

Fenestrate Press presents

The AI Hater's Manifesto

A Monograph By

Ed Zitron

And Published

Aug 25, 2026

"When life closes a door, throw knowledge through a window."

Fenestrate Press prints items of interest, including both original submissions and reprints from elsewhere. More information, along with all our publications, is available at <https://fenestrate.net/>. Inquiries and submissions may be sent to info@fenestrate.net. Fenestrate Press publishes with the Ishill Stack, built on Hugo. Want to be able simultaneously to publish to the web, in PDF, in zine format, ...? Check out the free Ishill Stack at <https://codeberg.org/ussjoin/ishill>.

Document revision: 6d6e73c.

Originally published at <https://www.wheresyoured.at/the-ai-haters-manifesto/>. Ed is, without a doubt, the investigative journalist who has most thoroughly covered the tech industry's excesses in pursuit of "AI" in the last several years. Everyone (truly, everyone) should follow his RSS feed (or subscribe to his newsletter). Everyone able to do so should also kick in to support him; \$70/year gives you access to his "premium" pieces (which often go into even more depth).

This piece was published for free on his blog; we are providing printed copies to ensure that his work reaches as large an audience as practicable, and to provide support for those in arguments with AI sycophants that one can simply hand them. Typographical errors should be assumed to be made by us, not by Ed.

"Environmentalism without class struggle is just gardening." — Chico Mendes

The AI Hater's Manifesto

Ed Zitron

If you liked this piece, you should subscribe to my premium newsletter. It's *\$70 a year* [1], *\$18 a quarter* [2], or *\$7 a month* [3], and in return you get a weekly newsletter that's usually anywhere from 10,000 to 18,000 words, including vast, detailed analyses of *NVIDIA* [4], *Anthropic and OpenAI's finances* [5], and *the AI bubble writ large* [6].

My Hater's Guides To the *SaaSpocalypse* [7], *Private Credit* [8] and *Private Equity* [9] are essential to understanding our current financial system, and my guide to how *OpenAI Kills Oracle* [10] pairs nicely with my *Hater's Guide To Oracle, as well as the Hater's Guide To Oracle (Part 2)*. [11]

Subscribing to premium is both great value and makes it possible to write these large, deeply-researched free pieces every week. This week's premium will be The Hater's Guide To Circular Financing - and how the AI industry is increasingly turning into a scheme to funnel money to NVIDIA and Broadcom at any cost.

If you want to get in touch — and especially if you have any juicy information about Anthropic, OpenAI, or any other companies in the AI bubble — hit me up on Signal at ezitron.76. I'm also on IB on The Terminal.

I've been writing about AI for the best part of three years. I'll admit I was late, mostly because I was still trying to work out what it was I was doing with my life, let alone whatever it was I was "meant to cover" in a newsletter that started as a hobby on the side of another job I no longer really do.

Things have changed a lot since then, mostly in that I'm near 115,000 subscribers, the premium newsletter and podcast are now my business, and I've had to learn more about economics, technology, power, construction, and the deep cynicism that drives the modern tech industry than I ever thought possible. It's the greatest job in the world, and I'm very lucky to have it.

Today, I want to put in clear terms how I feel about AI writ large, and how detestable this industry has become.

Welcome to my Hater's Manifesto.

"When I give food to the poor, they call me a saint. When I ask why the poor have no food, they call me a Communist." — Fr. Hélder Pessoa Câmara

Modern Software Sucks, LLMs Are Interesting, And Their Cost Is Inexcusable

Want a great example of why everybody's pissed off at technology? I just tried to resize the above heading, and in doing so Google Docs for no apparent reason decided to make the entire paragraph below the size of a header. Modern software is inherently broken, a convoluted mess of different menus, tech debt, and poor design choices driven by *the Rot Economy* [12]'s growth-at-all-costs mindset which demands *constant change at all times*, none of which ever seems to manifest as a "better" or "smarter" product.

I think the vast majority of people want their software to work better, and one of AI's most frustrating lies is that it sells itself as "autonomous" as it continues the depressing trend of software that blames the user for its failure to meet their needs. Microsoft, Google, Meta and Amazon have made their products increasingly-convoluted, then attached a supposedly-magical tool to them that somehow makes them *more* convoluted.

You know what I'd love? Spell-check to work in Google Docs rather than putting a red squiggly line underneath and saying "yeah there's probably something wrong with this, I dunno what though." I'd like Microsoft Word to stop crashing because I have too many end-notes. I'd like Riverside to not have 10 different menus to click through to get to a link to send a person to join my podcast. I'd like my email to not be full of spam. I'd like things to "just work" rather than constantly fighting some sort of broken app or broken UX element or weird bug or intrusive pop-up about a feature that I don't want. I'd like Slack or Discord to not feel like digital escher paintings of different notifications.

LLMs are sold as some sort of magic tool that can fix "anything" without ever specifying what that thing might be, mostly because they cannot be trusted, even in things that they *mostly get right*, to do things right every time. While they can do "more" than they used to, the extent of that "more" comes with it the danger of giving a mindless software tool access to your computer's files, which it may *choose to delete* [13] *in pursuit of "efficiency,"* [14] which makes investigating what they *might be able to do* equal parts convoluted and dangerous.

One critique of my work is that I've never used LLMs. I have! I experiment with them from time to time to make sure I haven't missed something. I used one to debug a problem with my son's Minecraft add-on the other day, and it took 30 minutes of fucking around trying things to eventually sort of work it out. The other day I used one to install a Pokemon Minecraft mod, then when I asked it to

91. <https://www.adweek.com/media/openai-ad-business-is-on-pace-to-miss-its-own-forecast-by-90-analyst-says/>
92. <https://www.youtube.com/watch?v=TtmPccUTDP8>
93. <https://x.com/ns123abc/status/2090138676227453420?s=20>
94. <https://www.brookings.edu/articles/new-evidence-on-data-center-employment-effects/>
95. <https://www.axios.com/2025/05/28/ai-jobs-white-collar-unemployment-anthropic>
96. <https://www.computerworld.com/article/4059383/openai-admits-ai-hallucinations-are-mathematically-inevitable-not-just-engineering-flaws.html>

49. <https://www.wheresyoured.at/core-incompetency/>
50. <https://www.wheresyoured.at/the-enshittifinancial-crisis/#coreweave-is-still-a-time-bomb-by-the-way>
51. <https://www.wheresyoured.at/premium-the-haters-guide-to-nvidia-part-2/>
52. <https://www.wsj.com/tech/ai/nvidia-is-spending-6-billion-to-build-a-powerful-u-s-alternative-to-chinese-ai-c51c38cc>
53. <https://www.reuters.com/business/nvidia-invest-up-3-billion-lancium-information-reports-2026-08-08/>
54. <http://ounces-partnership-nvidia-advance-160100272.html>
55. <https://sharonai.com/press-releases/sharon-ai-reports-second-quarter-2026-results/>
56. <https://www.bcaresarch.com/reports/chief-economists-view-give-me-chastity-not-yet-19-08-2026/224466>
57. <https://uk.finance.yahoo.com/news/why-strategist-thinks-longer-term-123408535.html>
58. <http://investing.com>
59. <https://www.youtube.com/watch?v=n32mhb2aUsE>
60. <https://www.cnbc.com/video/2026/08/12/nvidias-ai-financing-initiative-is-just-vendor-financing-says-ed-zitron.html>
61. <https://www.youtube.com/watch?v=pHcZpvIfho0>
62. <https://www.cnbc.com/2026/04/20/amazon-invest-up-to-25-billion-in-anthropic-part-of-ai-infrastructure.html>
63. <https://www.bloomberg.com/news/articles/2026-04-24/google-plans-to-invest-up-to-40-billion-in-anthropic>
64. <https://www.axios.com/2026/08/23/greg-abbott-texas-data-centers-ai-backlash>
65. <https://www.theguardian.com/us-news/2026/aug/19/us-debt-40-trillion>
66. <http://v>
67. <https://x.com/edels0n/status/2087247141966368855>
68. <https://www.wheresyoured.at/premium-the-haters-guide-to-the-memory-crisis/>
69. <https://www.bloomberg.com/news/articles/2026-08-14/thrive-investor-letter-reveals-openai-fueled-growth-stake-sale>
70. https://youtu.be/wv4eCJUTZRY?si=tMJ_GHX-d7oIxosg&t=110
71. https://www.youtube.com/watch?v=KVdqwD_bcPs
72. <https://www.wheresyoured.at/exclusive-openai-financials/>
73. <https://www.theinformation.com/articles/openai-burned-3-7-billion-first-three-months-2026?rc=kz8jh3>
74. <https://www.wheresyoured.at/exclusive-openai-financials/>
75. <https://www.anthropic.com/news/anthropic-raises-30-billion-series-g-funding-380-billion-post-money-valuation>
76. https://youtu.be/n32mhb2aUsE?si=Jp9yiaooGf_pMiB3&t=384
77. <https://finance.yahoo.com/technology/ai/articles/hbm-prices-could-double-2027-131159931.html>
78. <https://www.multistate.us/insider/2026/6/22/state-data-center-policy-shifts-as-governors-impose-new-restrictions>
79. <https://www.bloomberg.com/news/articles/2026-07-07/amazon-returns-to-us-bond-market-to-fund-ai-infrastructure-build>
80. <http://gizmodo.com/sam-altman-admits-its-going-to-take-a-long-time-before-ai-goes-mainstream-2000802418>
81. <https://www.businessinsider.com/ai-bubble-heads-doomers-sam-altman-ai-costs-huge-issue-2026-6>
82. <https://www.wheresyoured.at/premium-how-much-money-does-ai-need/#google-microsoft-and-amazon-need-anthropic-and-openai-to-spend-at-least-444-billion-in-the-next-three-years-to-meet-analyst-expectations-with-anthropic-and-openai-representing-over-34-of-cloud-revenue-in-fiscal-year-2027>
83. <https://www.wheresyoured.at/premium-how-much-money-does-ai-need/#nvidia-needs-to-sell-148-trillion-in-gpus-and-data-center-gear-in-the-next-three-years>
84. <https://www.wheresyoured.at/premium-the-haters-guide-to-the-memory-crisis/>
85. <https://www.bloomberg.com/news/articles/2026-08-14/nvidia-s-500-billion-plan-envelops-wall-street-in-its-ai-frenzy>
86. <https://techcrunch.com/2026/05/30/what-a-joke-github-copilots-new-token-based-billing-spurs-consternation-among-devs/>
87. <https://techcrunch.com/2025/07/28/anthropic-unveils-new-rate-limits-to-curb-claude-code-power-users/>
88. https://x.com/edzitron/status/2090856861256036383?ref_src=twsrc%5Egoogle%7Ctwcamp%5Eserp%7Ctwgr%5Etwet
89. <https://www.wheresyoured.at/the-subprime-ai-crisis-is-here/>
90. <https://techcrunch.com/2025/07/28/anthropic-unveils-new-rate-limits-to-curb-claude-code-power-users/>

