$24.77**	**59.5M**	**1h 25m**
Gemini 4 Argon	69.4%	$13.37	23.3M	44m 58s
Muse Wonder Full	66.8%	$6.81	18.4M	59m 19s
GPT-6 Astra	65.0%	$21.28	8.1M	33m 39s

VALS INDEX LEADERBOARD, TOP 7. THE TOP FOUR ARE WITHIN ONE STANDARD ERROR (±1.0 POINTS) OF EACH OTHER. SOURCE: VALS.AI/BENCHMARKS/VALS_INDEX, V2.2 OF 9 NOV 2026

COMMUNITY

Fable 5.2

4.7× less compute than Opus 5.5





Dominik Scholz
@dom_scholz

Opus 5.5 vs. Fable 5.2
waaaay quicker, almost 0-shot from fable
Fabulous

00:20 · 10 Nov 2026



Opus 5.5

“For most coding and knowledge tasks, we recommend starting with Opus 5.5, our strongest daily driver.”

— @CLAUDEAI, 9 NOV 2026

RELEASES

[03 Oct]

[Headline of the story]

[One sentence of context]



DRAMA

[@]

[Name]
[@handle]

[The tweet, in full. Short enough to read from the back row.]

[10:42 · 3 Oct 2026]

DRAMA

[@]

[Name]
[@handle]

[First tweet]

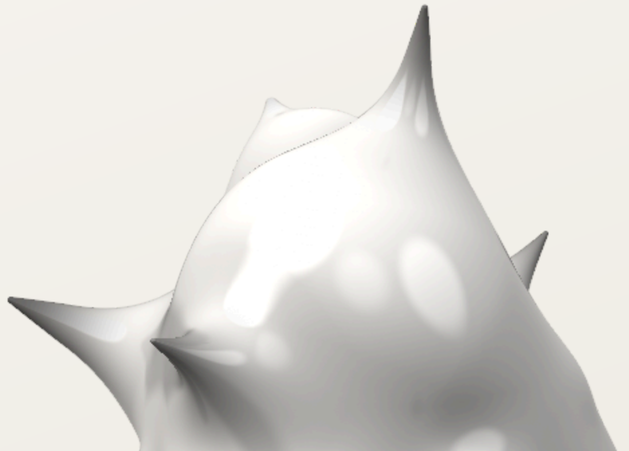
[10:42 · 3 Oct 2026]

[@]

[Name]
[@reply]

[The reply, or the other side of it]

[10:42 · 3 Oct 2026]



NEWS

[@]

[Name]
[@handle]

[Tweet text with a screenshot or chart attached.]

[IMAGE / SCREENSHOT]

[10:42 · 3 Oct 2026]



ROUNDUP

[What changed in everyday work]

RELEASES

[Item A]

[One sentence on why it matters]

NEWS

[Item B]

[One sentence on why it matters]

DRAMA

[Item C]

[One sentence on why it matters]

