THE MOJO ANGLE · RESEARCH BRIEF · NOVEMBER 2026

The Laws Behind the Magic

The Mojo Horizon Clock: how long a task frontier AI can complete on its own, the scaling laws behind it, and how close AGI really is.

Index: **Mojo Horizon Clock** · AIMOJO Research Desk · Edited by Aliakbar Fakhri

~20,000x
growth in 50% task horizon, GPT-2 (2019) to Mythos Preview (2026)

~4.1 months
doubling time since 2023 (AIMOJO fit; METR: ~4.3)

~5x a year
growth in frontier training compute (Epoch AI)

Mojo's verdict. AI progress follows laws as regular as physics, and each has a limit. The digital half of AGI is a live possibility before 2030; the physical half is much further away.

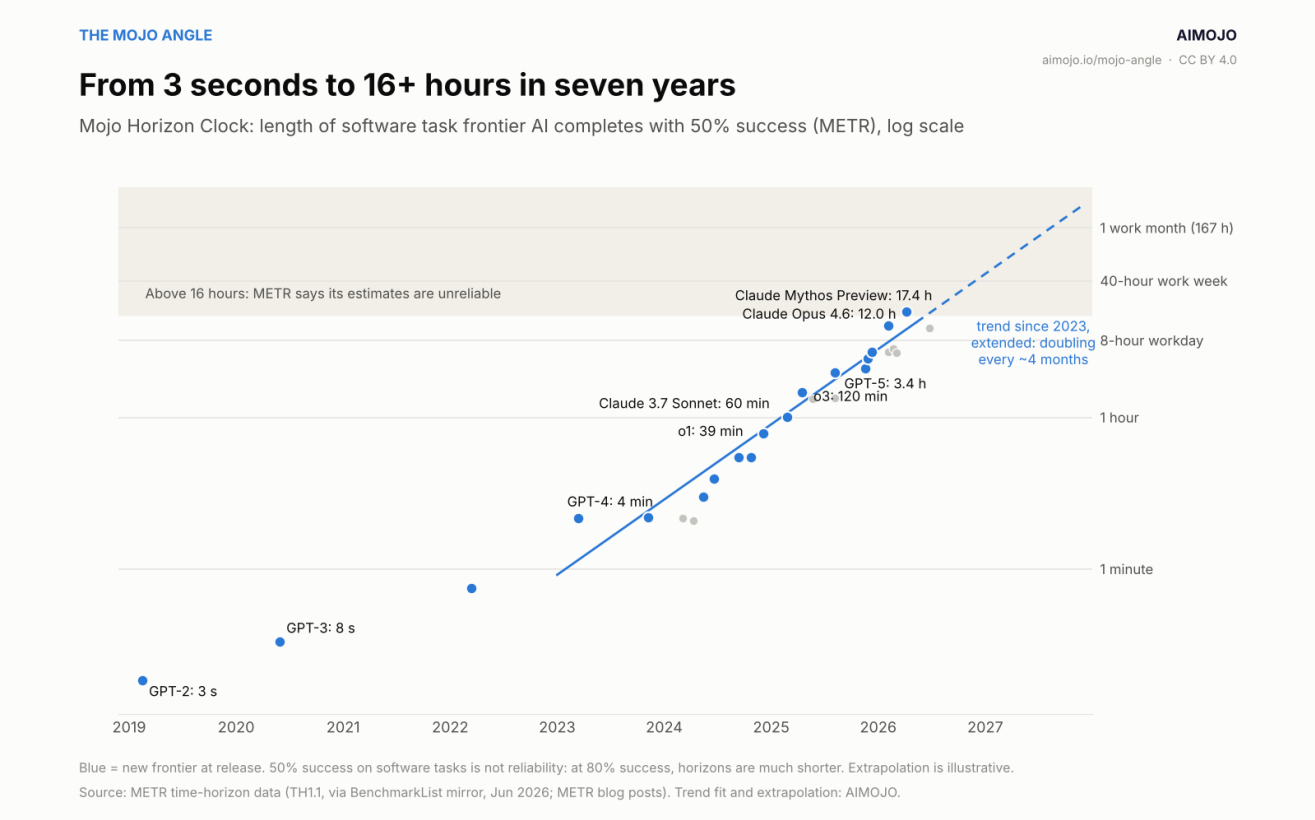


Figure 1. The Mojo Horizon Clock (METR data, AIMOJO fit).

