

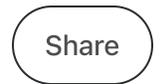
AI's Environmental Cost: Why Should I Care, What Is the Impact, What Is Being Done?

Where I explore the environmental impact that AI is having on us and our world.



TOM MARTIN

JUL 07, 2024



Thanks again to my loyal readers! Appreciate you coming back. 👍

As for you newbies, enjoy the spread, and be sure to come back for more. 🍲

LawDroid Manifesto is a reader-supported publication. To receive new posts and support

my work, consider becoming a free or paid subscriber.

This substack, LawDroid Manifesto, is here to keep you in the loop about the intersection of AI and the law. Please share this article with your friends and colleagues and remember to tell me what you think in the comments below.

This week we're going to switch gears a bit. I'm not going to talk about the whiz bang wizardry of Generative AI and how efficient it can make us, etcetera. Rather, I'm going to talk about GenAI's environmental impact. As it turns out, this is a problem hidden in plain sight. Research underscores the energy-intensive nature of AI operations, greenhouse gas impacts, secondary and tertiary effects that we haven't yet begun to assess. Suffice it to say that with AI barreling down the tracks like a runaway freight train we're going to be living with the consequences for generations. I'm going to explore what that all means and if there's anything we can do about it.

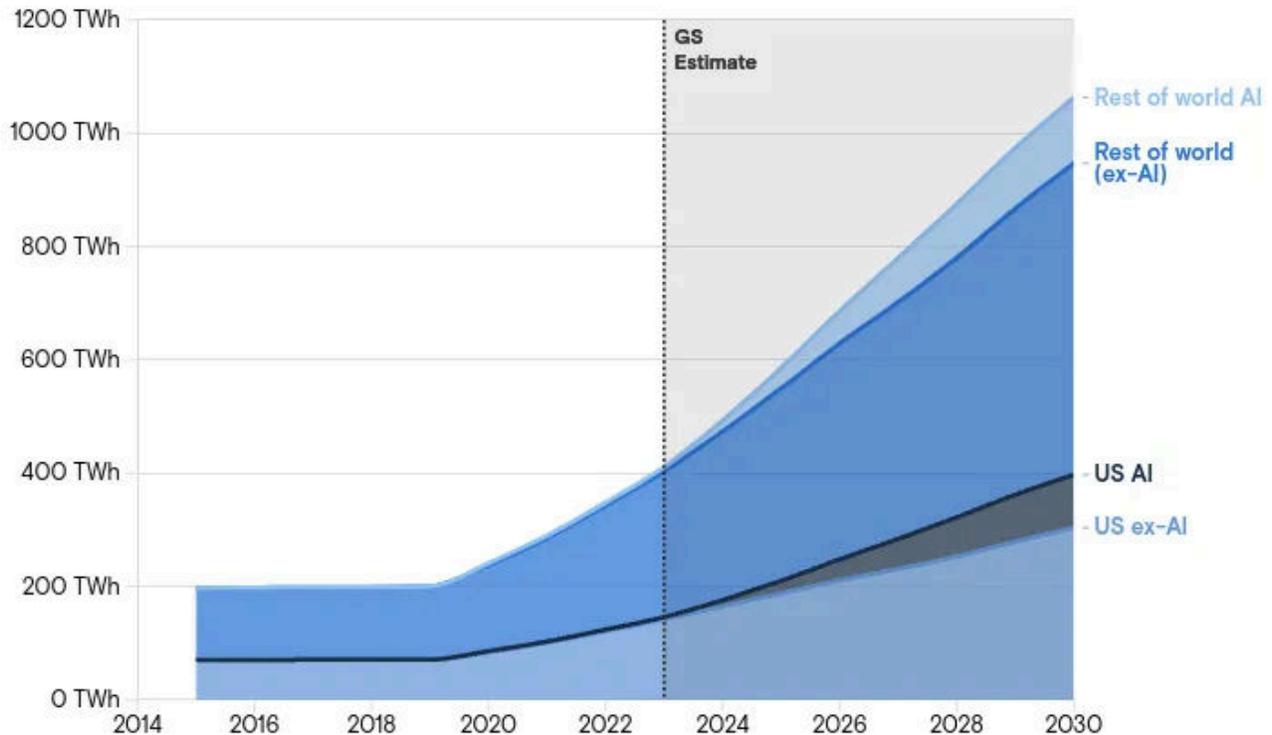
If this sounds interesting to you (and I promise to not to be too dour about the outlook), please read on...