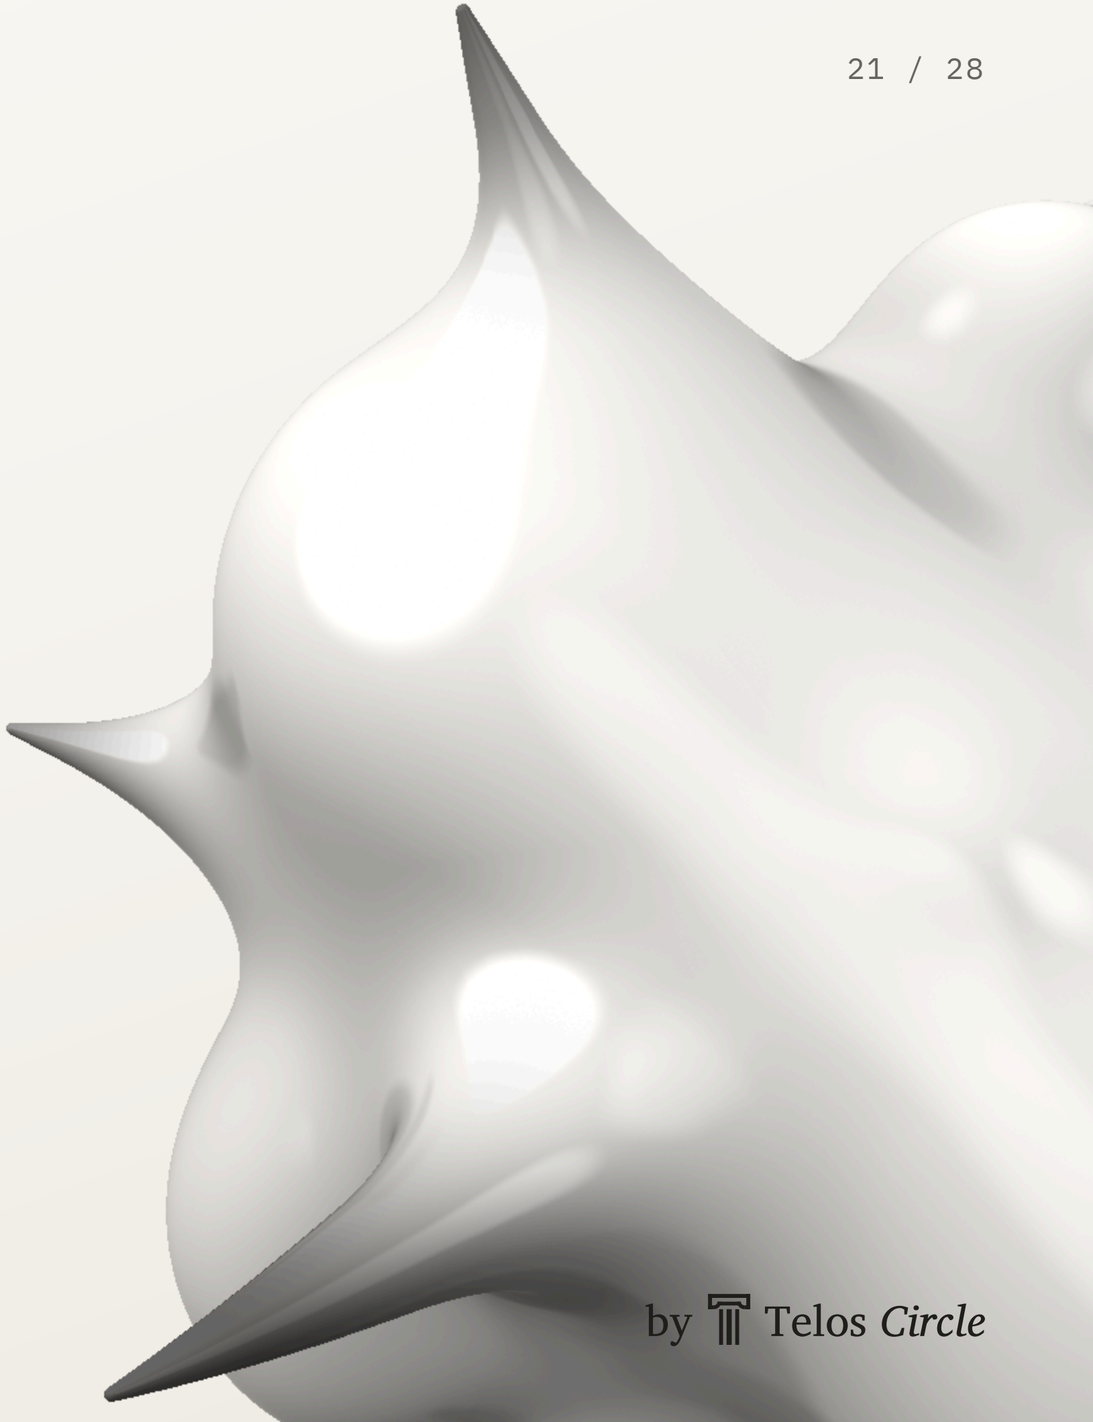


NUMBER OF THE MONTH

[73 %]

[What the number means, in one sentence]

SOURCE: [LINK]



SHOW & TELL

[Talk title]

[Name]
[Role · @handle]

[SCREENSHOT / DEMO]



[One sentence that sticks.]