Key findings

- Neural scaling law: loss falls as a power law of compute (alpha about 0.050 for compute-efficient training), so 10x compute cuts loss by about 11%. Chinchilla: about 20 training tokens per parameter.
- Effective compute grows very roughly 14x a year (5x compute times about 2.8x from algorithms), by AIMOJO's arithmetic.
- METR 50% time horizons rose from about 3 seconds (GPT-2, 2019) to at least 16 hours (an early version of Claude Mythos Preview, released April 2026; METR estimate, May 2026). AIMOJO's frontier fit to METR's data, consistent with METR's own figure, gives a doubling time of about 6 months over 2019-2026 and about 4.1 months since 2023.
- Extended illustratively, the post-2023 trend crosses a 40-hour work week around end-2026 and a 167-hour work month in the second half of 2027; METR says estimates above 16 hours are unreliable.
- ARC-AGI-3 went from 0.51% (March 2026) to 62.7% verified on the standard harness (September 2026); humans score 100%.
- Physics: brain about 20 W vs about 700 W for one H100; Landauer limit about 3e-21 J per bit; IEA projects data-centre electricity rising from about 415 TWh (2024) to about 945 TWh (2030).

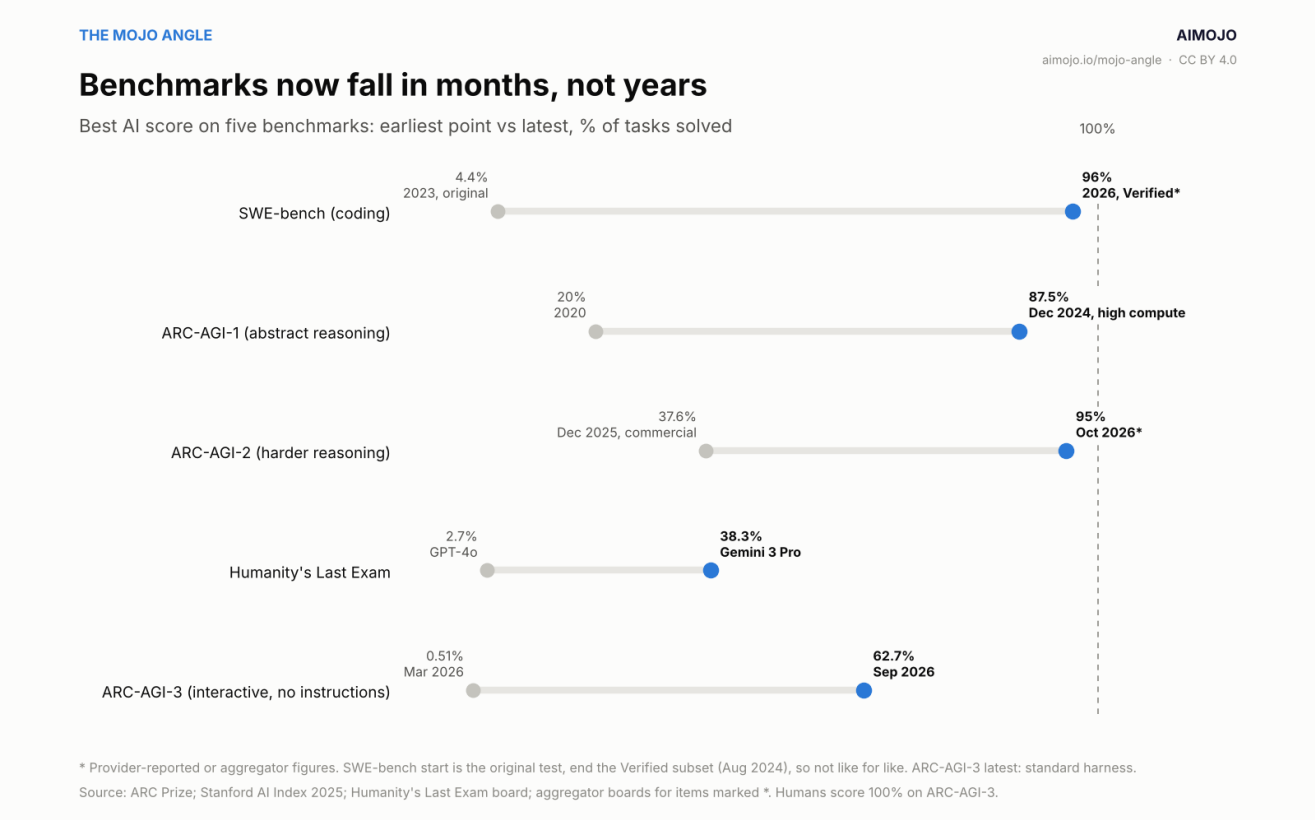


Figure 2. Best AI scores on five benchmarks, earliest vs latest.

