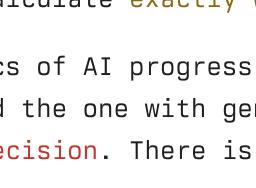


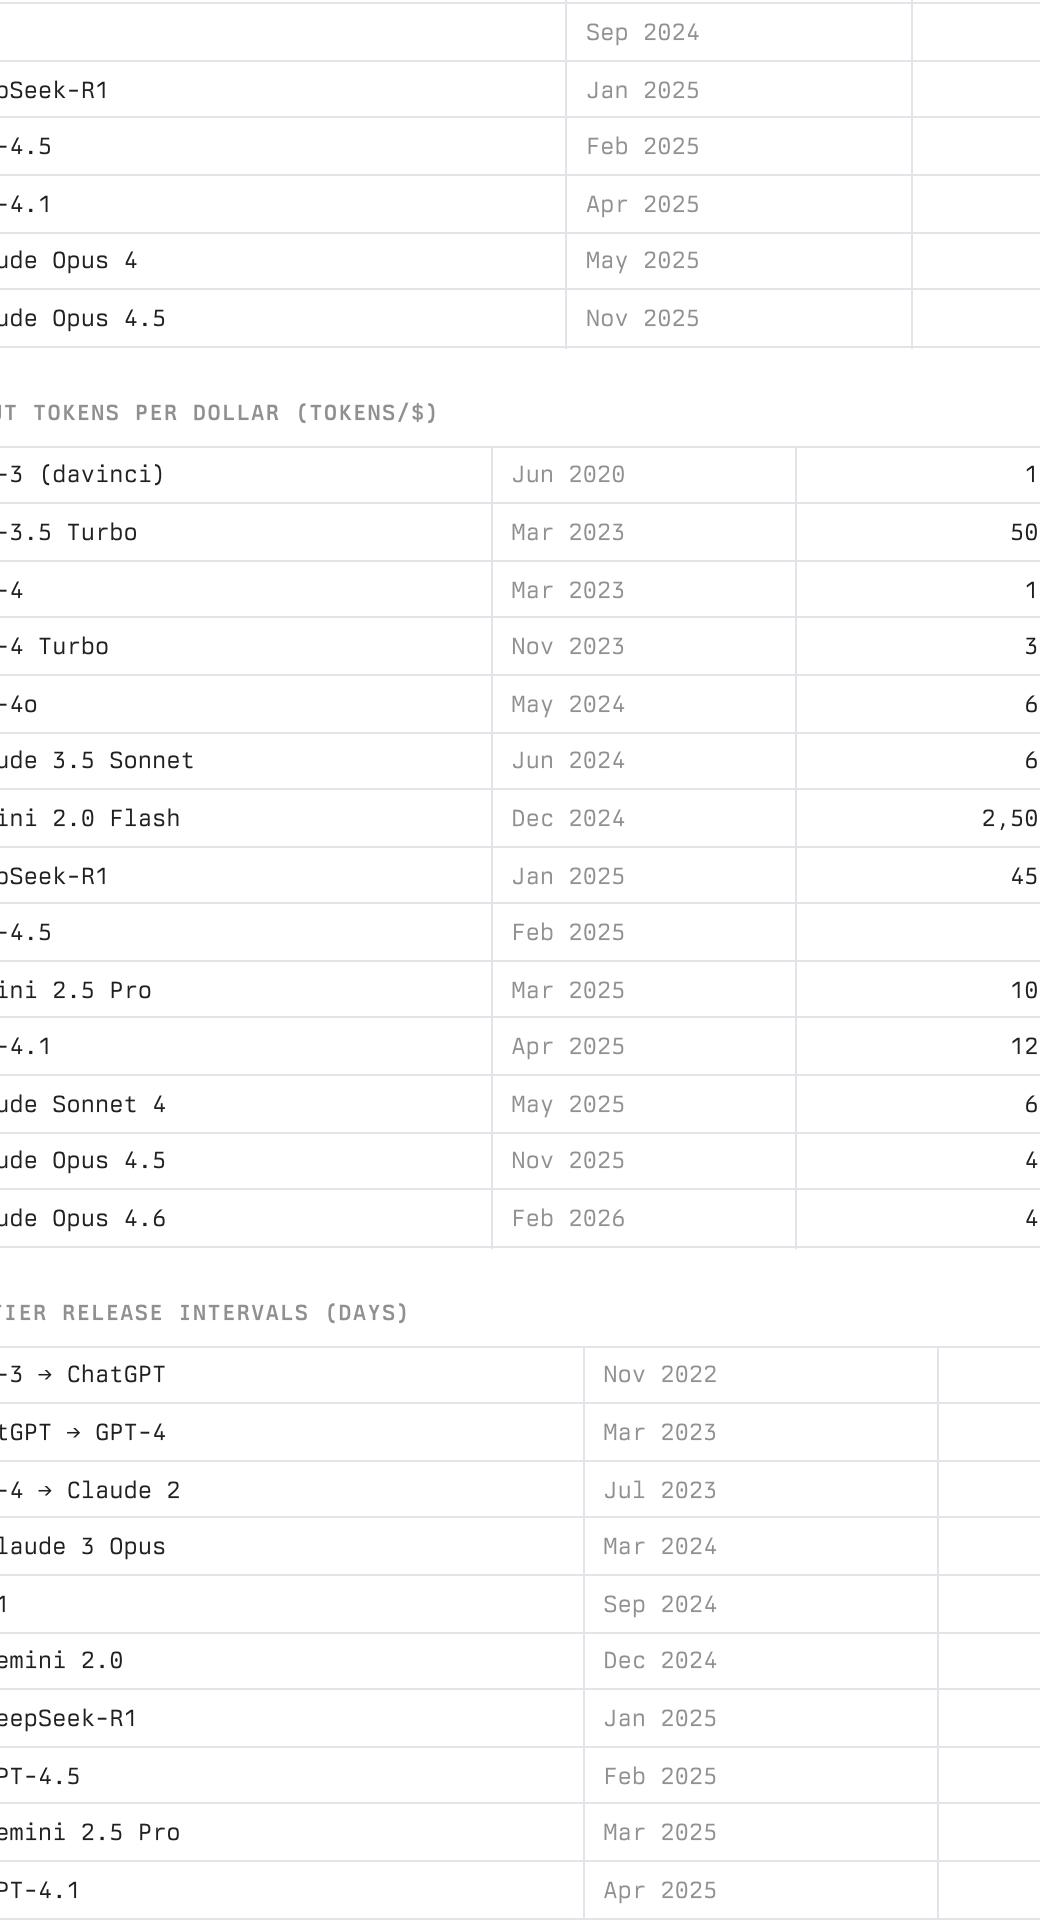


Cam Pedersen  
Engineer, Optimist



## The Singularity will Occur on a Tuesday

February 10, 2026



"Wait, the singularity is just humans freaking out?" "Always has been."

Everyone in San Francisco is talking about the singularity. At dinner parties, in coffee shops, at the OpenClaw meetup where Ashton Kutcher showed up for some reason. The conversations all have the same shape: someone says it's coming, someone says it's hype, and nobody has a number.

This seems like the wrong question. If things aren't accelerating, it's when. And if it's accelerating, we can calculate exactly when.

I collected five real metrics of AI progress, fit a hyperbolic model to each one independently, and found the one with genuine curvature toward a pole. The date has millisecond precision. There is a countdown.

(I am aware this is unhinged. We're doing it anyway.)

### The Data

Five metrics, chosen for what I'm calling their *anthropic significance* (anthropic here in the Greek sense ("pertaining to humans"), not the company, though they appear in the dataset with suspicious frequency):

1. MMLU scores: the SAT for language models
2. Tokens per dollar: cost collapse of intelligence (log-transformed, because the Gemini Flash outlier spans 150x the range otherwise)
3. Frontier release intervals: shrinking gap between "holy shit" moments
4. arXiv "emergent" papers (training 12mo): field excitement, measured memetically
5. Copilot code share: fraction of code written by AI

### MMLU BENCHMARK SCORES (% CORRECT)

GPT-3	Jun 2020	43.9
Chinchilla	Mar 2022	67.5
GPT-4	Mar 2023	86.4
Gemini Ultra	Dec 2023	83.7
Claude 3 Opus	Mar 2024	86.8
Claude 3.5 Sonnet	Jun 2024	88.7
o1	Sep 2024	90.8
DeepSeek-R1	Jan 2025	90.8
GPT-4.5	Feb 2025	89.6
GPT-4.1	Apr 2025	90.2
Claude Opus 4	May 2025	88.8
Claude Opus 4.5	Nov 2025	90.8