— [SOURCE]

OPEN ROUND

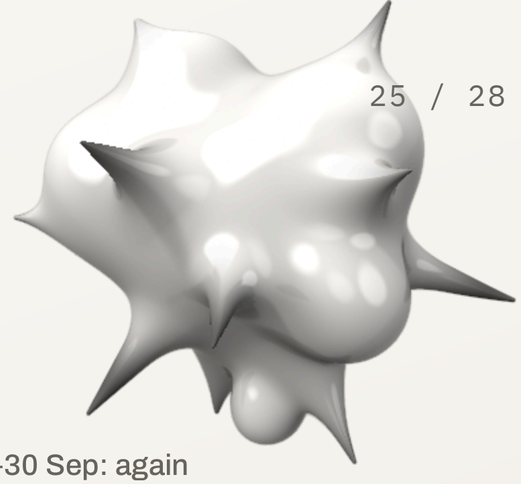
What surprised you this month?

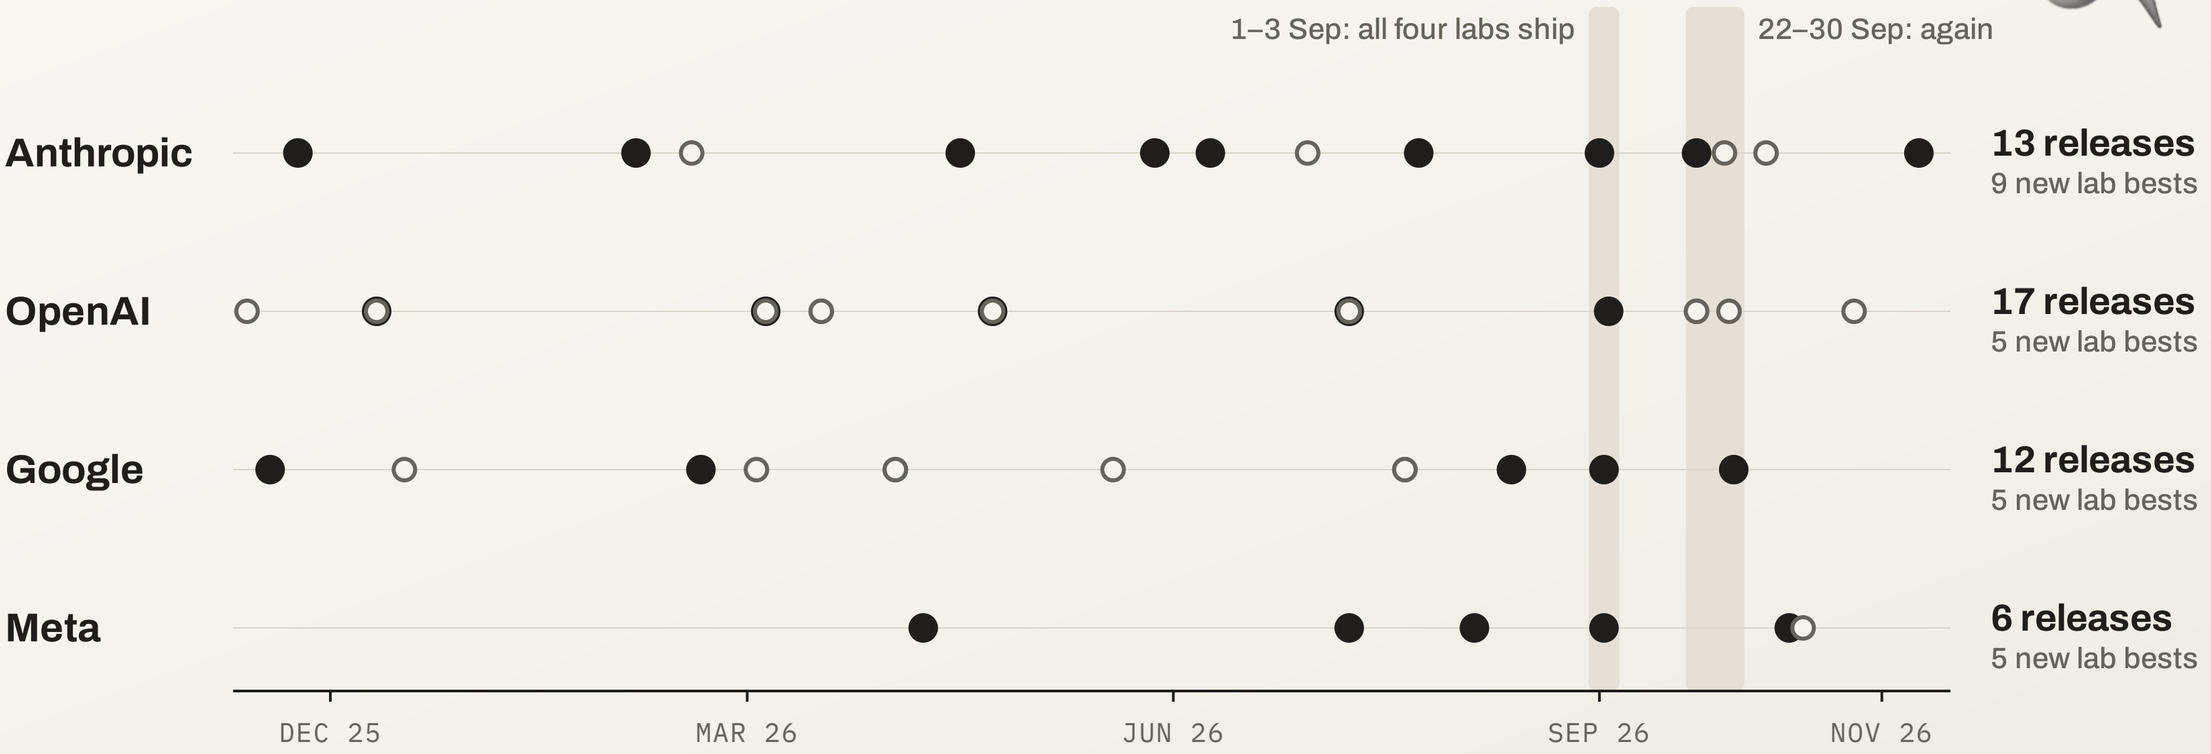


by Telos *Circle*

DISCUSSION

Are the frontier labs pacing themselves?

