RESEARCH PAPER

When Intelligence Costs the Climate The Hidden Environmental Impact of AI

Bhavna[1]*, Priyank Sirohi[2]

[1]*M.tech Final Year (Computer Science), Sir Chhotu Ram Institute of Engineering and Technology Chaudhary Charan Singh University Meerut, U.P. (India) Bhavnawalia1234567@gmail.com
[2]Assistant Professor, Sir Chhotu Ram Institute of Engineering and Technology Chaudhary Charan Singh University Meerut, U.P. (India) Priyanksirohi01@gmail.com

Abstract
We're building AI faster than we can track the mess it's making. To give you an idea of the scale, training just one big model uses as much power as a small town and spits out as much exhaust as over a hundred cars do in a year. If we keep going at this rate, data centers are going to suck up more electricity than the entire country of Japan by 2030. It's a huge drain on water, too—billions of gallons are used for cooling, and usually in places that are already bone-dry. We also can't forget that just building the computers in the first place creates as much pollution as running them does. On top of that, these massive facilities are mostly popping up in minority neighborhoods, which isn't a coincidence and isn't fair. The weird part is that AI could actually help fix the climate by making planes and factories run better, but right now, it's more likely to do more harm than good. Unless we start forcing companies to use clean energy, save water, and be honest about their numbers, we're headed for trouble.

Index Terms: Artificial Intelligence, Carbon Footprint, Data Centers, Energy Consumption, Environmental Justice, Life Cycle Assessment, Sustainability, Water Resources

How to cite this article: Bhavna, Sirohi P. When Intelligence Costs the Climate The Hidden Environmental Impact of AI. Int J Drug Deliv Technol. 2026;16(32s):1126-1131. DOI: 10.25258/ijddt.16.32s.125

INTRODUCTION

1) Background and Motivation

AI has taken over faster than almost any other tech in history—it's hit 40% of homes in wealthy countries in just three years, which is even faster than the jump from flip phones to the mobile internet. But there's a massive physical cost to that speed. The "brains" behind AI don't run on normal computers; they need massive data centers that eat up 10 to 50 times more power than a standard office building. Because the way these models are built (the "transformer" design) means that as the AI gets smarter and bigger, the amount of power and hardware it needs keeps climbing right along with it.

2) Problem Statement

The environmental damage from AI happens in a few big ways: it's expected to gulp down massive amounts of electricity by 2030, and it's already pumping out millions of tons of extra carbon every year—mostly because a huge chunk of that power still comes from coal and gas. On top of that, it's draining billions of gallons of water in the U.S. just to keep the equipment from overheating. The problem is that our current laws can't keep up. Since companies aren't forced to report these numbers and the tech is moving at lightning speed, policy makers are basically flying blind.

3) Research Objectives

This study addresses critical knowledge gaps through five objectives:

RO1: Quantify AI's environmental footprint using lifecycle assessment methodology
RO2: Characterize spatial distribution patterns relative to climate vulnerability
RO3: Evaluate environmental justice implications through demographic analysis
RO4: Model optimization scenarios comparing costs against mitigation benefits
RO5: Develop evidence-based policy recommendations

4) Contributions

This research does four main things: first, it looks at the entire life of the technology, including the pollution created during manufacturing, which shows that most people are actually underestimating AI's carbon footprint by about half. Second, it maps out where these facilities are built to show how they are unfairly placed in certain communities. Third, it creates a math-based system to try and balance using AI with the reality of our limited natural resources. Finally, it uses all this data to suggest specific, practical rules for the government to follow.

RELATED WORK

1) Background and Motivation

Early studies by Strubell and others showed that training an older, simpler model like BERT used about 1,507 units of electricity, but more recent work by Patterson's team estimated that the much larger GPT-3 required over 1.2 million units. While training takes a lot of power up front, new research shows that actually using these models (called inference) can be even more taxing; for popular apps, the daily energy used to answer user questions can overtake the initial training energy in just three to six months. Looking ahead, the International Energy Agency predicts that the world's data centers will need 945 billion units of electricity by 2030, with the massive growth in AI being the main reason for that surge.

2) Water Footprint Analysis

---

*Author for Correspondence: Bhavnawalia1234567@gmail.com

Li et al. developed methodology incorporating direct cooling water and indirect consumption through thermoelectric generation. Ren et al. calculated that GPT-3 training consumed 700,000+ liters. Lawrence Berkeley National Laboratory projects U.S. data center water consumption growing from 17.5 billion gallons (2023) to 35-70 billion gallons (2028). Spatial analysis reveals systematic construction in high water-stress regions.

3) Carbon Emissions and Climate Impact

Carbon Emissions and Climate Impact Lacoste and their team built a calculator to show that the carbon footprint of AI can be five times higher or lower depending on where you plug the computers in and how that local area gets its power. Even though big tech companies keep promising to reach "net-zero" emissions, their own reports show their actual pollution is jumping up by 25% to 50% every year. Furthermore, if you count the pollution caused by manufacturing the hardware and not just running it, the total carbon footprint basically doubles.