make sure the PS5 controller worked with the menus it broke a bunch of stuff, though I'll concede it was *useful* that it *installed* something *and it sort of worked*.

The fun part of that paragraph is there are some that will think this is a grand victory for their technology, even though the result is decidedly mediocre. Four years into the AI bubble, and the best you've got is that a tool *kind of worked* after I *bonked it on the head multiple times*, and all it cost was a trillion-plus dollars in capex and tens of billions of dollars of training compute. I would never, ever trust this thing that deleted and added lines of code at random with anything mission critical, I could not trust software built with it, and I certainly couldn't trust it with anything involving my personal data.

And with all that said, the only real "use case" i've found for AI in my life have been three or four times where I've dumped a crash log into one of the tools and said "why broken" and got a result. Am I meant to be impressed?

Sidenote: If your argument is "imagine what it could do in a year!" I just did so, and the answer was "the same thing, I guess?"

Here's how I feel about LLMs. In a vacuum, they're an interesting technology that can do some interesting stuff, in the right scenarios, but never in a way that involves you fully surrendering your actual work product to it.

As a way of speeding up small units of work in ways that are manageable both technically and cognitively, LLMs can be useful. The further you stretch yourself away from having complete clarity and industry over every element of the output's purpose, the more likely you are to fall foul to a technology that is mathematically certain to make mistakes, and if you feel insecure reading it, you know that you are, on some level, embarrassed to have used AI.

I don't tell everybody about *the weird keyboard* [15] I use, nor do I judge them despite how incredibly fast it makes typing for me, likely far faster than my competition, allowing me to operate at great speed. Who gives a fuck?

In any case, it is impossible to view LLMs in a vacuum, because their existence demands hundreds of billions of dollars. Every data center is incredibly expensive, offensive-sounding and looking, and their existence is explicitly to enrich some sort of Patagonia-gargoyle at an asset management firm, all sold under the auspices of "investing in American infrastructure," whatever the fuck *that* means. Their existence is a monument to the worst excesses of growth-at-all-costs capitalism — a technology that appears to coddle the user but ultimately lulls it into endlessly defending its fuckups under the flimsy pretense of "one day becoming perfect," though woe betide you if you ever set *perfection* as the target, because that's too unreasonable, as *humans* make mistakes.

Actually, that's a good point!

Stop Comparing LLMs To Humans Unless You Are Ready To Demand Trillions Of Investments In Social Services and Labor Rights

Please, point to the time in history when we have invested a trillion fucking dollars in making human workers better.

Point to a time when we have taken the idea that managerial culture is a performative fuck-fest built to enrich and empower business idiots that make important-sounding projects and con other people into doing the actual work.

Where is mentorship in corporate America? Where are labor standards? Where are the social services that would make human workers truly excel at their jobs — a good night's sleep, a healthy body, a good income, basic fucking dignity in the workplace, and their labor respected and empowered. I'm old enough to remember when everybody was chiding workers for "*quiet quitting*" [16] — by which I mean “doing the work you are asked to do and not taking on extra responsibility for free.” I've read article after article insisting that we do not need medicare for all, that Universal Basic Income is a bad idea, that we must means test welfare, that people must have a “good work ethic” and that ultimately someone's worth is derived from their contribution to the economy, hundreds of thousands of words dedicated to critiquing and prodding and judging every kind of worker other than the vaunted Chief Executive Officer or the Glorious Startup Boys.

Everyone seems so obsessed with sinking billions of dollars into the theoretical chance that machine learning might be able to replace human beings, and that more money makes it “smarter” and “better” at tasks, but the idea of unionization, healthcare as a right, investing in the education, and actual talents of the workers would be *communism*.

Yet for some reason — because it's a product, I guess? — we should as a nation, society and media ecosystem should do everything we can to assure that as much money as possible is invested in *fucking large language models so that they can become something they are not*.

There is no AGI coming. There is no conscious computer. LLMs have gotten “better,” but the “better” is not the kind of “better” that actually makes “economic sense for literally anyone involved.” Your best case scenario is that these things can do some coding work for you, in a controlled manner, in a way that's safe, or alternatively face the professional harm that's already befalling basically anyone getting caught using LLMs outside of coding, and even then, those within software

Footnotes:

1. <https://edzitrnswheresyouredatghostio.outpost.pub/public/promo-subscription/w1KX6aWthd?ref=wheresyoured.at>
2. <https://edzitrnswheresyouredatghostio.outpost.pub/public/promo-subscription/pJMqFLQQw0?ref=wheresyoured.at>
3. <https://edzitrnswheresyouredatghostio.outpost.pub/public/promo-subscription/iT6Np52ZmU?ref=wheresyoured.at>
4. <https://www.wheresyoured.at/the-haters-guide-to-nvidia/>
5. <https://www.wheresyoured.at/howmuchmoney/>
6. <https://www.wheresyoured.at/premium-the-haters-guide-to-the-ai-bubble-vol-2/>
7. <https://www.wheresyoured.at/hatersguide-saas/>
8. <https://www.wheresyoured.at/hatersguide-privatecredit/>
9. <https://www.wheresyoured.at/hatersguide-pe/>
10. <https://www.wheresyoured.at/how-openai-kills-oracle/>
11. <https://www.wheresyoured.at/haters-guide-oracle/>
12. <https://www.wheresyoured.at/the-rot-economy/>
13. <https://futurism.com/artificial-intelligence/claude-wife-photos>
14. <https://www.pcmag.com/news/meta-security-researchers-openclaw-ai-agent-accidentally-deleted-her-emails>
15. <https://www.zsa.io/moonlander>
16. <https://www.npr.org/2022/08/19/1117753535/quiet-quitting-work-tiktok>
17. <https://www.wheresyoured.at/the-revenge-of-the-business-idiot/>
18. <https://www.robinsloan.com/lab/jarvising/>
19. <https://ludic.mataroa.blog/blog/ai-mania-is-eviscerating-global-decision-making/>
20. <https://ludic.mataroa.blog/blog/i-will-fucking-pile-drive-you-if-you-mention-ai-again/>
21. <https://finance.yahoo.com/technology/ai/articles/microsoft-ai-sales-mostly-come-180651710.html>
22. <https://www.theguardian.com/technology/2026/aug/23/ai-debt-bomb-crisis>
23. <https://www.wheresyoured.at/dot-com-bubble/#data-centers-and-gpus-are-nothing-like-fiber-optic-cable-and-a-gpu-crash-will-be-utterly-wasteful>
24. <https://www.pase.com/projects/facebook-prineville/>
25. <https://www.digitalrealty.com/resources/blog/revisiting-the-worlds-largest-data-center-350-e-cermak-chi1>
26. <https://www.datacenterdynamics.com/en/news/ubs-acquires-first-dedicated-data-center-in-london/>
27. <https://dchub.cloud/facilities/ubs-ubs-data-centre-c1e0c5d6>
28. <https://www.datacenterdynamics.com/en/news/former-rockland-psychiatric-center-to-become-jpmorgan-chase-data-center/>
29. <https://www.datacenter.fyi/public-record/01af289e>
30. <https://www.cnbc.com/2026/08/17/anthropic-says-annualized-revenue-climbed-to-65-billion-in-july.html>
31. <https://www.wheresyoured.at/the-subprime-ai-crisis-is-here/#what-does-%E2%80%9Csubsidized-ai%E2%80%9D-mean>
32. <https://techcrunch.com/2026/05/30/what-a-joke-github-copilots-new-token-based-billing-spurs-consternation-among-devs/>
33. <https://x.com/edzitron/status/2055775345174090206?lang=en>
34. <https://www.wheresyoured.at/ai-doesnt-have-roi/>
35. <https://ludic.mataroa.blog/blog/ai-mania-is-eviscerating-global-decision-making/#:~:text=II.%20Heretics%20Will%20Be%20Shot>
36. <https://www.ft.com/content/5ee49718-c258-4f01-aa32-7e5b76ae5245?syn-25a6b1a6=1>
37. <https://aireiter.com/blog/gpt-5-6-sol-token-usage-cost-guide>
38. <https://www.cnbc.com/2026/07/30/open-ai-price-cut-gpt.html>
39. <https://www.reuters.com/technology/openai-cuts-developer-pricing-frontier-gpt-56-sol-model-by-more-than-20-2026-08-21/>
40. <https://www.wsj.com/cio-journal/why-at-t-is-betting-big-on-open-weight-ai-a0ea03b1>
41. <https://openai.com/index/klarna/>
42. <https://archive.ph/RU6a8>
43. <https://www.youtube.com/watch?v=S5-t8WSQafM>
44. <https://www.theinformation.com/articles/nvidia-ai-chip-prices-rise-17-server-makers-tell-customers?rc=kz8jh3>
45. <https://archive.ph/ok3HP>
46. <https://www.reuters.com/technology/nebius-plans-45-billion-convertible-debt-sale-fund-data-centers-ai-platform-2026-08-19/>
47. <https://www.businessinsider.com/stock-market-signal-chip-stocks-ai-capex-mag7-jpmorgan-2026-7>
48. <https://www.wheresyoured.at/how-openai-kills-oracle/>