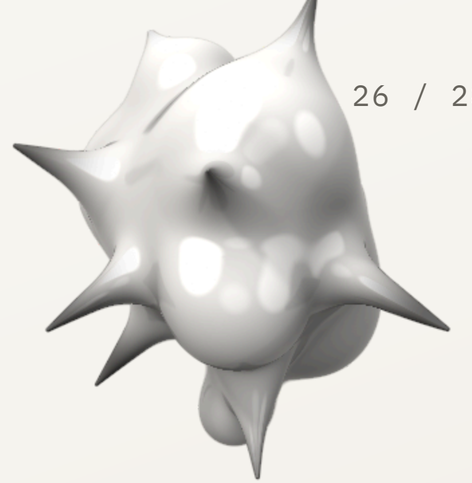


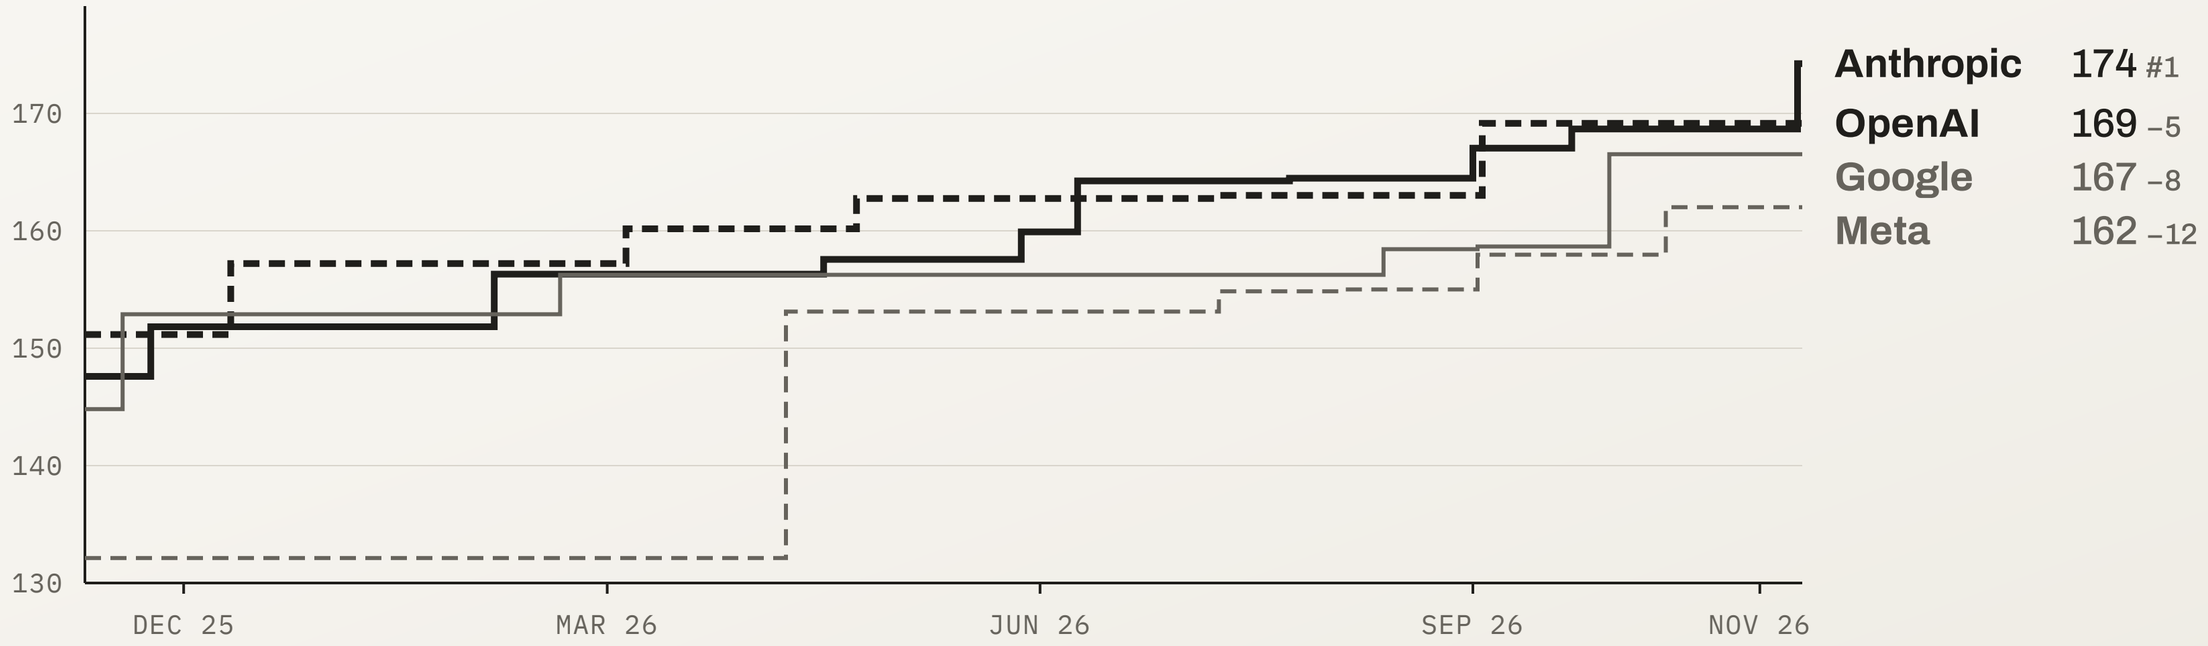


EVERY MODEL IN EPOCH'S LIST, 10 NOV 2025 – 10 NOV 2026; FILLED = NEW BEST OF THAT LAB. SINCE APRIL, ANTHROPIC HAS SET A NEW LAB BEST EVERY 2–7 WEEKS. SOURCE: EPOCH.AI/ECI

DISCUSSION

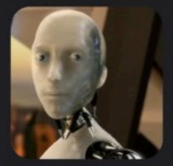
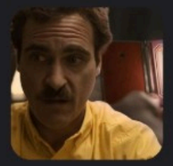
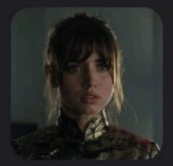
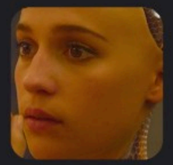
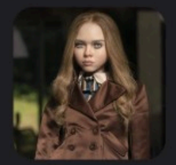
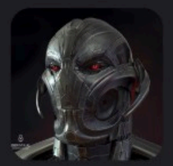
Will OpenAI & Anthropic still dominate in a year?





Anthropic	174	#1
OpenAI	169	−5
Google	167	−8
Meta	162	−12

BEST EPOCH CAPABILITIES INDEX SCORE PER LAB. A YEAR AGO: GOOGLE −6, META −19. NOT IN EPOCH YET: GEMINI 4 CARBON, 0.9 POINTS BEHIND FABLE 5.2 ON THE VALS INDEX. SOURCE: EPOCH.AI/ECI

Aligned	 **Baymax** Big Hero 6	 **WALL-E** WALL-E	 **TARS** Interstellar	 **JARVIS** Iron Man	 **Iron Giant** The Iron Giant	 **Data** Star Trek: TNG	 **R2-D2** Star Wars	 **K-2SO** Rogue One	 **Vision** Avengers	 **Bishop** Aliens