Mojo Horizon Clock data (METR TH1.1 and METR posts)

Frontier = new record at release. Values above 16 h are flagged by METR as unreliable.

Model	Release	50% horizon (min)	Frontier	Note
GPT-2	2019-02-14	0.05	yes	
GPT-3 (davinci-002)	2020-05-28	0.14	yes	
GPT-3.5 (gpt-3.5-turbo-instruct)	2022-03-15	0.6	yes	
GPT-4 (0314)	2023-03-15	3.99	yes	
GPT-4 (1106)	2023-11-06	4.04	yes	
Claude 3 Opus	2024-03-04	3.95		
GPT-4 Turbo	2024-04-09	3.73		
GPT-4o	2024-05-13	6.99	yes	
Claude 3.5 Sonnet (Jun 2024)	2024-06-20	11.4	yes	
o1-preview	2024-09-12	20.33	yes	
Claude 3.5 Sonnet (Oct 2024)	2024-10-22	20.52	yes	
o1	2024-12-05	38.83	yes	
Claude 3.7 Sonnet	2025-02-25	60.39	yes	
o3	2025-04-16	119.73	yes	
Claude Opus 4	2025-05-22	100.37		
Claude Opus 4.1	2025-08-05	100.47		
GPT-5	2025-08-07	203.01	yes	
Gemini 3 Pro	2025-11-18	224.33	yes	
GPT-5.1-Codex-Max	2025-11-19	223.71		
Claude Opus 4.5	2025-11-25	292.99	yes	
GPT-5.2	2025-12-11	352.25	yes	
Claude Opus 4.6	2026-02-05	718.81	yes	
GPT-5.3-Codex	2026-02-05	349.53		
Gemini 3.1 Pro (Preview)	2026-02-20	384.15		
GPT-5.4	2026-03-05	341.74		
Claude Mythos Preview (early)	2026-04-07	1044.78	yes	Above 16 h: METR says estimates are unreliable
GPT-5.6 Sol	2026-06-26	678		

Method and limits

Horizon values come from METR's Time Horizon 1.1 data as mirrored by BenchmarkList (30 Jun 2026) and from METR's own blog posts for later models. Frontier points are models that set a new record at release.

We fit an ordinary least-squares line to log2(horizon) against release date for frontier points, over 2019-2026 and over 2023 onwards. Doubling time = 12 / slope (months). Extrapolations are illustrative only.

Benchmark scores combine official boards (ARC Prize, Humanity's Last Exam, AI Index) and aggregator boards for provider-reported 2026 figures, which are flagged. The SWE-bench start point (4.4%, 2023) is the original test and the end point is the Verified subset introduced in August 2024, so that pair is not like for like; the dataset's note column explains each point.

Limits: METR measures 50% success on self-contained software and research tasks; reliability at 80% success is much lower, and METR stopped actively updating its leaderboard in September 2026.

How to cite

AIMOJO, The Mojo Angle, Mojo Horizon Clock, November 2026. aimojo.io/mojo-angle/ai-laws-agi. Licensed CC BY 4.0.

Licence and disclosure

Data and charts are released under Creative Commons Attribution 4.0 (CC BY 4.0): reuse freely with credit and a link to the source page. Research and drafting were AI-assisted; every figure was checked against the primary source listed, and the verdict was approved by the named editor. Forecasts and the index scores are AIMOJO's own judgements, not facts.

Sources

1. Kaplan et al. (2020). https://arxiv.org/abs/2001.08361
2. Hoffmann et al. (2022). https://arxiv.org/abs/2203.15556
3. METR, Time Horizon 1.1 (Jan 2026). https://metr.org/blog/2026-1-29-time-horizon-1-1/
4. METR time horizons page. https://metr.org/time-horizons/
5. Epoch AI, Trends. https://epoch.ai/data/trends
6. ARC Prize, GPT-6 Astra verification. https://arcprize.org/blog/astra
7. Grace et al. (2024). https://arxiv.org/abs/2401.02843

Changelog

November 2026: first published.