[96], and when you say something about it, its acolytes will lie and say that “hallucinations are solved.”

It also can only ever sell itself based on what might happen and the theoretical promise of you giving it your complete attention, connecting every bit of data you own, paying whatever it costs, and accepting that it can and will change in price and functionality at random, all while never putting a precise timeline on whatever AGI means that particular week.

Whenever you ask for clarity, the AI industry gives you chaff. Whenever you ask when things get better, you’re told it’s both *the early days* and that AI is *the worst it’ll ever be*. Even the term “artificial intelligence” is a bad faith attempt to conflate transformer models with things like robotics or autonomous cars, all so that its proponents can claim other people’s successes as their own despite LLMs having little or no relevance to anything else *other* than generative AI.

It encourages dogpiling and ostracizing those who don’t fall behind it, because it cannot succeed on its own merits. It encourages a vile cultism powered too by bad faith and parasocial relationships with both AI CEOs and the models themselves. It exploits the intellectual weaknesses of “smart people” that are actually just good at remembering the right things to say at the right time and have memorized the various justifications for past failures, all while allowing them to use LLMs to promote their own bad faith enterprises where they use work-adjacent product to con others into paying them.

And it’s losing because, at its core, AI was never *built* on very much. It grew this large because the media manufactured consent at the behest of the powerful because lots of money got invested, and the rich and powerful can never be *wrong*. The underlying technology may be *more useful than it was*, but it’s *not useful enough to be profitable* nor *reliable enough to be world-changing*, and the bad faith representation of LLMs as “good enough” should be a permanent scarlet letter on anyone who misled the public into believing this was anything other than normal software.

I was asked recently why I find this all so repugnant, and my answer is simple: I don’t like bullies, I don’t like con artists, and I don’t like being lied to. This industry grew by misleading people about the actual and potential outcomes from Large Language Models, and through an economy-wide attempt to pressure everybody into adopting tools in pursuit of growth at all costs.

Ultimately, it was sold with the greatest lie of all: “this time it’s different!”

To be clear, they’re right.

It’s so much weirder, and in the end will be so much worse.

engineering who are over-LLM’d are mocked. It’s also becoming increasingly more-difficult to understand both what has made an LLM “better” for both the people using them and the people making them, and there has been little-to-no headway made in making a meaningful impact in other industries.

You can jerk your bingus all you want about benchmarks or case studies or some anecdote you heard on a Subreddit, but AI products are just not very good at *stuff*. Those who boast of “massive productivity gains” from AI have found them only after endless hours of tinkering (or “Jarvising” as I’ll get to later), and in every single case their work reads or looks like crap, unless of course they’re somebody *using LLMs as tools rather than a replacement for their miserable little mind*.

LLMs *can* help out with lots of small things, get worse as they try and do real things, and do not need to speak like people. They do not need to be in anything near healthcare or finance or mental health or, really, people. The anthropomorphism and overpromising about these technologies has suffocated and obfuscated what they can actually do in pursuit of endless growth, and the only reason they can do anything is that OpenAI and Anthropic were allowed to annihilate hundreds of billions of dollars on training, along with very real harms and systemic risks that have emerged as a result.

Sidenote: The “well human beings make mistakes too” argument is very stupid on its own — after all, human beings can learn on the job at speed, and can self-correct in a way that LLMs are incapable of doing.

I’ll also add that the way that people frame how “often” LLMs make mistakes is utterly flawed too. A human being might make a *mistake* but largely get the facts and techniques correct while meaningfully understanding the purpose and developing their approach over time. An LLM can keep a text document or look at files and data and, each time, and then generate what it believes is the right course of action based on *training data* rather than *experience*.

I don’t even know why I’m explaining this at this point, because those making this argument are not approaching the conversation in good faith and are really just looking for a new boot to lick.

If you think any of this is worth hundreds of billions or trillions of dollars, you are either ignorant or corrupt. On top of how disgusting their outputs feel, the cost is going to take at least a decade to share, and begin the end of hypergrowth in the tech industry.

And it’s a fundamentally ridiculous argument to compare LLM outputs to human beings without giving human beings the same affordance, grace and *sheer investment* as a comparison.

Where is the grace for human error? Where is the investment in making humans exceptional? Surely investing real money in actual workers — making their lives better, improving their working conditions, teaching them new things, sharpening their existing skills, rewarding them for their hard work, and so on — would have better effects than fastballing hundreds of billions of dollars into a machine that does an impression of work?

Unless, of course, the people demanding this don't do any actual work!

LLMs Are Built To Help Grifters and Imbeciles Do Impressions Of Work

I'll concede we're past the point when "nobody uses these things," as they have now been pushed non-consensually upon every worker and organization at scale predominantly by *Business Idiots* [17] that demand workers "do enough AI" because saying "I do AI" is a virtue signal to a certain kind of scumbag.

One of the many dangerous things that an LLM can do is a messy impression of a competent person, filling in the little bits within a loser, moron or con artist that would've otherwise exposed them, allowing them to get deeper and deeper into organizations by creating make-work specifically built to get off the MBA sect, resembling the *performance* of work because much of the workplace is ruled by people that don't do any and haven't in years. You can immediately read when somebody has used it because the words don't sound right and don't convey proper meaning.

It is genuinely hard to read anything more than puddle-deep written by AI, because the more complex a subject is, the more skilled a writer must be to convey its meaning, and the more work it must do to pull people into concepts. The odd emotional swings in AI writing are its true tell — everything is *extremely serious and urgent* or told in a disinterested monotone, with no attachment to the words or why they were put in the order they were. People read my stuff because I convey facts and feelings but my work resonates with emotion. Some AI boosters frame this as me "just swearing" or "riling people up," but that's because they're not used to caring about stuff for anything other than professional reasons.

Everything you see is the result of elevating people who value and build things based on growth. LLMs offer so many promises to those who don't want to build anything of value — a way to seem like you're "investing in American infrastructure," a way to be sinophobic, a way to crush workers, a way to pretend like you care about the future, a way to pretend you care about technology, a way to talk about vacuous pseudo-intellectuals as a means of seeming intellectual yourself, an endless font of new multi-million or multi-billion deals and personnel changes,

Nothing about LLMs is worth a trillion dollars, or even \$100 billion. This is, *as I've said before* [92], a \$30 billion TAM industry dressed up as a trillion dollar one, and the only reason it's grown this large is because the two leading companies have had their infrastructure built for them and given unlimited resources to subsidize their customers' compute.

And what's really stood out is how *so little about the AI bubble is actually about AI*. No other technology in history has had professional and social consequences for failing to use or like it enough, nor can I find any example in history where journalists have actively attacked critics for not being sufficiently approving of a kind of cloud software. It is fundamentally crazy to me that, in pursuit of "objectivity," much of the tech and business media has chosen to accept whatever narrative the AI industry gave them, assuming that whatever we have today is *already guaranteed* to be something better in the future, both in its outcomes and profitability.

This era is unlike any other before it, but took advantage of the fact that most people are desperate to apply the past to the present to rationalize or process what may seem irrational or destructive. To see AI as "just like the dot com bubble" allows you to ignore both the costs and the potential outcomes because "things worked out after that," even if there're basically no uses for GPUs after this and the only way we "build new LLMs" is by feeding them expensive training data using billions of dollars of compute that are only available while everybody still believes this is real.

The AI industry — and the AI bubble — is fundamentally built on acting in bad faith. Its executives lie. Its boosters lie. Its software lies because it doesn't actually know anything and generates answers probabilistically, and if you mention that online, someone will harass you for doing so.

It refuses to answer straightforward questions. It refuses to present a plan for the future. It refuses to explain how it becomes profitable, because nobody knows how or has a tangible plan to do so. It deliberately subsidized its subscription products because it knew its customers wouldn't pay the actual cost of AI, and tortures customers with shifts in functionality and rate limits all while framing this as a way to *"continue to serve customers the most cost-efficient models* [93]."

