

Congratulations, You Saved the Planet by Not Asking ChatGPT

Unless you have been living under a rock, you have probably seen the discussions about AI, data centers, and how much electricity and water they consume.

First, a terminology complaint before I lose my mind: when people say “I do not use AI,” they usually mean “I do not use large language models.” LLMs are AI, but AI also powers search engines, recommendation systems, social-media feeds, fraud detection, navigation, photo processing, translation, and countless other services people use every day.¹

Now imagine an ordinary weekend.

You wake up and check Instagram. You eat granola, or perhaps a traditional Turkish breakfast with *sucuklu yumurta*. You drive to a café, have a couple of coffees with friends, walk around some shops, buy groceries, and return home. You open Amazon, Trendyol, or whatever else you use and order something. Perhaps you are environmentally conscious, so you buy it second-hand from a shop in Kadıköy or through Dolap. You eat dinner, stream something, scroll a little more, and go to sleep.

You are looking forward to your trip to Cologne—or Japan, which appears to be the new default destination—and perhaps you would not be opposed to having a child if the economy were less catastrophic.

Was that day environmentally innocent?

More importantly, would you feel proud because you did not make a single ChatGPT query?

I suspect many people would. However, something about that does not add up, so I decided to look at both sides of the comparison: the global environmental impact of AI infrastructure and the environmental impact of the ordinary habits we rarely moralize.

Everyone Hates AI. Everyone Also Uses It.

The obvious expectation is that people who are deeply concerned about AI would avoid using it. In reality, adoption keeps rising while public trust remains low.

A 2026 Pew Research Center survey found that 49% of American adults had used an AI chatbot, up from 33% in 2024. Among adults aged 18–29, usage reached 66%. At the same time, 63% of Americans said AI was advancing too quickly, and only 16% expected its effect on society over the next twenty years to be positive.^[1]

This is not automatically contradictory. You can use something while distrusting it. You can believe that a technology is useful while opposing the companies, labor practices, copyright rules, politics, or social consequences surrounding it. I believe that those are valid discussions. However, if someone says that using an LLM is environmentally unacceptable, continues using it, and then blames other users for aggregate demand, they are not standing outside the system they criticize.

At the simplest level:

$$\text{Total inference demand} = \text{number of users} \times \text{queries per user} \times \text{computation per query}$$

Demand rises when more people adopt AI, when existing users make more requests, or when each request requires more computation. If people who claim not to use AI actually abstained, total demand would be lower unless the remaining users dramatically increased either their query volume or the complexity of their requests.

Of course, adoption is not driven only by individual chatbot users. AI is increasingly embedded in search engines, office software (I am looking at you, Google Docs), customer service, advertising, social media, and corporate workflows. But aggregate usage does not appear from nowhere. It is composed of individual, corporate, and automatic uses, including the use of people who say they dislike it.

¹Throughout this article, I still use “AI use” as shorthand for deliberate use of generative-AI tools, because that is how the term is commonly used in these discussions.

The Global Picture (Is Not Small)

Obviously, a small footprint per query does not mean the AI industry has a small total footprint.

According to the International Energy Agency, data centers consumed roughly 415–460 TWh of electricity worldwide in 2024. The IEA projects their electricity consumption to reach approximately **945 TWh by 2030**, just under 3% of global electricity demand.^[2] That is more than double today's level.

AI is the largest driver of this growth, although data centers also support cloud computing, streaming, storage, cryptocurrency, conventional websites, and other digital services. Electricity consumption by the hardware most associated with AI workloads is projected to grow by around **30% annually** through 2030.^[2]

In addition, the global percentage can also be misleading. Three per cent of global electricity does not sound apocalyptic. However, data centers are usually concentrated in particular regions, while grids, power plants, and water systems are local. A cluster of facilities can create meaningful pressure even when its global share remains modest.