Why Should I Care About AI's Environmental Impact?

It's hard to make people care about something that doesn't affect them personally. And, the purpose of this article is not to make you care nor is it to scare you. Its purpose is to inform. Because AI's energy use and its effects are not something that I've heard talked about at any of the legal tech conferences and webinars that I've attended or presented at as a speaker. Many of the facts I've learned from writing this article have been eye-opening. Yet, framing the discussion in terms of "environmental impact" is unhelpful. Because that perspective is indirect and intangible. How does this affect me? Why should I care? And, there are a few good reasons to care:

(1) **electrical supply is limited.** Although the world's total electrical capacity has increased nearly 300% since 1980 ¹ and is expected to double by 2050, electrical supply is not infinite and brownouts and rolling blackouts are an unwelcome solution to reaching capacity. Imagine asking your favorite AI to summarize the key points of a witness's testimony when your laptops blinks off during trial due to a power failure.

Data center power demand



Source: Masanet et al. (2020), Cisco, IEA, Goldman Sachs Research

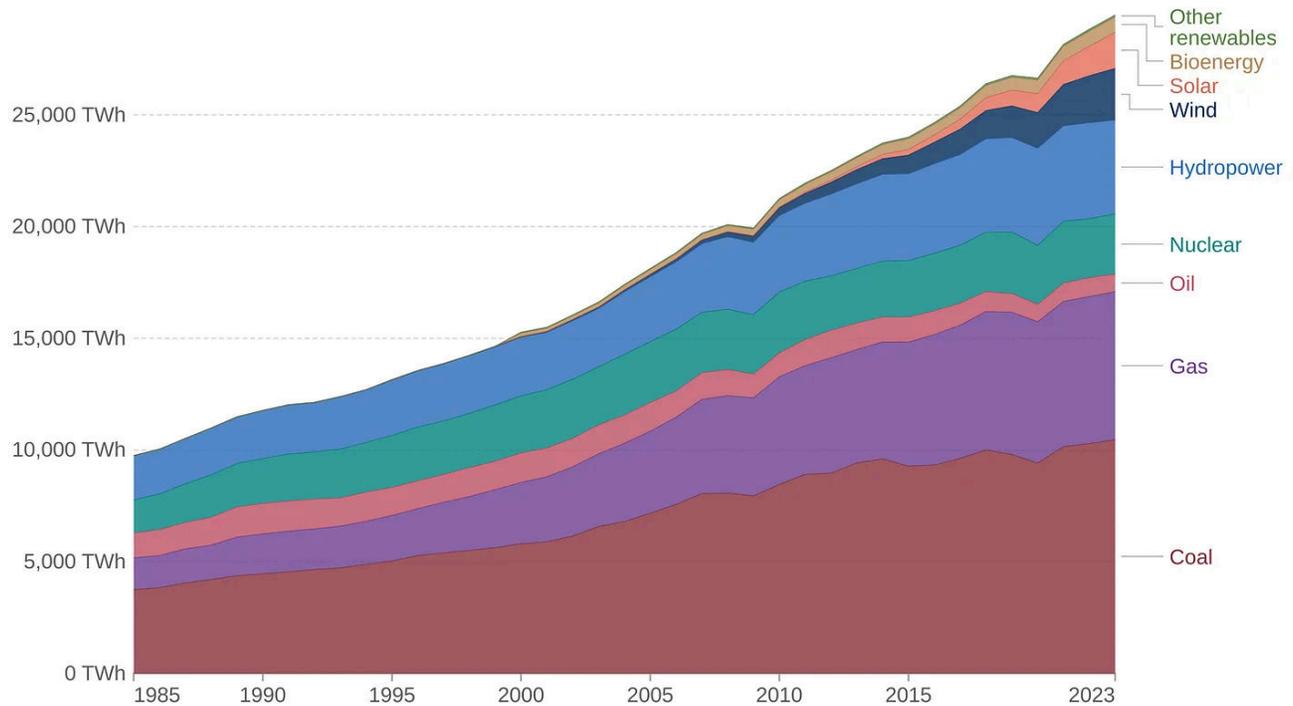
**Goldman
Sachs**

(2) **increased demand taps low quality sources of electricity.** Long thought to be relegated to history's trash bin of outmoded, dirty energy sources, coal plants are being reinvigorated because of the AI boom. ² Imagine going for a jog only to find yourself short-breathed from the 1980's era smog hanging in the air.