OUTPUT TOKENS PER DOLLAR (TOKENS/\$)		
GPT-3 (davinci)	Jun 2020	16,667
GPT-3.5 Turbo	Mar 2023	500,000
GPT-4	Mar 2023	16,667
GPT-4 Turbo	Nov 2023	33,333
GPT-4o	May 2024	66,667
Claude 3.5 Sonnet	Jun 2024	66,667
Gemini 2.0 Flash	Dec 2024	2,500,000
DeepSeek-R1	Jan 2025	456,621
GPT-4.5	Feb 2025	6,667
Gemini 2.5 Pro	Mar 2025	100,000
GPT-4.1	Apr 2025	125,000
Claude Sonnet 4	May 2025	66,667
Claude Opus 4.5	Nov 2025	40,000
Claude Opus 4.6	Feb 2026	40,000

FRONTIER RELEASE INTERVALS (DAYS)		
GPT-3 → ChatGPT	Nov 2022	902
ChatGPT → GPT-4	Mar 2023	104
GPT-4 → Claude 2	Jul 2023	119
→ Claude 3 Opus	Mar 2024	89
→ o1	Sep 2024	84
→ Gemini 2.0	Dec 2024	90
→ DeepSeek-R1	Jan 2025	40
→ GPT-4.5	Feb 2025	38
→ Gemini 2.5 Pro	Mar 2025	26
→ GPT-4.1	Apr 2025	20
→ Claude Sonnet 4	May 2025	38
→ Claude Opus 4.5	Nov 2025	186
→ Claude Opus 4.6	Feb 2026	73

Each metric normalized to  $[0,1]$ . Release intervals inverted (shorter = better). Tokens per dollar log-transformed before normalizing (the raw values span **five orders of magnitude**; without the log, Gemini Flash at 2.5M tokens/\$ dominates the fit and everything else is noise). Each series keeps its own scale, no merging into a single ensemble.

### Why Hyperbolic

Most people extrapolate AI with exponentials. Wrong move!

An exponential  $f(t) = ae^{bt}$  approaches infinity only as  $t \rightarrow \infty$ . You'd be waiting forever. Literally.

We need a function that hits infinity at a **finite time**. That's the whole point of a singularity: a pole, a **vertical asymptote**, the **math breaking**:

$$x(t) = \frac{k}{t-t_s} + c$$

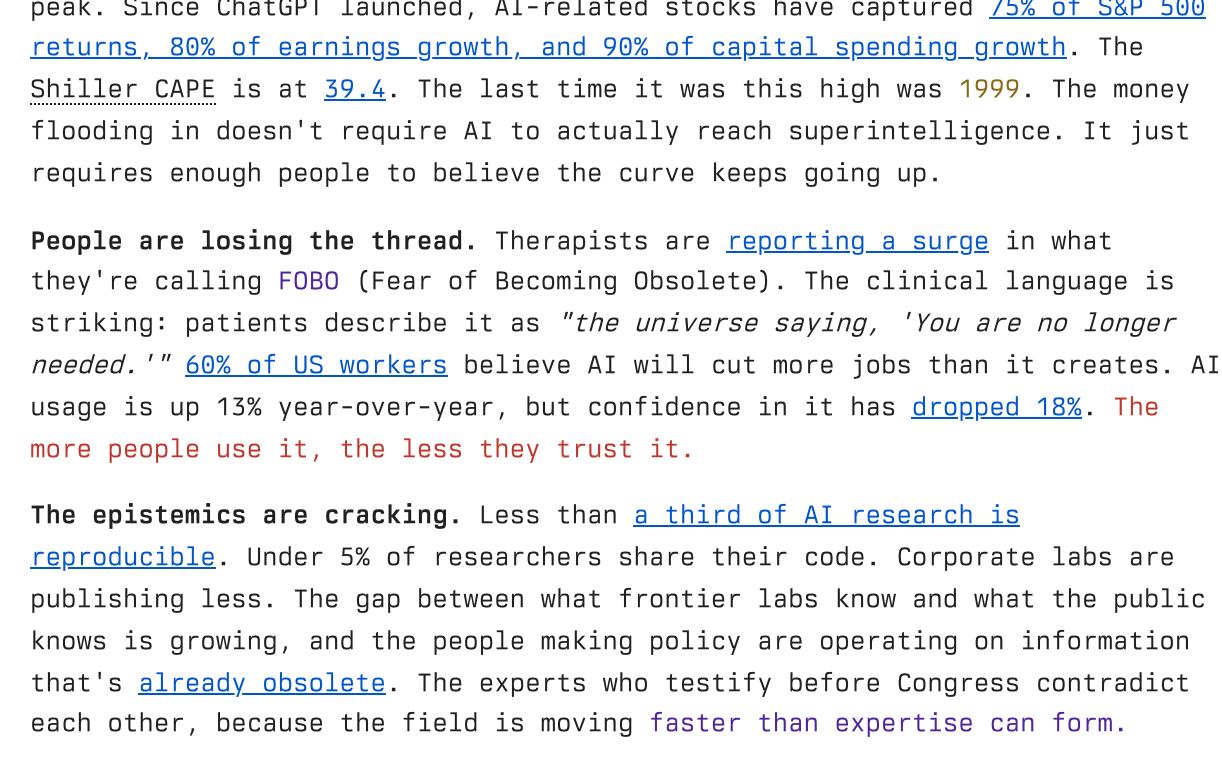
As  $t \rightarrow t_s$ , the denominator goes to zero.  $x(t) \rightarrow \infty$ . Not a bug. **The feature**.