The IEA expects electricity generation for data centers to exceed 1,000 TWh by 2030. Renewables are projected to meet almost half of the additional demand, but natural gas and coal together are still expected to supply more than 40% of the increase through 2030. Data-centre-related electricity emissions are projected to peak at around **320 million tonnes of CO₂** in 2030, which is less than 1% of global emissions but still a substantial amount in absolute terms.^[3]

So yes: AI infrastructure matters. This is why companies should disclose:

- electricity consumption;
- direct and indirect water consumption;
- the locations of their facilities;
- local grid and water impacts;
- cooling technologies;
- lifecycle emissions from hardware;
- training and inference energy separately;
- emissions avoided through efficiency improvements.

They should also be regulated.

But this standard should not be invented only for AI. Fashion companies, airlines, meat producers, oil companies, electronics manufacturers, delivery platforms, streaming services, and social-media companies also benefit from weak reporting standards, selective accounting, greenwashing, and insufficient environmental regulation.

The answer is consistent environmental governance, not treating one new industry as uniquely sinful while accepting opacity everywhere else.

Demand and Corporate Overbuilding

It is easy to imagine data-center construction as a direct response to the number of questions ordinary people ask. But the relationship is not that simple.

Technology companies build ahead of demonstrated demand because they fear becoming compute-constrained. Infrastructure decisions are driven by:

- expected future demand;
- model training;
- enterprise contracts;
- cloud-service ambitions;
- competition with other technology companies;
- investor expectations;
- strategic fear of not owning enough computing capacity.

In July 2026, Meta was reported to be discussing a deal under which Anthropic could lease as much as \$10 billion of Meta's computing capacity over two years.^[4] The talks were preliminary, and the report did not prove that Meta had built a completely unused data center. It did, however, show that computing capacity can be accumulated for one company's strategic plans and later redirected or commercialized for another purpose.

Meta is therefore not merely responding passively to current consumer demand. It is also attempting to create a cloud-computing business, monetize infrastructure, and position itself for a future in which access to computing capacity becomes strategically valuable.

This does not make users irrelevant. Demand still matters. But it means a new data center is not necessarily being constructed because you asked ChatGPT where to place a comma.

What Does One AI Query Cost?

I am sorry for disappointing you, but there is no universal footprint for an “AI query.”

A short text response, a long reasoning task, an image, and an AI-generated video do not consume the same amount of energy. The model, hardware, data center, electricity grid, response length, batching, utilization, and cooling system all matter.

I know, not the most reliable source, but OpenAI CEO Sam Altman stated that an average ChatGPT query uses approximately:

- **0.34 watt-hours of electricity;**
- **0.000085 gallons of water**, or about **0.32 millilitres**.

OpenAI did not publish a full methodology with these numbers, so they should not be treated as ground truth.^[5]

However, an independent 2025 bottom-up analysis estimated a median of **0.34 Wh** for an ordinary frontier-model query, with an interquartile range of **0.18–0.67 Wh**. A longer test-time-compute query using approximately fifteen times as many tokens reached a median of **4.32 Wh**.^[6]

Google separately reported that the median Gemini Apps text prompt consumed **0.24 Wh** and approximately **0.26 mL of water** in its production environment.^[7]

These results suggest that ordinary text prompts are relatively small, while long reasoning, image generation, agents, and video can be much more demanding.

Suppose someone makes **25 ordinary text queries per day**:

$$25 \times 365 = 9,125 \text{ queries per year}$$

Using the 0.18–0.67 Wh range:

$$9,125 \times (0.18\text{--}0.67) \text{ Wh} = 1.64\text{--}6.11 \text{ kWh per year}$$

Using an illustrative electricity carbon intensity of 0.4 kg CO₂e per kWh:

$$1.64\text{--}6.11 \times 0.4 = 0.66\text{--}2.45 \text{ kg CO}_2\text{e per year}$$

The central estimate remains:

$$9,125 \times 0.34 \text{ Wh} = 3.10 \text{ kWh}$$

$$3.10 \times 0.4 = 1.24 \text{ kg CO}_2\text{e per year}$$

If every request were instead a long reasoning query at 4.32 Wh:

$$9,125 \times 4.32 \text{ Wh} = 39.4 \text{ kWh per year}$$

$$39.4 \times 0.4 = 15.8 \text{ kg CO}_2\text{e per year}$$

That is much larger, but it is still useful to distinguish it from ordinary text use.

