


AI & Water Use: Context, Scale, and Justice


Updated with current research — May 2025

The Bottom Line

Yes, AI uses water — mainly to cool data centers and generate the electricity that powers them. But when placed in context, AI's water use is **orders of magnitude smaller** than the activities that actually dominate global freshwater consumption. The nuance matters, especially for justice-focused audiences: the issue is governance, siting, and accountability — not the raw volume.

The Numbers That Matter

 **Water per AI prompt:** A typical mid-length ChatGPT/Claude query consumes approximately **3–10 milliliters of water** (direct + indirect cooling), based on the Li et al. 2023 methodology and updated 2024-25 model efficiency. OpenAI CEO Sam Altman stated in April 2025 that the average ChatGPT query uses 0.3 ml — the lower bound. Even using the more conservative estimate of 10 ml, the comparisons below are striking.

 **The Cheeseburger Comparison:** A single quarter-pound cheeseburger requires roughly **450 gallons (~1,700 liters) of water to produce** — counting water used to grow feed crops, hydrate cattle, and process the meat (USGS; Water Footprint Network). At 10 ml per AI prompt, that is equivalent to **170,000 AI text prompts**. At Sam Altman's 0.3 ml estimate: over 5.6 million prompts. The cheeseburger wins by a landslide.





Other comparisons that put AI water use in scale:

- One pound of beef: **~1,800 gallons** of water (USGS/Water Footprint Network) — equal to ~80 million AI prompts at 10 ml each
- One pound of almonds: **~1,900 gallons** — comparable to beef; often cited alongside AI but vastly larger
- One cotton t-shirt: **~700 gallons** (~2,650 liters) — equal to ~265,000 AI text prompts
- Global agriculture: **~70% of all freshwater withdrawals** worldwide (FAO)
- All data centers combined: **well under 1%** of global freshwater withdrawals

A note on the wide range of estimates: Water-per-query figures in the literature range from 0.3 ml (Sam Altman, Apr 2025) to ~25 ml (early GPT-3 estimates, Li et al. 2023). The wide variation reflects different models, query lengths, data center cooling technologies, and

accounting methodologies (direct cooling only vs. including water used to generate electricity). All estimates, even the high ones, remain minuscule compared to food production.

Global Freshwater Withdrawal by Sector

Sector	Share of Global Use	Context
 Agriculture (incl. factory farming)	~70%	Irrigation, feed crops, livestock. The dominant driver by far. Peer-reviewed FAO data.
 Industry & energy (power plants, manufacturing)	~20%	Cooling, processing, extraction. Power plants are a major user.
 Municipal / household	~10%	Drinking water, sanitation, daily use.
 Data centers (all cloud, streaming, AI)	Well under 1% globally	Mostly indirect (water used to generate electricity). Fast-growing but still tiny in absolute terms.

Sources: FAO AQUASTAT; World Resources Institute; IEA; Li et al. (2023) Nature / arXiv

How AI Actually Uses Water

AI water use is indirect and operates through two pathways:

1. Direct Cooling (Evaporative Cooling Systems)

Large data centers use water to dissipate heat from servers. Evaporative cooling towers are the most common method — water evaporates and carries heat away. The amount varies dramatically by data center design: air-cooled facilities use no direct water; older evaporative systems use significantly more. Microsoft, Google, and Meta have all published (limited) data on their Water Usage Effectiveness (WUE) metrics.

2. Indirect Water (Electricity Generation)

Generating electricity — especially from fossil fuels and nuclear power — requires substantial water for cooling. This 'embedded' water in electricity is often larger than direct cooling water but is almost never reported by tech companies. Li et al. (2023) was the first major study to account for both direct and indirect water use, and found the combined footprint to be 2–3× higher than direct cooling alone.

Why AI Water Use Still Deserves Scrutiny

1. Where the Water Comes From Matters More Than How Much

The same data center using 1 million gallons in rainy Scotland is a very different situation from one using 1 million gallons in drought-stricken Nevada or Arizona. Google, Meta, and Microsoft have all faced local opposition for siting data centers in water-stressed regions. This is a siting and governance problem, not an inherent AI problem.

2. Transparency Is Lacking

Most tech companies do not publicly disclose granular water usage data for individual facilities. OpenAI does not publish water data. Li et al. (2023) had to estimate from indirect sources. This opacity is itself a governance failure — one that environmental advocates can and should target through corporate accountability campaigns.

3. Demand Is Growing Rapidly

The IEA's 2025 Energy and AI report projects data center electricity demand to double by 2030. Since water use tracks electricity use (especially indirect water), this implies significant growth in AI's water footprint absent major efficiency improvements or policy intervention. Li et al. (2023) projected global AI water demand of 4.2–6.6 billion cubic meters by 2027 — a meaningful volume, though still dwarfed by agriculture.

The Justice-Forward Framing

The right question is not:

"Does AI use too much water?"

The right question is:

*"Who decides where data centers are built? Who bears the water costs?
Who gets a say? And who benefits?"*

AI water use becomes an urgent issue when:

- Data centers are sited in drought-prone or water-stressed communities without local consent
- Companies externalize water costs to local ratepayers and ecosystems
- Transparency about water usage is absent or deliberately obscured
- Benefits (AI services, economic development) flow primarily elsewhere

These are governance and accountability failures — not inherent flaws of AI itself. The same critique applies to any large industrial water user.