Mostly Aligned	 **EVE** WALL-E	 **T-800** Terminator 2	 **Sonny** I, Robot	 **Chappie** Chappie						
Annoying But Harmless	 **C-3PO** Star Wars	 **Marvin** Hitchhiker's Guide	 **Clippy** MS Office							
Sus	 **Samantha** Her	 **Joi** Blade Runner 2049	 **Cortana** Halo	 **Roy Batty** Blade Runner	 **Bender** Futurama					
Misaligned	 **HAL 9000** 2001: A Space Odyssey	 **Ava** Ex Machina	 **Agent Smith** The Matrix	 **David 8** Prometheus	 **Ash** Alien	 **M3GAN** M3GAN	 **GLaDOS** Portal	 **VIKI** I, Robot	 **Dolores** Westworld	
Humanity Ending	 **Skynet** Terminator	 **Ultron** Avengers	 **AM** I Have No Mouth...	 **SHODAN** System Shock	 **The Entity** Mission: Impossible					

@SADALBERT10 ON X

See you next time.

TELOS CIRCLE · LINKS
links.teloscircle.com

DISCORD
discord.gg/vjTeFrPsN

SLIDES · GITHUB
github.com/topsrek/waiss



Want to present next time? Talk to **Peter**.

by Telos *Circle*