Also, the water comparison is less secure. Yes, OpenAI’s 0.32 mL estimate and Google’s 0.26 mL estimate are similar, but water use varies with geography, cooling, weather, and electricity generation. It is also easy to compare incompatible things: a full agricultural water footprint on one side and direct data-center water on the other.

Therefore, the water numbers in this article should be read as approximate operational estimates, not universal lifecycle totals.

Training Is Missing

The per-query figures above describe inference: what happens when an already-trained model generates an answer.

They do not fully include:

- initial pretraining;
- post-training and reinforcement learning;
- failed experiments;
- architecture searches;
- model updates;
- manufacturing of chips and servers;
- construction of data centres;
- supporting storage and networking equipment.

These costs matter. However, there is no transparent public dataset that allows the full development cost of current proprietary frontier models to be defensibly divided by their lifetime number of queries. Training is largely a fixed cost, while inference grows with use. For a model serving billions of requests, inference can eventually dominate its operational energy consumption.

So the honest conclusion is not that training is negligible. It is that its amortized contribution per request is unknown.

This is another reason companies should be forced to disclose lifecycle information instead of asking the public to choose between marketing claims and speculative estimates.

The Scale Issue

Even if individual queries become highly efficient, total energy use will rise if the volume of requests grows faster than the efficiency gains.

AI is highly vulnerable to this because individual usage feels completely free. Users never see a fuel gauge, a utility meter, or an invoice. Furthermore, AI is shifting from active chatbot interactions into automatic background processes, summaries, and search features.

The same 2025 inference study estimated that a deployment serving one billion ordinary queries per day would use approximately **0.8 GWh daily**. If only 10% of those requests became long test-time-compute queries, demand could rise to approximately **1.8 GWh daily**.^[6]

This creates a serious environmental challenge, but the pattern is not unique to AI. Lowering the cost or friction of a resource consistently drives up total consumption, seen everywhere from fast fashion and fuel-efficient cars to fast shipping and cheap flights.

The core takeaway is that a low impact per use does not guarantee a low impact overall. However, this systemic growth does not mean that every individual use is inherently irresponsible.

Now Compare It with the Rest of Your Life

The following comparisons are not meant to prove that AI is harmless. They are meant to show the scale of familiar activities that rarely receive the same personal moral scrutiny.

Transport

An average petrol car emits roughly **250 grams of CO₂ per kilometre** from the tailpipe, although the exact figure depends on fuel economy, vehicle size, and occupancy.^[8]

Walking has essentially zero direct operational emissions. Cycling also has no tailpipe emissions, although lifecycle estimates sometimes account for bicycle production and the additional food consumed by the rider.

Rail travel is typically far lower-carbon than driving. Our World in Data, using UK government conversion factors, reports that switching from a car to a train for a medium-distance journey can reduce travel emissions by around **80%**.^[9]

Approximate comparisons per passenger-kilometre vary by country and occupancy, but a useful illustration is:

A ten-kilometre petrol-car trip can therefore produce around:

Transport mode	Approximate emissions
Walking	0 direct g CO ₂ e/passenger-km
Train	≈19–40 g CO ₂ e/passenger-km
Coach / well-used bus	≈25–100 g CO ₂ e/passenger-km
Small or medium car	≈148–250 g CO ₂ e/passenger-km

Table 1: Illustrative passenger transport emissions

$$10 \times 250 = 2.5 \text{ kg CO}_2$$

The alternative is not necessarily “never leave home.” It can be walking for short trips, taking rail where available, using a bus, sharing a car, or combining several errands into one journey.