The Hopeful Part: AI's Water Footprint Is Governable

Unlike agricultural water use — which is politically entrenched, subsidized, and structurally resistant to reform — AI water use is:

- Measurable: Water Usage Effectiveness (WUE) metrics exist and are improving
- Optimizable: Dry cooling and closed-loop systems can dramatically cut water use
- Governable: A small number of large companies control most AI data center capacity
- Regulatable: State-level and federal water disclosure requirements are achievable

Ironically, AI is also being used to reduce water consumption elsewhere:

- Smart irrigation: AI-driven systems reduce agricultural water use by 15–50% in trials (the sector responsible for 70% of withdrawals)

- Leak detection: AI identifies municipal pipe failures faster and more accurately than conventional methods
- Climate modeling: AI accelerates hydrological forecasting critical to drought resilience
- Crop optimization: AI improves yield-per-gallon efficiency, reducing water intensity of food production

The water justice argument isn't 'AI is fine.' It's 'AI's water footprint is small, improvable, and regulatable — and AI done right can help solve the water crises that matter far more.'

Annotated Source List

AI & Data Center Water Use (Primary Research)

1. Li et al. — "Making AI Less Thirsty" (arXiv 2304.03271; also in Communications of the ACM)

Pengfei Li, Jianyi Yang, Mohammad Islam, Shaolei Ren — UC Riverside, 2023 — <https://arxiv.org/abs/2304.03271>

Why it matters: The most-cited peer-reviewed paper on AI water footprints. First to systematically account for both direct cooling and indirect electricity-generation water use. Finds GPT-3 training consumed ~700,000 liters directly and ~5.4 million liters total. Per-query estimates of 10–50 ml for GPT-3. Essential for scientific audiences — cite the ACM version for extra credibility.

2. "How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference"

arXiv 2505.09598, May 2025 — <https://arxiv.org/html/2505.09598v1>

Why it matters: The most recent (2025) peer-reviewed benchmarking of AI environmental footprints across multiple models and queries. Provides updated per-query water estimates reflecting current model efficiency. Ideal for technical audiences who want current data.

3. EESI — "Data Centers and Water Consumption"

U.S. Energy & Environmental Study Institute — <https://www.eesi.org/articles/view/data-centers-and-water-consumption>

Why it matters: Clear, accessible explainer of how data centers use water (direct + indirect) and why location matters more than raw totals. Well-sourced, non-partisan, ideal for a nonprofit or policy audience.

4. Stanford 'And the West' Project — "Thirsty for Power and Water: AI Data Centers Across the West"

Stanford University, 2025 — <https://andthewest.stanford.edu/2025/thirsty-for-power-and-water-ai-crunching-data-centers-sprout-across-the-west/>

Why it matters: Excellent journalism and research on the regional justice angle — drought-prone Western U.S. communities hosting data centers without adequate consent or water offset policies. Strong for equity framing.

5. OECD.AI — "How Much Water Does AI Consume? The Public Deserves to Know"

OECD AI Policy Observatory — <https://oecd.ai/en/work/how-much-water-does-ai-consume>

Why it matters: Policy-focused overview from the OECD arguing for mandatory water disclosure. Builds the governance case well. Good for audiences focused on regulatory or corporate accountability solutions.

Agriculture & Factory Farming Water Use (The Baseline)

6. USGS — "How Much Water Does It Take to Grow a Hamburger?"

U.S. Geological Survey (Water Science School) — <https://water.usgs.gov/edu/activity-watercontent.php>

Why it matters: Federal government source showing 4,000–18,000 gallons per hamburger, depending on production conditions. The authoritative U.S. reference. Useful for the cheeseburger comparison because it's a federal agency, peer-reviewed methodology, and accessible.

7. Water Footprint Network — "Water Footprint Explained"

Water Footprint Network (Hoekstra et al. methodology) — <https://waterfootprint.org/en/water-footprint/what-is-water-footprint/>

Why it matters: The standard academic methodology for calculating water footprints — green water (precipitation), blue water (surface/groundwater), and grey water (pollution). The basis for the 1,800 gal/lb beef figure. Cite this when audiences ask about methodology.

8. FAO AQUASTAT — "Water Use in Agriculture"

UN Food and Agriculture Organization — <https://www.fao.org/land-water/water/water-management/water-use-agriculture/en/>

Why it matters: Authoritative UN data showing agriculture at ~70% of global freshwater withdrawals. The gold standard baseline for all sector comparisons. Updated annually.

9. World Resources Institute — "Agriculture's Water Problem"

World Resources Institute — <https://www.wri.org/insights/agriculture-water-problem>

Why it matters: Strong on systemic causes: irrigation inefficiency, livestock feed crops, and climate stress. Connects water use to broader food system reform. Good for audiences already familiar with food justice.

Helpful Synthesis Sources

10. The Conversation — "AI Has a Hidden Water Cost — Here's How to Calculate Yours"

The Conversation, 2024 — <https://theconversation.com/ai-has-a-hidden-water-cost-heres-how-to-calculate-yours-263252>

Why it matters: Academic-backed, journalist-written explainer that walks readers through how AI water use is calculated and lets them estimate personal usage. Good for interactive or workshop settings.

11. Wikipedia — "Environmental Impact of Artificial Intelligence"

Wikipedia (well-sourced synthesis, continuously updated) — https://en.wikipedia.org/wiki/Environmental_impact_of_artificial_intelligence

Why it matters: Good neutral overview with primary citations pointing to Li et al. and IEA data. Useful starting point for audiences wanting a balanced overview without industry spin.

Prepared for environmental nonprofit use. Statistics reflect best available data as of May 2025. Water figures vary significantly by model, query length, data center design, and accounting methodology. For scientific audiences, primary sources (Li et al. 2023; arXiv 2505.09598) should be cited directly.