Electricity production by source, World

Measured in terawatt-hours¹.

Our World
in Data



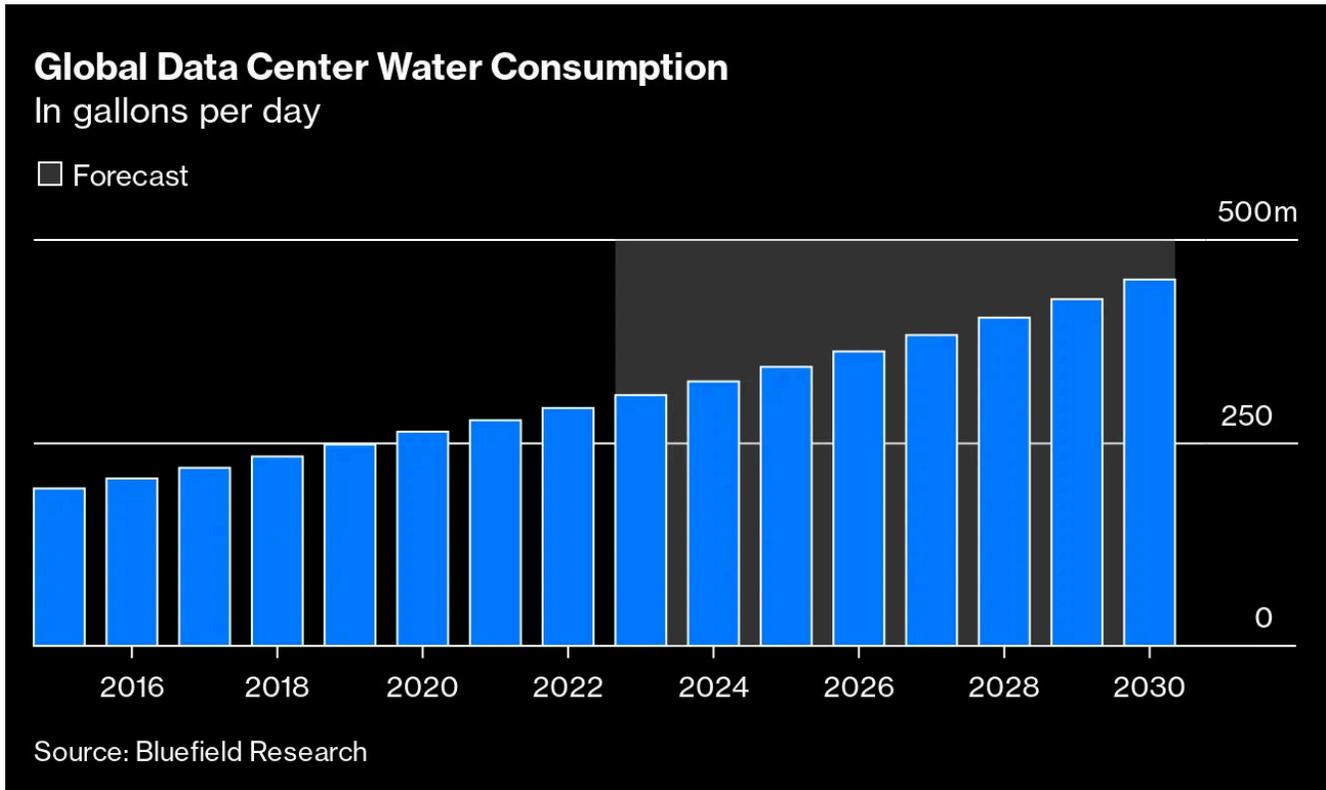
Data source: Ember (2024); Energy Institute - Statistical Review of World Energy (2024)

OurWorldInData.org/energy | CC BY

Note: "Other renewables" include waste, geothermal, wave, and tidal.

1. Watt-hour: A watt-hour is the energy delivered by one watt of power for one hour. Since one watt is equivalent to one joule per second, a watt-hour is equivalent to 3600 joules of energy. Metric prefixes are used for multiples of the unit, usually: - kilowatt-hours (kWh), or a thousand watt-hours. - Megawatt-hours (MWh), or a million watt-hours. - Gigawatt-hours (GWh), or a billion watt-hours. - Terawatt-hours (TWh), or a trillion watt-hours.