Food

Food production accounts for approximately **26% of global greenhouse-gas emissions**.^[10]

The largest global food-system dataset shows a consistent hierarchy: beef and lamb generally produce far more emissions than poultry, eggs, or plant-based foods. This remains true whether foods are compared by kilogram, calorie, or protein.^[10]

Average emissions per kilogram include approximately:

Food	kg CO ₂ e per kg
Beef from beef herds	99.5
Lamb and mutton	39.7
Cheese	23.9
Pork	12.3
Poultry	9.9
Eggs	4.7
Rice	4.5
Tomatoes	2.1
Wheat and rye	1.6
Peas	1.0
Potatoes	0.46

Table 2: Global average food emissions

A 150-gram beef portion corresponds to roughly:

$$0.15 \times 99.5 = 14.9 \text{ kg CO}_2\text{e}$$

A 150-gram serving of peas:

$$0.15 \times 1.0 = 0.15 \text{ kg CO}_2\text{e}$$

A salad is not automatically carbon-free. Greenhouse-grown produce, air-freighted ingredients, cheese-heavy salads, and food waste all matter. Tomatoes also look relatively high when measured per calorie because they contain few calories.

Still, replacing some beef meals with legumes, vegetables, grains, eggs, or poultry can reduce emissions substantially without requiring everyone to become vegan overnight.

Coffee

One estimate places a 12-ounce black coffee at approximately **0.258 kg CO₂e**, while a dairy latte reaches approximately **0.844 kg CO₂e**, largely because of milk.^[11]

Two drinks therefore produce roughly:

The alternative does not have to be “never drink coffee again.” It can mean choosing black coffee sometimes, using plant-based milk with a lower footprint, avoiding disposable cups where possible, or simply appreciating that coffee is another discretionary pleasure with an environmental cost.

Order	Estimated emissions
Two black coffees	0.52 kg CO ₂ e
Two dairy lattes	1.69 kg CO ₂ e

Table 3: Approximate emissions from two coffee orders

Clothing

Levi Strauss estimated that one pair of 501 jeans produces approximately **33.4 kg CO₂e** over its lifecycle.^[12]

A fast-fashion study estimated that low utilisation can push jeans to approximately **2.50 kg CO₂e per wear**, eleven times the per-wear impact of a traditional longer-use pattern.^[13]

The obvious alternatives are:

- buy fewer items;
- wear each item for longer;
- repair clothing;
- buy second-hand;
- avoid expedited air delivery;
- wash less often and at lower temperatures.

Buying second-hand is generally better than buying a comparable new product, but buying a second-hand item you did not need is still consumption.

Streaming and Social Media

The Carbon Trust estimates that one hour of video-on-demand streaming in Europe produces approximately **55 grams of CO₂e**. The viewing device is the largest determinant: a large television generally consumes much more electricity than a phone or laptop.^[14]

A two-hour film therefore produces roughly:

$$2 \times 55 = 110 \text{ g CO}_2\text{e}$$

Platform-specific social-media estimates are less reliable, but they commonly place video-heavy platforms above text-heavy services. The exact ranking should not be treated as sacred because device, network, video quality, electricity source, and measurement boundary all differ.

The lower-impact alternative is not to delete every app and stare at a wall. It may simply be reducing autoplay, using Wi-Fi instead of mobile networks where appropriate, lowering video resolution when high resolution provides no benefit, or spending less time scrolling when you do not even enjoy it.

Flying

Aviation produces approximately **2.5% of global CO₂ emissions**, and its total warming effect is larger when non-CO₂ effects are included.^[15]

For individuals who fly, a single long-distance journey can dominate many smaller annual habits.

Possible reductions include:

- taking trains for feasible regional trips;
- flying economy rather than business class;
- travelling less often but staying longer;
- combining purposes into one trip;
- avoiding unnecessary connecting flights.

Again, the point is not that nobody should ever visit Japan. It is that a flight remains environmentally real even when it produces a meaningful, joyful, or memorable experience.

Activity	Approximate climate impact
One ordinary AI text query	0.00007–0.00027 kg CO ₂ e
25 ordinary text queries/day for one year	0.66–2.45 kg CO ₂ e
25 long reasoning queries/day for one year	≈15.8 kg CO ₂ e
One hour of video streaming	0.055 kg CO ₂ e
Two black coffees	0.52 kg CO ₂ e
Two dairy lattes	1.69 kg CO ₂ e
Ten-kilometre petrol-car trip	≈2.5 kg CO ₂
150 g serving of peas	≈0.15 kg CO ₂ e
150 g beef meal	≈14.9 kg CO ₂ e
One pair of Levi’s 501 jeans	33.4 kg CO ₂ e

Table 4: Illustrative environmental comparisons

The Full Comparison

The following table combines illustrative individual actions. Boundaries differ, so these values should not be interpreted as perfect equivalents. They are intended to provide scale.