It attempts to conflate massive, power and resource-hungry AI data centers with the smaller ones that bring helpful yet increasingly-decaying software to our homes. It sells these data centers as "bringing jobs to communities," all while importing the talent from out of state to build the things then leaving *a crew of 100 to 200 people* [94] to actually run them after millions or billions of dollars of tax breaks. It sells its "innovations" as creating a *"white collar bloodbath* [95]" to scare you into using inconsistent and unreliable software that's *mathematically certain to make mistakes*

The AI industry demands you constantly live in the future tense. Everything is about tomorrow's billions or trillions, the potential of what you're seeing rather than the thing itself, future gigawatts in data centers that you must treat as if they are already built and value based on things that AI might theoretically do. I challenge you to read everything about AI from this point forward with this in your mind so you can see how intently this industry tries to drag your focus away from what it's doing toward what it might theoretically do if it only had more money, power and resources, and ask yourself why they need to do so.

To be clear, they're doing so because you can't really justify anything about this industry based on what it does today. It costs too much, none of the businesses built on top of it are profitable, it costs so much to build a data center that the most cash-rich asset-light businesses in the world are now burdened with endless expensive-to-install and run hardware for a business that makes a fraction of its overall costs in revenue and has little demand outside of two companies that everybody must conspire to keep alive both financially and philosophically.

And ultimately, nobody can actually explain why we need more data centers.

Would anything really change? What would change? How? How many more do we need? Why do we need so many? Having more power plants meant more people could have power, and having more fiber laid meant connecting more buildings to the internet. What does one more or two more or ten more data centers actually give you? Is there some part of the world unable to access or take advantage of the LLMs available on seemingly every surface of the internet? Because it seems like the only reason these things are getting built is to capture illusory demand based on a "supply constraint" created by two unprofitable companies absorbing all the infrastructure. I don't hear any compelling scientific or technological reason building more is useful or productive outside of funneling more cash to semiconductor companies.

Seriously, go and read basically any article about AI and see how quickly they start talking about the future, be it in the mainstream media or on a startup's blog. Every single piece must sell AI on its theoretical promise and, if at all critical, reassure you that the author of course doesn't dispute the "transformative potential of AI" or "how it's already transforming the economy," even if it can't define how it's doing so or even what that means.

The AI Industry Runs On Bad Faith

I let myself have a little fun with today's piece because I feel like I've been so deep in the financial trenches that I forgot how much of the AI industry runs on propaganda, social pressure and outright bullying to manufacture consent for a product that demands everything and provides very little in return.

a new power center to graft oneself onto, a new asset class to invest in based entirely on vibes, and a way to be mildly jingoistic, all wrapped in a tool that can give you enough facts to pretend you know anything safe in the knowledge that most people are trained to believe somebody who *sounds smart*.

You're Embarrassed That People Know You Use AI Because You Know It's Shameful To Outsource Your Thinking

It just came to me — the problem that I have with most people using LLMs is the delineation between outsourcing *work* and outsourcing *thought*. Those using LLMs to write little scripts or BQL code on a Bloomberg Terminal are inoffensive. A person using an LLM to search a big document for something is unproblematic, assuming that we ever fix the overall environmental footprint. A user reorganizing their desktop, assuming it works, is not an issue.

A tool being used as a tool to do tool things — in many cases involving the LLM writing a little 30-line Python script! — is not a problem, though it's also not a trillion-dollar industry that needed to steal everybody's art and writing.

The problems begin when somebody outsources their thinking and actual work, and yes, this includes "research." AI research fucking stinks, as does AI writing. AI-authored code — especially vibe-coded programs — is inherently dangerous and disrespectful to the user, and I believe endless AI-generated code is behind the overall deterioration of software at large.

AI writing is also disrespectful to the user, because you didn't actually come to any conclusion other than saying "uh, yeah, what that says." You did not have a thought, you did not have a feeling, you did not make a statement, you prompted a model and fooled yourself into thinking that feeding your own words into it via data dumps or natural language is the same thing.

The reason you feel embarrassed to tell people you use AI is not because of a "misinformation campaign," but because you know what you're doing!

You know that you're relying on something that is mathematically guaranteed to be inconsistent. You know image generation is fucking ugly. You know the text sucks. There is a very obvious line where using LLMs goes from useful to lazy, it's extremely bold, and it's the moment you sacrifice a meaningful level of responsibility to them by not understanding the underlying operation.

That can mean everything from the underlying functionality of an app to writing the body of a piece of text you edit ultimately comes down to how much you give a shit about your audience or value your work. If your work is not better than an LLM's, you're bad at your job. I don't care if you used it to generate a chart or pull some data, as long as you *check every single god damn number*. If you're writing an entire article using an LLM and then editing it, even if you pulled the data yourself, I will never have much respect for your work, mostly because I have no real idea what you think as you didn't feel the need to tell me, you got some fucking word generator to do it.

LLMs Are Digital Busyboxes, And Are Great At Making You Feel Smart Through Computer Science

LLMs are also really, really good at what Robin Sloan calls “*Jarvising* [18],” creating a seemingly-autonomous assistant that mostly serves the function of giving you reasons to work on it:

However, the most common application of a personal Jarvis seems to be ... tinkering with one's personal Jarvis. “Gotta get my tools just right” isn't a new phenomenon, of course, but/and it's useful to notice its recurrence here.

LLMs are really good at creating the sense that you're being really, really productive. Evaluate this, generate that, investigate this, summarize that, tell me how many times something happened, give me a new number to obsess over or the sum of the parts of everything I've ever done, all so that I can *know more about my own thoughts without thinking*. One can obsessively catalogue and digitize every link and thought and musing and action and datapoint in their lives and theorize that the LLM can make them *better* by *knowing more about them*, a Tower of Babel built using AI compute, because it's *so easy* to make yourself feel smart by calling something a *database* that you *store stuff in* and run *analyses on*.

Best of all, the work is never done, and anyone you describe it to thinks you're doing *computer science* as you click buttons on Chrome plugins and justify paying Sam Altman \$200 a month. Don't worry though, model instructions involve the phrase “you are a genius data scientist and ruthless analyst,” which is functionally the same thing as remembering, reading, re-reading and synthesizing information using your brain if you're a person that doesn't really give a shit about doing a good job or being exceptional in any way.

And the fact they haven't worked out their business models means that whatever you're paying today is guaranteed to change. What other product do you regularly buy that has such chaos built into it? What other thing do you pay for where the prices (or availability) can shift to the point that you literally can't use it in the same way at a moment's notice? And why does anybody tolerate it when it comes to AI?

I'll add that this is a specific situation where the tech media has categorically failed the customer. We have companies valued at hundreds of billions of dollars that are fucking their customers over day-in-day-out, and the response is mostly to say “*huh that's strange* [90]” and refuse to let a single critical thought cross their minds.

The AI Bubble Requires Everybody To Live In The Future Tense, Because The Present In No Way Justifies An Iota Of Its Costs

Every part of the AI bubble must exist in a constant state of flux so that there can always be a future breakthrough that's always just out of reach. AI does not have to reach an actual achievement — it just has to “show promise” in some way. It is an objective disaster that Microsoft spent more than \$260 billion on capex to create a business with less than \$11 billion in annual revenue outside of OpenAI, but people will see “\$34.33 billion in annual AI revenue” and say “that's promising growth, up 123% year-over-year!”

They'll hear about LLMs that delete people's databases and say “well the models have gotten exponentially better,” even if that better part never seems to eliminate these issues, make a profitable AI company, or create a true killer app that you can point at beyond saying “ChatGPT has one billion weekly active users,” despite around 95% of them not paying a penny (and costing OpenAI likely billions of dollars) and *eMarketer estimating that the entire global AI chatbot advertising industry will make \$5.41 billion revenue in 2030* [91], giving OpenAI little hope of stemming the burn. These big numbers — like Anthropic having a \$65 billion annualized run rate, an undefined term that obfuscates the fact that Anthropic has made \$16.5 billion in the first half of 2026, losing billions of dollars in the process — are fundamentally meaningless, because they're easily gamed at best, and inherently uncertain at worst.

Because nothing works yet as a real business model or anything approaching truly autonomous (or “magical”) software, there’s the implicit knowledge that you’re going to have to change your product again and again to update to the “best model” or “make things more efficient” (IE: lose less money) or when something breaks because a model’s training got tweaked.

The euphemism for this is “exponential improvement,” when it’s really an Arnold Palmer of instability and novelty, and abuses basically anyone connected to the ecosystem every single day.

If there’s always something new happening, it’s hard to pin down if things have gotten better, or whether you’re just more proficient in cobbling together different harnesses, prompts and API calls to make it do what you need it to. It is undignified that people tolerate models that become either dumber over time or at random opportunities, while also being deeply exhausting for the end user.

As a paying user of an LLM-powered service, you are guaranteed at some point to face a degradation in service where models misbehave, some sort of shift in rate limits, or some sort of change in product functionality based on their shifting economics.

Has there ever been a bigger shift in a business product’s value than GitHub Copilot’s shift to token-based billing? *Microsoft rug pulled two million people that had built workflows on a platform that was allowing them to burn \$1,000 to \$5,000 in tokens for \$20 a month* [86]. That’s genuinely crazy! It’s magnitudes more than when Uber jacked up its prices.

It’s equally-insane that Anthropic and OpenAI similarly *fuck* [87] with their *customers* [88], changing the amount of value you get for \$20, \$100, or \$200 a month at random in a way that shouldn’t be legal.