(3) data centers consume large amounts of water. The electronic equipment found in data centers throw off a lot of heat and keeping them cool is critical to their operation. A large data center can use between 1 to 5 million gallons of water *per day*, equivalent to the water consumption of a town of 10,000 to 50,000 people.³ Imagine having to fight with big tech for drinking water; some small towns already have.⁴



These are all real impacts, but is AI's impact a difference in kind or degree?

How is AI's Environmental Impact Different Than Other Tech?

There's been much discussion about industry and its impact on the environment. It's all a bit overwhelming and the debate has drawn on for decades without improvement. How then can something as seemingly abstract as artificial intelligence have a real impact on the environment we live in? The action of entering a prompt and getting a response from a computer seems innocuous and routine.

"So what? My toaster and computer use electricity. How is AI's environmental impact any different?"

"How is this any different than running a Google search (that uses AI, right?) or watching a YouTube video?"

These would not be unreasonable reactions. Our use of energy increases as our population increases. Some increases are seasonal, as with air conditioning. Sure, some energy use is spurred by innovation, as was the case with electric vehicles. What makes AI so different that I should care about its impact on the environment?

The difference is the *scale* and the *acceleration* of AI's energy use.

What Is AI's Environmental Impact?

Generative AI is much more energy-intensive than previous technology.

Alex de Vries, a data scientist at the central bank of the Netherlands and a Ph.D. candidate at Vrije University Amsterdam, where he studies the energy costs of emerging technologies, famously raised the alarm about the energy-intensive nature of crypto-currency. de Vries has this to say about GenAI's energy consumption:

[I]f you were to fully turn Google's search engine into something like ChatGPT, and everyone used it that way—so you would have nine billion chatbot interactions instead of nine billion regular searches per day[.] Google would need as much power as Ireland just to run its search engine⁵

Four main components contribute to AI's energy usage: (1) model creation, (2) usage, (3) infrastructure, and (4) the multiplier effect. To get our heads around the scope of AI's environmental impact, let's consider a few key tangible facts and examples:

AI Model Creation

As AI models have grown larger and more powerful, so have their energy needs. An AI model, such as GPT-4o, was created by training a neural network with data. That training is not instantaneous nor energy free.

⚡ Creating a generative AI model called BERT with 110 million parameters consumed the energy of a round-trip flight from Los Angeles to New York for one person.

⚡ Creating GPT-3, which has 175 billion parameters, consumed 1.287 gigawatt hours of electricity and generated 552 tons of carbon dioxide equivalent, the equivalent of

123 gasoline-powered passenger vehicles driven for one year.

⚡ GPT-3 was trained on 10,000 V100 GPUs and each Nvidia V100 uses around 250W. For visual learners, like me, this means that there were 10,000 of these Nvidia cards (see below) running in a hot server room somewhere for 26 days straight!⁶