The exact values can be challenged (and you should challenge them)

AI companies must disclose more. Water estimates are especially uncertain. Training costs are missing. Query complexity varies dramatically. Regional electricity mixes matter. One prompt can trigger several model calls. Agents can repeat tasks automatically. Video generation is not comparable to a short text answer.

However, the uncertainty applies both ways. It is not reasonable to choose the highest speculative estimate for AI, the lowest convenient estimate for everything else, and then announce that the moral question has been settled.

I Am Sorry But “Just Google It” Is Not a Moral Philosophy

People opposed to LLM use often respond with some version of:

“Just search Google.”

“Just write it yourself.”

“Just do the task without AI.”

Fine. Now apply that logic consistently.

Why did you drive instead of walking or taking public transport?

Why did you buy another item of clothing rather than wearing what you already owned?

Why did you order something online?

Why did you eat beef instead of lentils?

Why did you drink coffee, stream a series, fly somewhere for a holiday, or replace a device that still worked?

Most of these things are not strictly required for survival. We do them because they save time, create comfort, entertain us, express identity, reduce effort, or make life more enjoyable.

That does not make them immoral.

It also means “you could have done without it” is not, by itself, a serious environmental framework.

The relevant questions are:

- How much harm does the activity create?
- How often is it performed?
- What benefit does it provide?
- What alternatives exist?
- Can the impact be reduced?
- Is the burden local or global?
- Is the user being judged proportionately?

Instead, people often treat old habits as natural and new habits as choices.

Nobody appears personally offended by the environmental footprint of paperback books. Introduce a Kindle, however, and suddenly everyone remembers mining, batteries, electricity, and electronic waste. A Kindle can be better or worse depending on how many books it replaces; that is not the point here.

The point is that familiar consumption becomes invisible, while technological novelty remains morally conspicuous.

This Is Not an Argument Against Joy

I am not arguing that everyone should stop drinking coffee, traveling, buying clothes, eating enjoyable food, watching films, or using technology. In fact, I think people should continue doing these. I do not think environmental ethics should require turning life into a punishment.

If coffee brings you joy, and caffeine, drink your coffee.

If traveling to Japan is meaningful to you, I am not going to pretend that staying home forever is a real alternative.

If a particular meal connects you to your family or culture, or simply you enjoy it a lot, its value cannot be reduced to one carbon number.

If an LLM helps you learn, work, create, communicate, be more productive, or manage a disability, that benefit is also real.

Environmental impact is one consideration among many. It is not the only value in human life. So, what I mean is that my point is not attacking your desires, my point is people show selective moral superiority.

If you accept the environmental cost of your own pleasures because they make your life better and easier, you should at least pause before treating someone else's useful or enjoyable activity as uniquely frivolous.

There Is Having a Child

I could not finish this article without mentioning this, right? IYKYK. Okay. One influential 2017 study estimated that having one fewer child in a developed country could reduce attributed emissions by **58.6 tonnes of CO₂e per year**.^[16]

This number requires a major disclaimer. It assigns parents a fraction of their descendants' future emissions, so it is not directly comparable to the immediate annual footprint of driving, food, or electricity. It is a long-term accounting method, not the emissions released by the biological act of having a baby.

Still, the underlying point is obvious: creating an additional lifetime of consumption has an environmental impact orders of magnitude larger than whether one adult occasionally asks ChatGPT to rewrite an email.

I am not arguing that people should not have children. Okay, maybe I kind of do that. At least, I prefer you doing that. But like, reproductive choices are deeply personal, and reducing a human life to a carbon calculation would be grotesque. However, apparently, I am kind of grotesque.

Also, I am saying that if you are planning to create an entirely new consumer of food, housing, transport, clothing, electronics, and energy, perhaps do not build your environmental identity around scolding someone for asking a chatbot three questions.