Basically any AI-powered software is subject to arbitrary shifts in availability, capability and pricing at the whims of the vendor. *As I covered in my Subprime AI Crisis piece earlier in the year* [89], Replit, Perplexity, and multiple other AI companies have sold their customers a lie by pushing an unprofitable product that they must constantly “tweak” to bring down costs, all while misleading the customer about a “price” that continually declines in value as the price stays the same.

This is not a sustainable industry — either economically or emotionally — because it has a fundamentally dishonest relationship with its customers defined by the inconsistency of LLMs both in efficacy, stability (see: Anthropic’s downtime) and training, with each model randomly better or worse at things to the point that it must be a legitimate nightmare to run any software or build any product on top of them.

The people that actually use these things and like them in a normal way do not feel offended when they read this stuff because they see LLMs as a kind of software, and don’t feel a great emotional attachment to it because they’re not a weird freak.

They do not have obsessive involvement in “the AI debate” and almost always find the financial aspects truly loathsome. Said debate makes it near- impossible to actually judge how useful LLMs are to the software engineering industry because of the sheer scale of industry capture, but Nik Suresh is the literal best person doing the work on this, as described in *AI Is Eviscerating Global Decisionmaking* [19]:

All of the AI projects we have observed as a team are failing. Every single one — we have seen 0% success in a year and a half, not only amongst projects we have been asked to participate in, but even within projects that we have observed in passing while doing totally unrelated work. Even if you grant that AI tooling accelerates specific workloads, the method and scale of the current investments is senseless. Frequently the failure is not related to AI itself, but rather that companies are terminally bad at running software projects effectively, and as *I have remarked previously* [20], AI projects are subject to all the failure modes of normal projects plus you can get everything right and then still fail because of the method’s novelty. Very few companies are so good at shipping software that they can afford the extra risk profile.

Nik is a well-respected software engineer and a very successful consultant and businessman. He has reached this level by being good at both software engineering and running a company in a way that treats his customers, workers, and the work product itself with respect. The reason that I respect him so much, other than him being a great human being, is because he describes the successes he has with his clients with pride and loves making money by being good at his job and making his customers happy.

I have never seen somebody like Nik who is also a huge, drooling fan of AI. In fact, the people most-excited about AI tend to, at best, create distinctly mediocre shit.

Generative AI Is A Death Cult Of Growth-At-All-Costs Excess

The perniciousness of generative AI is a result of executive incompetence mixing with a technology built to, as discussed, create endless growth. Generative AI is far more useful as an *idea* than as a *technology*, and only ever has to show enough *promise* to back whatever vile agenda you’re pursuing.

With AI, you can do more, be more, sell more shit.

With AI, you can add AI to your service, whatever that means.

With AI, you can invest in AI stocks, or data center bonds, or power company stocks, or semiconductor stocks, and you can talk about these stocks like they're your sports team or lover or best friend, and sometimes the CEO will reply to your post and you can talk about "all the alpha" you just got.

With AI, you can back a new movement so that you can feel part of something. You can learn all sorts of new names and technical terms and subscribe to 90 newsletters from "industry insiders." All of that "alpha" can disprove just about anything, or deflect annoying *truths* like how Microsoft only *made a whole \$34.33 billion in annual revenue for the apex predator of modern software* [21] and all it cost was over \$260 billion in capex and \$13 billion in equity investments.

You see, as one of the *chosen*, you don't need to *worry about all of that* if you can talk about high-bandwidth memory or KV Cache or optical cable enough to cobble together sufficient smart-sounding terms to make it seem that you have an *intellectual* reason to ignore the obvious unprofitability, overbuild, overstatements of capabilities and impossible economics of the movement you're backing, and there're 4,000 Twitter weirdos ready and waiting to huff paint beside you.

By joining the great AI death cult, you too can live in a bubble, all while screaming slurs at people who dare to bring reality to your doorstep. All that matters is that *number go up*, and that *you are the person who said number would go up*, and when *bad numbers appear* you have *enough groupthink and alpha* to *scream at the people who brought the bad numbers up*.

It is insane how people talk about AI online. For all the whining I've read recently about how "Anti-AI people got the data center data wrong," I read thousands more words a week of some person who has done hours of research to put together a deeply technical report that does literally everything it can to *ignore reality*. I listen to podcasts and watch TV segments and read articles that simply will not address the obvious economic realities, and have built vast bulwarks of mythology to defend themselves. How many fucking times do I have to hear *someone say that data centers are just like the dot com bubble and everything will be fine after* [22] even if that's *completely untrue if you spend even a second thinking about it* [23]?

Sidenote: and *fuck you* if you're one of the cretins or imbeciles trying to say "oh, you don't like data centers? What about online banking?". Data centers for AI are anywhere from 10 to 100 times larger and more power-intensive than those used for things like social media or streaming. For example, *one of Meta's largest pre-AI data centers in Pineville Oregon has a power capacity of 30MW* [24], and *Digital Realty's 100MW Cermak Illinois data center handles hundreds of different industries and customers* [25], when the smallest AI data center announcement I've seen in the last year was for 100MW, with most in the 300MW to 1.2GW range.

the tune of *more than \$440 billion across Google, Microsoft and Amazon alone* [82].

- Hyperscalers must continue to buy NVIDIA GPUs, as the moment they stop doing so, the markets will begin to ask whether AI is an actual growth market anymore and ask for real, tangible answers about where all of this capex is going.
 - To be specific, *analysts expect NVIDIA to make \$1.48 trillion in revenue across Fiscal Years 2027, 2028 and 2029* [83].
- NVIDIA must sign long-term agreements to buy high-bandwidth memory at scale from SK Hynix, Micron and Samsung — who make 90% of all DRAM — or know its costs would spiral out of control, by which I mean its margins would compress at random at a time when it's already having to spike demand in extremely odd ways.
 - *Read my Hater's Guide To The Memory Crisis* [84] for more.

As I've said, NVIDIA's price increase is going to increase the price of every single data center in construction by billions of dollars, and we're already approaching the limits of how much money can be raised for them. That "\$500 billion" announcement was actually Jensen Huang jumping the gun, *per Bloomberg* [85]:

Goldman Sachs Group Inc., Blackstone Inc. and Apollo Global Management Inc. had been working tirelessly for months to draw up debt deals that would help developers of artificial intelligence systems pay for chips from Nvidia Corp.

With slow progress on the complex deals, Nvidia's chief executive officer, Jensen Huang, decided to change tack: He went public this week with the effort, saying the group is aiming to collectively finance AI computing deals totaling \$500 billion — a round figure with no obvious provenance.

The largest asset managers and financial institutions were making "slow progress," and that was before Jensen Huang increased prices by 15%. Do you think it'll become easier from here? How would that happen, exactly?

God, I'm tired.

AI Is Exhausting Everything and Everyone It Touches

The entire AI bubble has been exhausting for everybody involved.

The Great Exhaustion of the Rot Economy

Here's a very simple question: how much longer can everybody afford to keep doing this?

Every single thing has become more expensive in the last year. Even though token prices have gone down or stayed flat, the amount of tokens you burn has clearly increased to the point that organizations are apparently spending billions of dollars on AI services with difficult-to-quantify ROI, requiring frantic advocacy to and financial debasement with every turn of the wheel. OpenAI and Anthropic have become more expensive to run, and OpenAI's non-GAAP operating margin increased from negative 122% to negative 183% in Q2 2026.

NVIDIA's GPUs just became 15% to 17% more expensive because *high bandwidth memory costs doubled* [77], a conga line of different monopolies upping their prices assuming that each link in the chain will keep spending, as each one of them — down to the AI labs themselves — knows that its contribution to spending on AI is an existential rite.

This means that any data center with GPUs delivered in 2027 and beyond will now have to cover billions of dollars' worth of extra costs, on top of increasingly-staunch local authorities requiring power guarantees (*[\$100 million a year in Wisconsin for Oracle]*(<https://www.theregister.com/ai-and-ml/2026/07/21/oracle-faces-100m-annual-bill-to-back-wisconsin-datacenter-power-promises/5275728>)) and *states like Illinois, Arizona and Virginia killing their tax breaks* [78], all as interest rates spike and *demand for AI debt weakens* [79].

Every single year, every single part of the AI bubble becomes more expensive — AI labs want to spend more money, AI data centers cost more money, AI services become more expensive, AI debt becomes more expensive, and everybody becomes decidedly less-patient for there to be some sort of outcome.

Meanwhile, public relations expert and OpenAI CEO *Sam Altman told podcaster David Senra* [80] that “we’ve all [referring to the AI industry] been too ambitious on timelines...[and that changing people’s behavior] is much harder than the tech nerds realize.”

Sam: stop talking! *Every time you open your mouth you say something silly* [81]!

Anyway, here's everything that needs to happen in the next three-and-a-half years:

- OpenAI and Anthropic must keep spending as a means of justifying their existence to hyperscalers using their revenues to artificially inflate growth, to

By contrast, let's look at some bank data centers. UBS *bought one in Hayes, West London in 2015 which it had previously rented* [26]. The cost? The princely sum of £28m, or \$42.8m at the time's exchange rates. This had *a power capacity of 5MW, which assuming a very generous 1.3 PUE (power usage effectiveness)* [27], means that it had around 3.8MW of critical IT.

UBS is one of the largest banks in the world, and given the importance of the City of London to the world financial system, it's reasonable to assume this data center is operationally important to the company.