Nvidia V100 GPUs hard at work

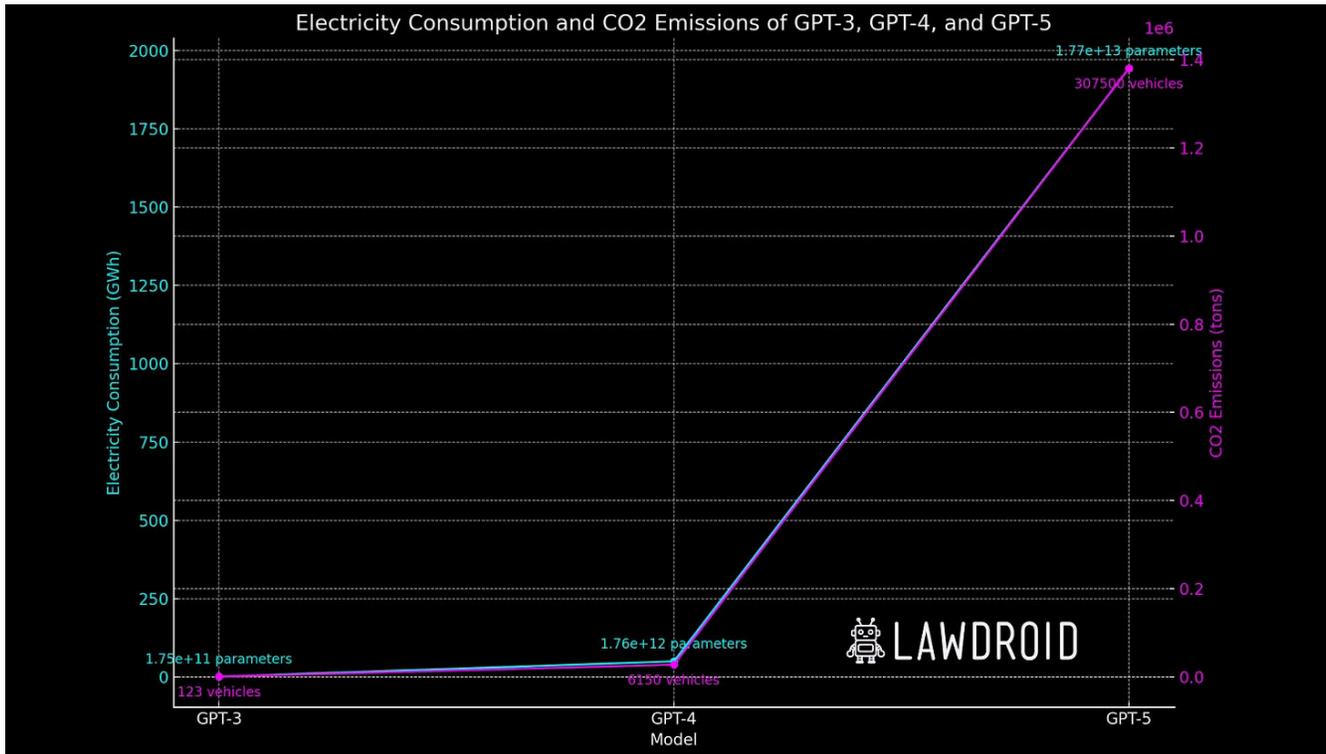
Creating GPT-4, which has 1.76 trillion parameters, required over 50 gigawatt-hours, the equivalent of powering 11,000 households for one year.

⚡ GPT-4 trained for 3-4 months, and generated an estimated 27,600 tons of carbon dioxide, 50 times what it took to train GPT-3.

If we assume that GPT-5's model creation scales in the same proportion as GPT-3 to GPT-4, a comparison and projection might look like this:

- **GPT-3:**
 - Parameters: 1.75×10^{11}

- Electricity Consumption: 1.287 GWh
- CO2 Emissions: 552 tons
- Time to train: 26 days
- Equivalent to 123 gas-powered passenger vehicles driven for one year
- **GPT-4:**
 - Parameters: 1.76×10^{12}
 - Electricity Consumption: 50 GWh
 - CO2 Emissions: 27,600 tons
 - Time to train: 3-4 months
 - Equivalent to 6,150 gas-powered passenger vehicles driven for one year
- **GPT-5 (Projected):**
 - Parameters: 1.77×10^{13}
 - Electricity Consumption: 1,942.5 GWh
 - CO2 Emissions: 1,380,000 tons
 - Time to train: 14 months
 - Equivalent to 307,555 gas-powered passenger vehicles driven for one year



Electricity consumption and CO² emissions of GPT-3, GPT-4 and GPT-5
(projected)

This chart illustrates the significant projected increase in parameters, electricity consumption, and CO² emissions for GPT-5 if it follows the same proportional growth as from GPT-3 to GPT-4.

AI Model Use

The energy used in model creation pales in comparison to model usage, which accounts for 90% of energy expenditure. Every interaction with an AI system comes with an energy cost. As AI models become more widely adopted and integrated into applications, their operational energy consumption is becoming a significant concern.

⚡ ChatGPT consumes 25 times more energy than Google's search engine.

ChatGPT consumes over half a million kilowatts of electricity each day.

⚡ ChatGPT's daily power usage is nearly equal to 17,000 U.S. households, each using about twenty-nine kilowatts.

⚡ ChatGPT consumes about 1/2 liter of water to cool equipment, equivalent to one plastic bottle of water, in a single conversation (20 to 50 questions).

ChatGPT creating a single image consumes as much energy as charging a smartphone.

ChatGPT has over 180 million monthly users; it costs an estimated \$700,000 per day to run and operate ChatGPT, with the cost per query being around \$0.36 cents.