So, Are You a Hypocrite?

AI has a real environmental cost, full stop. Data centers consume substantial electricity and water and their rapid expansion can be a real problem for local grids and water systems. More complex models and automated agents can, and will multiply demand. Companies are insufficiently transparent. And obviously, infrastructure should be regulated. So like, the global concern is legitimate.

But the article's argument is narrower: ordinary individual text use appears relatively modest compared with many normalized discretionary habits, and singling it out as a personal moral failure while ignoring larger optional consumption is inconsistent.

If you avoid these other activities too, good for you. Your position is coherent.

If you participate in them but do not judge others, also good for you. At least you recognize that modern life is built from compromises.

If you drink coffee, eat meat, drive, fly, buy clothes, stream television, and use AI yourself while still arguing for stronger environmental regulation, that is also coherent. Regulation and personal imperfection can coexist.

But if you drive to a café, order two lattes, eat beef, buy jeans, stream a film, plan a flight to Japan, consider having a child, continue using AI when convenient, and then announce that somebody else is destroying the planet by asking ChatGPT to summarize a PDF:

Congratulations.

You did not save the planet.

You just chose which emissions felt embarrassing.

Sources

1. Pew Research Center — *How Americans' Opinions and Use of AI Differ by Age* (2026).
<https://www.pewresearch.org/internet/2026/06/17/how-opinions-and-use-of-ai-differ-by-age/>
2. International Energy Agency — *Energy Demand from AI* (2025).
<https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai>
3. International Energy Agency — *Energy Supply for AI* (2025).
<https://www.iea.org/reports/energy-and-ai/energy-supply-for-ai>
4. Reuters — *Meta, Anthropic in Talks for Potential \$10 Billion Compute Lease Deal* (2026).
<https://www.reuters.com/technology/meta-talks-10-billion-anthropic-compute-deal-nyt-reports-2026-07-17/>
5. The Verge — *Sam Altman Claims an Average ChatGPT Query Uses Roughly One Fifteenth of a Teaspoon of Water* (2025).
<https://www.theverge.com/news/685045/sam-altman-average-chatgpt-energy-water>
6. Oviedo et al. — *Energy Use of AI Inference: Efficiency Pathways and Test-Time Compute* (2025).
<https://arxiv.org/abs/2509.20241>
7. Elsworth et al. — *Measuring the Environmental Impact of Delivering AI at Google Scale* (2025).
<https://arxiv.org/abs/2508.15734>
8. U.S. Environmental Protection Agency — *Greenhouse Gas Emissions from a Typical Passenger Vehicle*.
<https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle>
9. Our World in Data — *Which Form of Transport Has the Smallest Carbon Footprint?*
<https://ourworldindata.org/travel-carbon-footprint>
10. Our World in Data — *Environmental Impacts of Food Production*.
<https://ourworldindata.org/environmental-impacts-of-food>
11. CDP — *Brewing a Sustainable Future: The Carbon Footprint of Your Coffee*.
<https://www.cdp.net/en/insights/brewing-a-sustainable-future-the-carbon-footprint-of-your-coffee>
12. BBC Future — *Can Fashion Ever Be Sustainable?*
<https://www.bbc.com/future/article/20200310-sustainable-fashion-how-to-buy-clothes-good-for-the-climate>
13. Li et al. — *Carbon Footprint of Fast Fashion Consumption* (2024).
<https://www.sciencedirect.com/science/article/abs/pii/S0048969724016498>
14. Carbon Trust — *Carbon Impact of Video Streaming*.
<https://www.carbontrust.com/en-eu/our-work-and-impact/guides-reports-and-tools/carbon-impact-of-video-streaming>

15. Our World in Data — *What Share of Global CO₂ Emissions Come from Aviation?*
<https://ourworldindata.org/global-aviation-emissions>
16. Wynes and Nicholas — *The Climate Mitigation Gap: Education and Government Recommendations Miss the Most Effective Individual Actions* (2017).
<https://iopscience.iop.org/article/10.1088/1748-9326/aa7541>