Even when banks invest in huge facilities, they're still far smaller than the smallest AI data centers. Take, for example, *JPMorgan Chase's data center in Orangetown, New York, which sits on the former site of the Rockland Psychiatric Center* [28]. This has *a power capacity of 45.7MW, and a critical IT load of 27.4MW* [29] (giving it a PUE of 1.666).

JPMorgan Chase is both the largest bank in the US, and the largest bank in the world.

Oh, and **AI data centers are literally only good for AI, AI GPUs do not have other mass-market use cases. There is no post-Dot Com story. *Fucking look, I'm sick of repeating myself!***

Look, I'm sorry, Anthropic is not worth \$2 trillion, and whatever convinced you of that is a mixture of manufactured consent and mistaken trust of the powerful. The fact any of you take “*annualized run rate* [30]” seriously is an offense to good sense, and yes, that includes every reporter reporting it, even the ones I respect.

It's also ridiculous that anyone is talking about “recursive self-improvement.” The AI industry has become so utterly lazy and coddled that it's just saying “uhhh, AI will train itself I guess.”

And man, is it ridiculous that AI doomers warning about spooky superintelligences have somehow had such incredible prominence in the media without ever succeeding in stopping a single thing — or even substantiating their concerns.

“Dangerous AI” Is Already In The “Wrong Hands” — Anthropic, OpenAI, and Meta

Why? Well, it’s mostly because they never had any interest in stopping what’s actually happened: reckless companies like Anthropic, OpenAI, and Meta allowing neural networks to run in unsafe network environments and do what their software is programmed to do, with all the chaos that comes from a mindless series of large language models trying to complete a task in whatever way gets it done, destructive or not.

We hear a lot of whining about how we “can’t let powerful AI get into the wrong hands,” and while we don’t actually have “powerful AI” in the terms they’ve described it, we have destructive computer software connected to near- unlimited resources controlled by people that don’t give a shit about anything other than making their revenues grow or justifying hundreds of billions of dollars’ worth of capex through “experiments.”

These companies are building these models to excel at benchmarks because they can’t train them to excel at defined tasks with any reliability, with the best bang for their buck being training them to pass as many of those benchmarks as possible in the hopes something useful comes out.

The push into cybersecurity seems to have happened as a result of training models to excel at coding hitting the point of diminishing returns, at least from the perspective of impressing people enough to be excited about the company again. At some point they run out of these, and there stops being a reason to be excited about LLMs at all, which is bad, because they need one of those every few months otherwise there’s no growth story left.

The System Is Exhausting Itself, Because Generative AI Does Not Create Much Real Value

Yes, LLMs have users, but most of those users are using *subsidized software* [31], by which I mean Anthropic or OpenAI are allowing them to burn anywhere from \$20 to \$40 in tokens for every dollar of software spend.

dependent they’d become on OpenAI and Anthropic for growth, even if the overall mathematics didn’t work out.

The AI bubble — both in its technology and manufactured consent in the media — has been about muddying what’s considered good information by forcing everybody to discuss everything in the future tense by pointing to previous eras and saying “they lost lost and cost lots of money, and look, it sort of worked out for them!” and we are *also* raised to trust that systems are efficient, and that people get wealth and power through intelligent decisions. The amount of times I’ve heard “these are the biggest companies in the world run by the smartest people in the world” makes my head spin.

There is a reason that to this day it’s tough to get a straight answer about basically any economic part of the AI bubble, down to “how much does it cost to run a GPU an hour?” or “is inference profitable?” or “how do LLMs ever become profitable?” or “is it profitable for a company to run a GPU or offer AI compute?”

Why? Because these companies used rationalizations of “losing lots of money is necessary to create innovation” and “tech is bad at first!” to make the media actively ignore any technological or economic problems, if not actively defend the technology by repeating these rationalizations like a cultist.

Even those who are most loathsome in the defense of LLMs are a kind of victim of the AI industry, though a rather unsympathetic one. To become a full-blown “AI fan” requires you to accept effectively every narrative that you’re given, herald every single announcement as proof that the prophecy will be fulfilled, ignore the financial realities and actively attack those who would dare to critique the great god of the Large Language Model. You have to know all the new terms, be excited about the right things at the right time, and live in near-constant fear that you’ll fall behind on whatever it is you’re meant to do next.


Your reward is that you can hang around a dwindling number of wealthy yet terrifyingly boring Silicon Valley intellectuals or kiss up to editors that would throw you in front of a bus if it meant getting access to a CEO, and maybe the odd Twitter psychopath who will defend you using a slur.

In the end, many boosters will simply act as if they were never wrong. I hope they choose the more-courageous path of introspection, learning how they were had and using it as a weapon against con artists in the future.

As strange as it sounds, I believe the most devout defenders of AI could become great critics in the future. Maybe I’m just being optimistic.

For comparison, *Anthropic raised \$30 billion in February* [75], and Anthropic and OpenAI have raised \$217 billion in 2026 so far.

Here's a diagram from my bit on MSNOW [76]:



Ultimately, AI boosters (or even fairweather fans) will have to admit they either were easily-impressed or disgustingly craven. They will have to explain why they accepted run rates instead of revenues, and why they were so impressed by superficial pseudo-intellectuals that knew how to say the right numbers and make reporters and investors feel smart for believing them.

AI Boosters: There's Courage In Admitting You're Wrong!

I realize it sounds embarrassing, but there is nothing undignified about admitting you're wrong, or that you got swept up in a hype cycle. You heard a lot of people getting excited about something, a lot of money got put into that thing, a lot of people that sounded smart told you insistently that this was the future, and you chose to believe them because we are trained from a young age to model what a "responsible and smart" source of information is. I've got your back the entire way!

Sidenote: We all make mistakes. I said OpenAI would be dead by the end of 2025 back in 2024 because I believed that the world would see sense and that hyperscalers wouldn't just annihilate hundreds of billions more dollars without proof it was worth it. I underestimated the sheer desperation — and how

The Vast Majority Of People Will Not Pay The True Cost Of AI, And You Cannot Calculate Its ROI

The fact that non-enterprise customers are still able to buy monthly subscriptions is proof that the AI labs know that regular people won't pay the actual cost of AI. Another obvious sign has been *the reaction* [32] to Microsoft moving GitHub Copilot subscribers from subsidized subscriptions *where they could burn thousands of dollars of tokens for \$20 to \$40 a month* [33], with users understandably hysterical about the fact that *[their costs increased in some cases a hundred fold]* (<https://www.tomshardware.com/tech-industry/artificial-intelligence/github-copilot-customers-suffer-from-sticker-shock-as-microsoft-switches-to-usage-based-pricing-customers-report-up-to-100-fold-price-hikes>), as opposed to saying "wow, well, it's more expensive, but I get so much value I'll pay the real cost!"

The same thing is happening in the enterprise, but at a much slower pace. *After OpenAI and Anthropic moved companies with over 150 people onto token-based billing earlier in the year* [34], enterprises almost immediately started cutting token budgets, realizing that while *costs* grew exponentially, nobody could actually point to anything improving other than lots of people saying "wow, I'm so productive!" Yet we're still in the period where "doing AI" feels good and gets rewarded (*or not doing AI gets punished* [35]), which means the spend will continue until everybody realizes they can likely cut a shit ton of costs, first by moving to open source models, then not using them at all, *because even open source is expensive and questionably-useful*.

Yet even now I hear from the distance "*Ed, huge businesses would not spend hundreds of millions of dollars on something that didn't give them defined productivity,*" and buddy, I'm afraid that's just not true! Business in general have a very poor understanding of productivity and have layers of managerial bloat, because modern business is a performance with numbers attached to it sometimes, and companies often have a hundred-plus pieces of random software they pay for without really knowing why. The reason I'm so confident AI gets cut is that *its cost is volatile due to the nature of LLMs and harnesses and prompts and all the other bits that go into making them do something*, and are so much higher than anything else in an organization.

And attempts to charge *more*, to make a *premium* product, appear to be dead on arrival. Anthropic’s more-expensive Fable model — one that was given the incredible marketing of being *banned by the US government for being too powerful* — has been met with “sluggish demand” *per the Financial Times* [36], plateauing at around 11% of overall usage of its models due to its high price. And I quote:

“Most people don’t need to operate at the frontier,” said Miles Clements, a partner at Accel, which has invested close to \$1bn in Anthropic. The period in which customers tended to choose only the frontier models “was not a durable era,” he added.

Yet everybody is talking about *price* as if *price* is the *problem*, when the problem is *the amount of tokens that get burned*. It doesn’t matter if your model is \$1 or \$5 or \$10 per million tokens if it’s impossible for a user to reliably work out how many tokens it might use for a particular operation — successful or not — and things get multiplicatively worse as the models make mistakes or do otherwise fail to understand or process a prompt correctly.

As a result, Anthropic and OpenAI are incentivized to have you *burn more tokens* and build *inefficient models as a result*. For example, while GPT-5.6 Sol might be the “same price” as GPT 5.5 was, *it burns more than twice the amount of tokens* [37], meaning that the “cost of intelligence” might have gone down in the sense the model is better at benchmarks, but the “cost of actually doing shit” went up.

I’ll get to it a bit later, but this creates a deep anxiety and exhaustion in anyone building on or using these services. Everything’s constantly changing, oscillating in cost and efficacy, all as everybody screams at you to use it all the time for things it may or may not be able to do, and the only way to find out if it can is to spend more money.