AI Infrastructure

The infrastructure required to support the growing demand for AI is contributing to the its environmental impact. As the AI industry expands and more companies invest in AI infrastructure, the number of power-hungry servers is set to skyrocket.

AI servers consume 30-100 kilowatts whereas a rack of traditional servers in a data center runs on only 7 kilowatts of electricity.

Nvidia, the current leader in the AI server market, shipped 100,000 units last year expected to consume 7.3 times as much energy annually as traditional servers.

Predictions suggest that NVIDIA will ship 1.5 million AI server units annually by 2027 if current trends continue.

Nerdy Sidenote

Koomey's Law dictates that the amount of energy necessary to do a set amount of computing falls by half every two and a half years. This phenomenon might help to partially address energy usage concerns. But, as data processors become more efficient, they also run hotter due to the increased density of circuitry.

AI is predicted to consume twice as much energy as the whole of France by 2030, according to some calculations.

⚡ As data processors (GPUs) become more efficient, they naturally run hotter and need more water cooling.

⚡ Gartner predicts that by 2030, AI could consume up to 3.5% of the world's electricity. "AI consumes a lot of electricity and water. This negative impact should be mitigated," said Pieter den Hamer, VP Analyst at Gartner.

de Vries cautions "if we decide we're going to do everything on AI, then every data center is going to experience effectively a 10-fold increase in energy consumption."

AI Multiplier Effect

There are two multiplier effects.

Model Multiplier. The model multiplier accounts for the environmental costs of numerous AI companies and open-source models. Multiply OpenAI's environmental impact figures by contributions from GenAI competitors Amazon, Anthropic, Cohere, Google, Microsoft, Nvidia and a panoply of open source models.

⚡ **Usage Multiplier.** The usage multiplier considers the widespread adoption of AI across various industries and applications. Multiply GenAI usage across the vast universe of legacy products to incorporate GenAI features and applications plus countless GenAI startups using the technology.

We can represent this AI Multiplier Effect in the form of an equation:

$$E_{total} = E_{OpenAI} \times M_{models} \times M_{usage}$$

Example Calculation

Let's suppose:

- The environmental impact of OpenAI's models E_{OpenAI} is measured in CO² emissions or energy consumption.
- M_{models} is estimated as 7 (considering contributions from Amazon, Anthropic, Cohere, Google, Microsoft, Nvidia, and open-source models).

- M_{usage} is estimated as 10 (considering the widespread adoption of AI).

Then the total environmental impact would be:

$$E_{\text{total}} = E_{\text{OpenAI}} \times 7 \times 10$$

Let's assume E_{OpenAI} is 27,600 tons of CO_2 (from GPT-4):

$$E_{\text{total}} = 27,600 \times 7 \times 10 = 1,932,000 \text{ tons of } \text{CO}_2$$

This equation provides a framework to estimate the total environmental impact by considering contributions from various AI developers and the scale of AI adoption.

What Are AI Companies Doing About Their Environmental Impact?

In January 2024, OpenAI CEO Sam Altman admitted an energy breakthrough is necessary for future artificial intelligence, which will consume vastly more power than people have expected. "I think we still don't appreciate the energy needs of this technology...There's no way to get there without a breakthrough."⁷ Altman was referring to a nuclear fusion breakthrough (and he has invested in a startup pursuing that goal), but he has also suggested adding nuclear fission to the mix of supply.

AI companies are addressing the environmental impact of their technologies through various initiatives. Here are some of the key approaches being adopted:

1. **Efficient Waste Management:** AI is being used to make waste management more efficient, which is crucial since waste is a significant producer of methane and responsible for a large portion of greenhouse gas emissions. For example, AI systems are being developed to improve recycling and reduce waste.⁸
2. **Sustainable Data Centers:** Companies are increasingly focusing on controlling energy consumption in data centers. Sustainable data centers are being promoted as a way to offset the carbon footprint of AI. These centers use AI to track and

analyze data, helping to address environmental challenges like climate change, pollution, and deforestation.⁹

3. **Corporate Environmental Reports:** Companies, like Google, produce detailed environmental reports. However, critics argue that these reports sometimes avoid discussing the specific energy costs associated with AI, focusing instead on broader technological applications to environmental issues.¹⁰

These efforts show a growing commitment among AI companies to not only mitigate the environmental impacts of their technologies but also to harness AI's capabilities to tackle broader environmental challenges.