$$x(t) \rightarrow \infty, x'(t) \rightarrow \infty, x''(t) \rightarrow \infty, x'''(t) \rightarrow \infty$$

- MMLU
- Tokens/\$
- Release gaps
- arXiv "emergent"
- Copilot code share

### - singularity

• 95% CI



52 data points across 5 metrics, fit to  $x_j(t) = k_j / (t_s - t) + c_j$

Polynomial growth ( $t^n$ ) never reaches infinity at finite time. You could wait until heat death and  $t^{47}$  would still be finite. Polynomials are for people who think AGI is "decades away."

Exponential growth reaches infinity at  $t = \infty$ . Technically a singularity, but an infinitely patient one. Moore's Law was exponential. We are no longer on Moore's Law.

Hyperbolic growth is what happens when the thing that's growing **accelerates its own growth**. Better AI  $\rightarrow$  better AI research tools  $\rightarrow$  better AI  $\rightarrow$  better tools. Positive feedback with **supralinear dynamics**. The singularity is real and finite.

### The Fit

The procedure is straightforward, which should concern you.

The model fits a separate hyperbola to each metric:

$$y_j^{(0)} = \frac{k_j}{t - t_s} + c_j$$

Each series gets its own scale  $k_j$  and offset  $c_j$ . The singularity time  $t_s$  is shared. MMLU scores and tokens-per-dollar have no business being on the same y-axis, but they can agree on when the pole is.

For each candidate  $t_s$ , the per-series fits are linear in  $k_j$  and  $c_j$ . The question is: which  $t_s$  makes the hyperbola fit best?

Don't actually have nobody. If you minimize total RSS across all series, metrics don't care if  $t_s$  is always at infinity. A distant hyperbola degenerates into a line, and lines fit noisy data just fine. The singularity degenerates into being whatever fit noisy data just found. You're finding the edge of your search grid, not a singularity.

So instead, we look for the real signal. For each series independently, grid search  $t_s$  and find the  $R^2$  value where the date where the curve hits zero is closest to  $t_s$ . If  $R^2$  peaks at  $t_s$ , then  $t_s$  is the real signal. If not, peak at some finite  $t_s$  and then decline. If it's really just linear,  $R^2$  will keep increasing as  $t_s \rightarrow \infty$ . Not peak, no signal, no vote!

One series peaks! arXiv "emergent" (the count of AI papers about emergence) has a clear, unambiguous  $R^2$  maximum. The other four are monotonically better fit by hyperbolic. The singularity date comes from the one metric that's actually better fit by hyperbolic.

This is more honest than forcing five metrics to average out to a date that none of them individually support.

Not the model. Deterministic. The stochasticity is in the universe, not the model.

The date is the moment when the curve keeps going up.

People are losing the thread. Therapists are reporting a surge in AI research. They're calling FOMO (Fear of Becoming Obsolete). The clinical language is needed: **90% of AI workers** believe AI will cut more jobs than it creates. AI usage is 13% year-over-year, but confidence in it has dropped 18%. The

AI people are less afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI research is up.

The AI people are across the board more afraid of AI than the rest of us. The AI

behaviors. The other four metrics are linear. The machines are improving steadily. **We are the ones accelerating!**

The social consequences of that acceleration (labor displacement, institutional failure, capital concentration, epistemic collapse, political realignment) are not predictions for 2034. They are descriptions of 2026. The singularity in the data is a singularity in human attention, and it is already exerting gravitational force on everything it touches.

I see no reason to let epistemological humility interfere with a perfectly good timer.

**See you on the other side!**

Want to know when I post?



Enter terminal mode