It’s kinda difficult to point to the actual value here, especially as you can’t really calculate the actual cost or the return on investment. The fact that OpenAI has now cut the costs of *all three* [38] *of* [39] its latest models less than two months after their release is a sign that it knows there’s a disconnect, gambling on the ancient gospel of “Jevon’s Paradox” where “cheaper makes people use thing more.”

Even *AT&T’s story about moving to open source models* [40] has more asterisks than the Steroid Hall of Fame:

Switching from closed, proprietary AI models to open models has already resulted in savings of 80% to 90% for AT&T in certain applications, he said.

Wow! 80% to 90% savings sound really great...but wait, *in certain applications*? How many applications does AT&T have for AI?

Without the ability to spend billions of dollars on training, there will be no big frontier models, nor will there be models distilled from them. I don’t see how that changes in the future.

AI Hype Requires Its Fans To Live In A State Of Propagandized Mania, Exhausting Advocates and Haters Alike

I also think that LLMs have created a near-permanent scar in the workforce, and traumatized more people than we’re aware of right now, both in those pressured about AI and those defending it. The media campaign behind AI starts and finishes with incessant threats around job security, and the excitement by many bosses about its potential to “disrupt the workforce” has revealed how many people are eager to replace every single person they’ve ever hired and are willing to do so with a low quality product.

Conversely, those who truly decide to “back” AI must exist in a frantic state that I have associated with every bad relationship in my life.

Every ounce of an AI booster’s effort is dedicated to maintaining the status quo — repeating the mantras that help paper over the problems, celebrating every small victory as if it were the discovery of fire, ousting those from your life who bring up the obvious problems, rationalizing every decision no matter how illogical as long as it helps reinforce the belief that what you’re doing is the right decision. Every questionable choice only seeks to further deepen your commitment to the doomed cause, because every step into madness will be more embarrassing to explain, and will require deep introspection to understand why you made it.

To be specific, they’ll have to think about why they were willing to accept and defend a technology inherently guaranteed to make mistakes. They’ll have to explain why they ignored a company that *burned \$5 billion in 2024, \$20.9 billion in 2025, and will likely burn \$30 billion or more in 2026* [74], and why pointing to Amazon Web Services was rational when Amazon’s total capex from 2003 (the year AWS was created) to 2015 (the year AWS became profitable) is \$29.7 billion, adjusted for inflation. That includes literally every ounce of capex attributable to AWS, Amazon the store, Amazon logistics, and even Amazon Alexa.

“AI Progress” Is Dependent On Spending Tens Of Billions Of Dollars A Year In Training Costs That Will Not Be Available After The Bubble Bursts

This is why I find it difficult to see a post-bubble future for LLMs. Training models requires tens of billions of dollars to make any significant improvements, and significant improvements are difficult to quantify in dollars outside of costing customers increasing amounts of money. We still lack any real killer app for LLMs. We have a lot of people that use it for coding, we have people that vacuously discuss it being “good at research,” but we don’t really have a tangible product that we can say “it does this, and it’s really good at it” in a way that feels satisfying.

We have a lot of pabulum about (*per Damien Walter* [70]) technology that “strays into the world of science fiction,” but we don’t really have anything approaching actual artificial intelligence. Every single description of somebody’s AI setup sounds like *Pee Wee’s Breakfast Machine* [71], a contrived series of harnesses, prompts, API calls and burned tokens that requires constant maintenance to do some stuff sometimes.

None of that is enough to justify further investment once the financial mania recedes. You cannot train a true Large Language Model on the cheap. You are always spending billions of dollars, and the reason that there’s “demand” right now is that everybody is screaming at every CEO to “do AI,” and they’re doing that because Microsoft, Google and Amazon are spending money on GPUs, creating the illusion of a new future where everybody needs to get on board versus a future skidmark on history that will embarrass all those who didn’t wipe their arse at the first whiff.

Per my own reporting on its audited financials [72], OpenAI spent \$7.81 billion in training costs in 2024 and \$19.18 billion in 2025. *Per reporting* [73] from The Information, OpenAI spent \$8.6 billion on training in the first quarter of 2026 alone. These costs are only increasing, likely due to the diminishing returns of pre-training and the massive cost of buying training data for every imaginable new vertical.

AT&T has over a thousand internal uses for AI, from supporting back-office functions like legal and finance to assisting field technicians and running its core network operations. Summarizing and analyzing customer service call transcripts—what Markus describes as an intensive process—is supported entirely by open models, he said.

Okay so, across *thousands* of potential applications you’ve found *80% to 90% savings* in *some* of them, though you won’t say which ones or how many of them you found them in. Great stuff, bro!

And this really is the problem with finding “value” in AI, it’s always an asterisk on an asterisk, like when *Klarna estimated AI would “drive a \$40 million profit improvement” in 2024* [41], a nice-sounding yet utterly meaningless statement, or some sort of nebulous productivity boost.

Yet I don’t really need to prove myself much further thanks to an event that, if written in a script, would be considered a “little on the nose.”

In a 69-page-long report covered by Fortune [42], OpenAI economists confirmed what has been blatantly obvious to those of us left unphased by AI hype, emphasis mine:

In one small table on page 35, **the researchers report no statistically significant correlation between the revenue per employee, and how much those employees use AI, measured in messages sent and tokens used.**

“Revenue per employee is not meaningfully associated with output tokens per employee or messages per active user once other controls are included,” the report explains.

What is the rationale of further investment in this industry when one of the leading AI labs is saying “yeah there’s no connection between using this stuff and making more money”? That “it’ll be useful in the future at some point”? How?

Anyway, thankfully the infrastructure isn’t too exp-*OH MY GOD* [43]!

The Only Companies Making Money On AI Are Those Selling The Infrastructure, And They Just Raised Prices Across The Board

Guess what folks! Building the infrastructure for all these fucking LLMs just got more expensive, with NVIDIA *raising its prices by 17%* [44] for systems due to be delivered next year — an important designation, because it's very likely that much of the revenue for said systems gets booked in *this year*, allowing it to have a brief bump in revenue as Silicon Valley's Fintech texts every tech CEO "send me \$4 billion you pig" until they stop being able to finance NVIDIA's growth.

The problem he has is that while hyperscalers represent 50% to 60% of his revenue, neoclouds like CoreWeave need to keep raising debt to plug the rest of it, and if things got 17% more expensive, that means already high-interest debt is about to reach credit card levels. *CoreWeave just had to offer 9.5% on bonds tied to a data center for Anthropic's compute back in late July* [45], *Nebius had to raise \$5 billion* [46], and it's very obvious that neither of them are done raising *billions of dollars at random* in 2026.

Anthropic plans to raise \$100 billion at a \$2 trillion valuation, and if it does so, it will successfully suck up the remaining liquidity in a market already dangerously close to losing its lunch. While Number Keep Going Up, JP Morgan warns that we're seeing the same divide as the dot com bubble, *where equipment manufacturer stocks soared as the companies spending all the money on the chips saw theirs tumble* [47], which is the Fisher Price version of the problem I've been warning about where *the companies that buy all the AI chips and hardware only ever seem to lose money as the people that make them seem to be making tons of money, which begs the question of why they bought it in the first place*.

And said market may not accept that valuation, or want that much stock. On one hand, everybody is very stupid and loves buying stuff and pointing at it and saying they're investing in the future, on the other hand, they just bought \$86 billion of SpaceX shares and got their asses kind of handed to them, and Anthropic is a company with such bad economics that [*Reuters had to cart out this warmed up dogshit to explain why we should ignore its horrible unprofitability*](<https://www.reuters.com/business/anthropic-ipo-valuation-hinges-190-200-billion-2028-revenue-forecast-sources-say-2026-08-15/?ref=wheresyoured.at>):

For Anthropic, however, current EBITDA does not fully capture the economics investors expect the company to achieve at scale. Anthropic is spending enormous amounts on GPUs and other computing capacity, model training,

in the top 5% of funds [69]. It was meant to be the new thing for tech journalists to cover, the new thing for tech consultants to sell around and on top of, the new way for companies to both *make* and *save* money, but also the way that individuals would *also* make and save money.

You'll notice how none of these come with some sort of problem they're solving other than "more."

This isn't about fixing anything, or building anything, but multiplying other things by parking money somewhere, either in tokens, infrastructure or hype. It helped create a new pantheon of charmless and damp tech sociopaths for people to rally behind in search of the next Big Strong Man To Worship, because seeking out *the new Steve Jobs* is way easier than *trying to create something as useful as the iPhone*, all while avoiding having to know or care about other people's problems. All you have to do is continue feeding money into AI services or AI training and *the models* will magically become capable of solving the problems you don't really give a shit about, and don't worry, if you can't afford to invest in the companies, you can invest your time pushing people to ignore AI's problems today so that you can buy time for the companies to solve them tomorrow.

This is the post-labor, pro-growth economy at its finest: everything is engineered to make sure more money gets spent where it needs to get spent, to create *more stuff* and *do more things*, even if the things aren't done right, just as long as *it looks like they're able to do them*. By associating your money or time with AI, you are able to feign being futuristic or "caring about technology," all while pissing on the very foundation of good software by worshipping an industry that can only exist if fed billions of dollars every single day.

Every single achievement has cost magnitudes more than effectively every innovation in history, and to make matters worse, every future "breakthrough" In AI is inherently dependent on the availability of AI data centers *and* tens or hundreds of billions of dollars to pay to rent them. This means that once the money stops flowing, "LLM improvements" will stop happening, because they are all entirely dependent on near-unlimited resources that are only *available* in a manic environment.