However, there is a notable lack of transparency about the environmental costs of generative AI, which appear to be soaring. OpenAI has not disclosed the energy consumed or CO₂ produced to train GPT-4, let alone GPT-4o or the current training on GPT-5. The actual energy consumption and the carbon footprint of training and using large AI models makes it difficult to assess the true environmental impact.

In regions like the EU, new regulations such as the Corporate Sustainability Reporting Directive now require data center operators to disclose their environmental impacts. This is part of a broader push to make tech companies more accountable for their environmental footprint.¹¹ In my opinion, environmental impact reports should be legislatively mandated for all tech companies producing foundation models worldwide.

Provisos

This article focuses on CO₂ emissions and energy usage, but AI's environmental impacts go much further. Life cycle assessment is a solid methodology to evaluate not only global warming potential, but also other direct environmental impacts. LCA considers all the steps from production to use and end of life.¹²

This article also does not analyze the ways in which AI may contribute to solving the energy supply and CO₂ emissions problems.

Closing Thoughts

The rapid advancement and adoption of AI is undoubtedly transforming our world in profound ways. But as we marvel at the efficiency and capabilities unlocked by this technology, we must not overlook the hidden costs to us and our environment. The energy-intensive nature of AI, from training ever-larger models to running power-hungry servers and data centers, is contributing to a growing carbon footprint and straining finite resources like water and electricity.

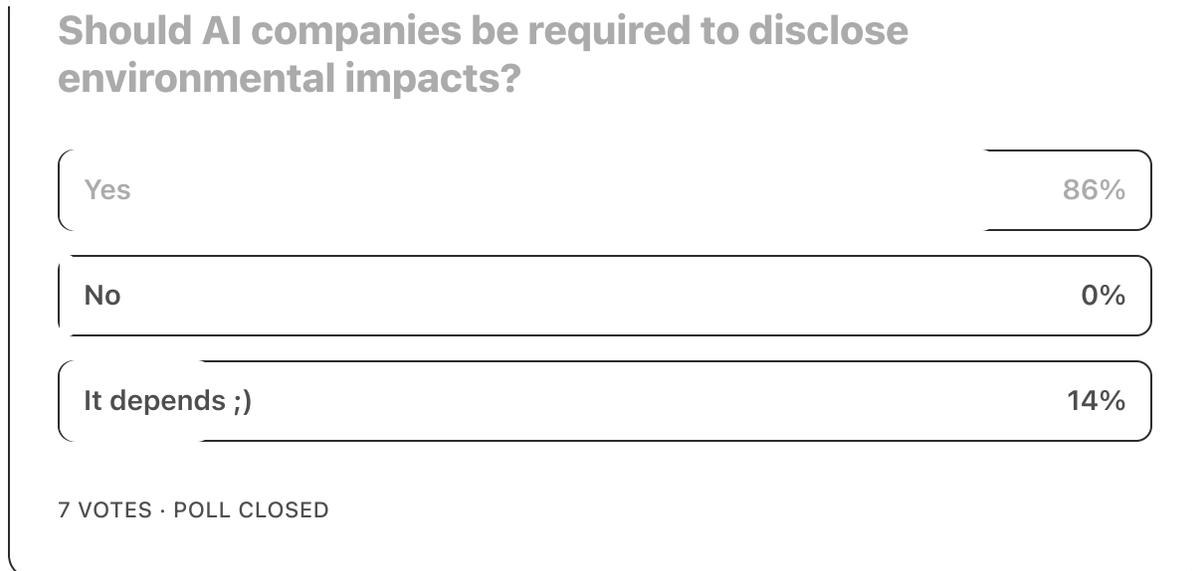
While some AI companies are starting to acknowledge these impacts and explore mitigation strategies, transparency is still sorely lacking. We need a clearer picture of the true costs associated with the AI boom. Efforts to optimize algorithms, utilize renewable energy, and boost sustainability are steps in the right direction, but they must be drastically scaled up to keep pace with AI's rapid growth and its demands upon our infrastructure.

The AI train has left the station, and it's transforming our world in ways we're only beginning to comprehend. But we cannot be passive passengers on this journey. Efficiency and innovation cannot come at the expense of our future. The path forward with AI must be one of sustainability, not just for our own sake, but for generations to come. This is a conversation we need to have now, before it becomes a runaway train leaves us all in its carbon dust.