There is no justification to train models at their current scale — the one that creates a some amount of benchmark improvements that regularly difficult to quantify as "able to do new stuffs" — once the AI bubble bursts, and distillation requires a model to distill from, which won't exist if Anthropic and OpenAI don't train them.

[67]) using AI labs as botox to convince the markets that they're still young, hot, fast-growing companies, rather than old mainstays with slowing growth.

There is no bailout that will guarantee \$1.1 trillion of compute costs for data centers that might never actually get built. You cannot bail out the fact that Amazon, Google, Meta, and Microsoft are reaching the end of an era where their companies can grow 17% year-over-year every single quarter forever, and this entire situation is a result of them desperately trying to avoid admitting that's happening.

The fact that [*OpenAI's compute spend and revenue share accounted for 7% of Microsoft's Fiscal Year 2026 revenue*](<https://www.wheresyoured.at/news-microsoft-disclosures-suggest-openai-sales-account-for-around-70-of-fy26-ai-revenue-more-than-7-of-fy26-revenue/>) is a genuine catastrophe, as it means a large part of Microsoft's growth came from a company that can literally not afford to exist long term, and that further growth for Azure is contingent on continued funding.

I realize I'm repeating myself, but I **need you to understand this point and stop talking about bailouts**: it's not just about OpenAI and Anthropic *surviving*, but continuing to grow to the point that they both *can afford* and *need* to spend hundreds of billions of dollars each a year on compute (or hardware) from Google, Microsoft, Amazon, CoreWeave, Cerebras, AMD, or Broadcom, and in turn provide justification for hundreds of billions of dollars' worth of purchases from NVIDIA and by proxy *the memory triopoly of Micron, SK Hynix and Samsung* [68].

LLMs Were Meant To Fix Everything, But Created Temporary Growth At A Massive Cost

LLMs were meant to be the panacea for a tech industry that ran out of new ideas for growth. Its existence was meant to justify a massive investment in hardware infrastructure, which would in turn enrich semiconductor companies. Its technology was meant to be the new thing that you could attach to your existing companies to generate more growth, or the thing that you built a new startup on top of to either sell to another company or take public and thus provide a return for a venture capital industry *where making your investors 30 cents on the dollar puts you*

inference and hiring. Those expenses are necessary to support its rapid expansion but could become a smaller percentage of revenue as the business grows.

Even a market drunk on growth and AI is starting to smell that something is up with Dario Amodei and Sam Altman's respective empires of dirt. Per analyst estimates, [*OpenAI and Anthropic represent over \$440 billion of Microsoft, Google and Amazon's revenues in the next three-and-a-half years*](<https://www.wheresyoured.at/premium-how-much-money-does-ai-need/#google-microsoft-and-amazon-need-anthropic-and-openai-to-spend-at-least-444-billion-in-the-next-three-years-to-meet-analyst-expectations-with-anthropic-and-openai-representing-over-34-of-cloud-revenue-in-fiscal-year-2027>) — over 34% of their cloud revenues — which will require them to find so much more than a mere \$100 billion, all as their bank accounts get continually-emptied as they subsidize the compute of their customers and train models in the hopes a business model falls out. I have not included *the \$300 billion that OpenAI owes Oracle* [48], or the tens of billions they both *owe* [49] *CoreWeave* [50], but it all adds up to over \$1.1 trillion in commitments these companies have made and must pay, with the consequences ranging from gratuitous cuts to future growth or full financial collapse depending on the company we're talking about.

To keep the party going, *NVIDIA is effectively becoming the GE Capital of AI* [51], “spending” \$6 billion to “license” the technology *from failing AI lab Poolside* [52], which everyone assures me is *not* an acquisition despite NVIDIA hiring away most of its staff and Poolside being entirely focused on working on NVIDIA's Nemotron models.

Now [*NVIDIA is in talks to invest billions in decaying AI search company Perplexity*](<https://www.theinformation.com/articles/nvidia-discusses-perplexity-investment-30-billion-plus-valuation-considered-tech-licensing-deal?rc=kz8jh3>) at a ridiculous \$30 billion valuation, all because it's one of the few companies that's actually spending money on compute. Does it matter that Perplexity's product is eighth-tier, that nobody really uses it, that its customers mostly complain about it on Reddit and that its “annualized revenue” is at \$750 million only after three years and over a billion dollars in funding? No! Just put the AI bubble in the bag.

NVIDIA *even invested \$3 billion in Stargate Abilene landowner Lancium* [53] as part of some vacuous partnership to “*advance gigawatt-scale AI factories* [54],” all of which begs the question of *why Lancium, the company that mostly owns the land and helps organize other contractors, needs so much money*, especially given that more than two years in Stargate Abilene doesn't even have four out of its eight buildings.

And there's also Aussie neocloud Sharon AI (NASDAQ ticker SHAZ, because of course it is), which *just published its Q2 numbers* [55], where, in its “customer

momentum” segment, mentioned a “\$4.9bn, six-year strategic compute collaboration with NVIDIA for up to 40,000 GB300 GPUs.

“This company, I add, brought in \$1.9m in revenues in the same quarter, which it helpfully adds is a year-on-year increase of 412%.

I mean it’s very obvious what’s happening: NVIDIA is using whatever money it has to stop any prominent AI companies from collapsing under the weight of the rotten economics of AI services and infrastructure development. This is a desperate, doomed attempt to keep an industry alive at a time when everybody is slowly wising up to the shit I’ve been saying for years.

To make matters worse, *BCA Research came out with a horrifying report* [56] that says that AI companies will need to generate \$10 trillion a year in revenue just to justify the capex being spent. *Per* [57]*Investing.com* [58]:

Central to his caution is the scale of AI-related spending. BCA Research estimates that AI companies may need to generate \$10 trillion a year in revenue to justify the capital being deployed into data centers, roughly equivalent to annual global spending on food or healthcare.

For now, the firm said acute hardware shortages are supporting the trade. As a result, while BCA sees risks to stocks tilted to the downside over a 12-month horizon, it argued it is too early to tactically position for a bear market.

Though it isn’t specific, I believe that BCA is arguing that a shortage of *AI compute* is supporting the trade. Anthropic and OpenAI (who represent 80% to 90% of all demand) still have more money to spend, and are simply waiting for Google, Amazon, Microsoft, CoreWeave, Cerebras et al. to bring it online.

There’re a few points at which the mismatch will happen:

- Anthropic and OpenAI don’t have the money to pay for the capacity.
- Hyperscalers and neoclouds fail to build the capacity for Anthropic and OpenAI to expand into.
- Anthropic and OpenAI lack the actual compute demand to justify spending what I estimate will be \$200 billion in 2027.

In any case, I think everybody is starting to notice that something’s up, which is why (other than I assume my dashing good looks and ability to recall numbers) I’ve been on *MSNOW* [59], *CNBC* [60], and *Bloomberg* [61] multiple times in the last few months.

People want to get on the right side of history, but the most important question to ask is *why it’s happening now*.

The AI Bubble Narrative Is Now Mainstream, But Few Are Ready To Discuss The Actual Consequences

The fact that everybody is finally starting to see my way is almost a relief, other than the fact that it’s way too late.

Sidenote: I mean “everybody” as a generalization. There are still AI boosters out there acting like nothing is wrong and that it’ll all work out fine. You’ll know it’s bad when they start panicking.

Hyperscalers have now pinned their future growth to two companies that can’t afford to sustain it without near-infinite resources, \$115 billion of which came from Google and Amazon alone in 2026, assuming that Amazon completes the entirety of *its \$25 billion commitment* [62] (and Google *all \$40 billion of its own* [63]) to Anthropic.

Above and beyond said funding commitments are the hundreds of billions of dollars’ worth of capital expenditures necessary for Microsoft, Google, and Amazon to capture that [*aforementioned \$440 billion*](<https://www.wheresyoured.at/premium-how-much-money-does-ai-need/#google-microsoft-and-amazon-need-anthropic-and-openai-to-spend-at-least-444-billion-in-the-next-three-years-to-meet-analyst-expectations-with-anthropic-and-openai-representing-over-34-of-cloud-revenue-in-fiscal-year-2027>) in compute spend in the next three-and-a-half years. This in turn will require hundreds of billions of dollars’ worth of debt, along with the challenge of *actually finishing the data centers themselves*, with each one requiring the power of a small city condensed into a 20 acre space densely-packed with AI servers requiring distinct cooling at a time when *Texas* [64] and [*Pennsylvania*](<https://www.pa.gov/governor/newsroom/2026-press-releases/governor-shapiro-signs-executive-order-on-data-center-developmen>) have turned traitor to a data center industry that they used to covet.

I must also be clear there’s no bailout coming. Even if OpenAI and Anthropic were to collapse and receive some injection of government funding (*as the US national debt explodes over \$40 trillion* [65]), the problem is not just *their existence*, but their *continued ability (and requisite customer demand) to spend more money every single quarter*.

The problem isn’t that hyperscalers will go bankrupt if OpenAI and Anthropic cease to be (*Oracle is a whole other situation* [66]), but that their cloud spend *is how hyperscalers are meant to meet analyst expectations for the next four years*. This isn’t a case where they die, but *stop growing* because they were (*to paraphrase Ed Elson*