The purpose of this piece is not to vilify AI but to bring awareness to an underappreciated issue. With eyes wide open, we can work towards solutions that allow us to reap the benefits of AI while being responsible stewards of our planet. It's a balance we must strike, and the first step is understanding the stakes.

The AI revolution is here, but it's up to us to ensure it doesn't come at the cost of our future. As legal professionals at the forefront of AI adoption, we have a unique opportunity and responsibility to shape the path forward – will you join the charge for sustainable, responsible AI innovation?

POLL



By the way, if you'd like to learn more about how AI works and how it will impact the legal profession, you should apply to LawDroid University!

My NEW 5-part webinar series, Generative AI for Lawyers: Empowering Solos and Small Law Firms, is now available at LawDroid University.

LawDroid University is available for free for everyone to use.

- **Free to use** - It's 100% free educational content for everyone, just sign up below.
- **Insightful** - Get educated about the intersection of artificial intelligence and the law as taught by experts.
- **Value Packed** - Filled with videos, summaries, key takeaways, quotable quotes, transcripts and more! Find sessions on AI and the State of the Art, Ethics, Access to Justice, Practice of Law, Education, and the Business of Law.
- **AI Q&A** - Ask a chatbot questions about the content and get fully informed answers immediately.

👉 To immerse yourself in this enriching educational voyage, learn more, or sign up, please visit <https://lawdroid.com/subscriptions/lawdroid-university/>.



5-PART WEBINAR SERIES

Generative AI for Lawyers

Empowering Solos and Small Law Firms

lawdroid.com

- 1 Net electricity consumption worldwide in select years from 1980 to 2022, Statista, <https://www.statista.com/statistics/280704/world-power-consumption>
- 2 AI is exhausting the power grid. Tech firms are seeking a miracle solution, The Washington Post, <https://www.washingtonpost.com/business/2024/06/21/artificial-intelligence-nuclear-fusion-climate>
- 3 A new front in the water wars: Your internet use, The Washington Post, <https://www.washingtonpost.com/climate-environment/2023/04/25/data-centers-drought-water-use>
- 4 A lot of data centers operate where water is scarce. *Id.* In The Dalles, Oregon, a local paper fought to unearth information revealing that a Google data center uses over a quarter of the city's water. In Los Lunas, New Mexico, farmers protested a decision by the city to allow a Meta data center to move into the area. In Cascade Locks, Oregon, residents are also pushing back against a proposed data center that they worry will raise electricity rates and suck up precious water. *Id.*
- 5 The AI Boom Could Use a Shocking Amount of Electricity, Scientific American, <https://www.scientificamerican.com/article/the-ai-boom-could-use-a-shocking-amount-of-electricity>

- 6 AI Chatbots: Energy usage of 2023's most popular chatbots (so far), TRG Datacenters, <https://www.trgdatacenters.com/resource/ai-chatbots-energy-usage-of-2023s-most-popular-chatbots-so-far>
- 7 OpenAI CEO Altman says at Davos future AI depends on energy breakthrough, Reuters, <https://www.reuters.com/technology/openai-ceo-altman-says-davos-future-ai-depends-energy-breakthrough-2024-01-16>
- 8 9 ways AI is helping tackle climate change, World Economic Forum, <https://www.weforum.org/agenda/2024/02/ai-combat-climate-change>
- 9 AI And Environmental Sustainability – A Symbiotic Relationship, Forbes, <https://www.forbes.com/sites/delltechnologies/2023/11/28/ai-and-environmental-sustainability--a-symbiotic-relationship>
- 10 Google's environmental report pointedly avoids AI's actual energy cost, Techcrunch, <https://techcrunch.com/2024/07/02/googles-environmental-report-pointedly-avoids-ais-actual-energy-cost>
- 11 EU Corporate Sustainability Reporting, https://finance.ec.europa.eu/capital-markets-union-and-financial-markets/company-reporting-and-auditing/company-reporting/corporate-sustainability-reporting_en
- 12 Ligozat et al., Unraveling the Hidden Environmental Impacts of AI Solutions for Environment Life Cycle Assessment of AI Solutions. Sustainability, 2022, 14, pp.5172. <https://amu.hal.science/hal-03650884/document>



11 Likes · 2 Restacks

Comments



Write a comment...

© 2024 Tom Martin · [Privacy](#) · [Terms](#) · [Collection notice](#)
[Substack](#) is the home for great culture