4) Environmental Justice

Environmental Justice Robert Bullard created the framework that proves there is a consistent link between who lives in a neighborhood and how much pollution they are forced to deal with. This is now happening with data centers; legal groups have documented cases where these facilities are breaking clean air laws specifically in neighborhoods where most residents are from minority groups.

5) AI for Climate Solutions

AI for Climate Solutions On the other hand, researchers like Rolnick have found 13 different ways that AI could actually speed up climate action. Another study by Vinuesa found that AI could help reach nearly 80% of the UN's sustainability goals. Experts believe this could lead to a massive 1,400 million-ton drop in carbon emissions by 2035, though there are still huge hurdles standing in the way of making that a reality.

METHODOLOGY

1) Research Design

To get the full picture, we used a few different methods. First, we tracked the entire life of the technology—from factory to scrap heap—using the standard international rules for measuring environmental impact. We then used digital mapping to see if data centers are being built in places that are already running out of water or in neighborhoods with specific demographic patterns. We also ran some math to see how different factors influence each other and built a model to figure out how to balance AI's growth with the reality of limited resources. Finally, we took a close look at a few real-world examples to see how these issues play out in the real world.

2) Data Sources

TABLE I: PRIMARY DATA SOURCES

Category	Source	Coverage	Reliability*
Energy	IEA, LBNL	Global, U.S.	9.0/10
Carbon	EPA, Academic	U.S., Model-specific	8.7/10
Water	USGS, Corporate	National, Major tech	8.2/10
Spatial	WRI, Census	Global, National	9.2/10

*Based on methodology rigor, peer review, and data completeness

3) Metrics and Calculations

Energy Consumption:

$E\_train = P\_GPU \times n\_GPU \times t\_train \times PUE$

$E\_infer = E\_query \times N\_queries \times T\_period$

Carbon Emissions:

$C\_total = C\_operational + C\_embodied$

$C\_operational = E\_total \times I\_grid \times (1 - R\_renewable)$

Water Footprint:

$W\_total = W\_direct + W\_indirect$

$W\_indirect = E\_total \times W\_grid$

where W_grid represents water intensity of electricity generation (typical values: coal/gas 1.9 L/kWh, solar/wind 0.03 L/kWh).

Environmental Justice Metrics:

$Exposure\_Index = \Sigma(P\_i \times D\_i) / \Sigma P\_i$

$Disparity\_Ratio = (P\_minority,exposed / P\_minority,total) / (P\_majority,exposed / P\_majority,total)$

4) Validation

Data validation through cross-validation across multiple independent sources, outlier detection using Tukey's method, and consistency checks. Monte Carlo simulation with 10,000 iterations quantifies uncertainty: energy ±15%, carbon ±20%, water ±25%, projections ±30%. Confidence intervals reported at 95% level.

ANALYSIS AND RESULTS

1) Energy Consumption Analysis

Training Phase Requirements:

Analysis of 15 representative LLMs reveals mean training energy of 847 ± 512 MWh. Regression analysis yields: $\log_{10}(E\_train) = 0.89 \log_{10}(N\_params) + 2.34 R^2 = 0.91, p < 0.001$ This indicates approximately linear scaling, with energy growing at N^0.89 relative to parameters.

TABLE II: TRAINING ENERGY FOR SELECTED MODELS

Model	Parameters	Energy (MWh)	CO_2 (MT)	Grid Location
GPT-3	175B	1,287	552	U.S. mixed
BLOOM	176B	433	25	France (nuclear)
PaLM	540B	2,870*	1,231*	U.S. mixed

*Estimated from disclosed compute budgets

The 3× energy difference between GPT-3 and BLOOM despite similar parameter counts demonstrates hardware efficiency and grid carbon intensity impact.

Aggregate Demand:

Global data center electricity consumption reached 575 TWh in 2024, with AI workloads comprising 28% (161 TWh). Projections indicate 945 TWh by 2030 (95% CI: 380-475 TWh for AI portion), representing 9.5% CAGR overall and 17.6% CAGR for AI specifically.

2) Carbon Emissions Assessment

Operational Emissions:

Carbon footprint scales with energy modulated by grid intensity. Regional variation creates 5-10× differences:

- Wyoming: 843 g CO_2/kWh (95% coal/gas)
- Washington: 127 g CO_2/kWh (71% hydro/nuclear)
- Iceland: 18 g CO_2/kWh (100% hydro/geothermal)

Embodied Carbon:

Hardware manufacturing contributes substantial emissions. GPU production requires silicon wafer fabrication (2,000 kWh/wafer), rare earth mining (15-30 kg CO_2/kg), and precious metals (10,000-20,000 kg CO_2/kg). Estimated embodied carbon per GPU: 150-300 kg CO_2.

For hyperscale training:

C_hardware = 10,000 GPUs × 200 kg = 2,000 MT

C_total = 552 MT (operational) + 2,000 MT (embodied) = 2,552 MT

This approximately doubles reported operational emissions.

Corporate Trends:

TABLE III: MAJOR TECH COMPANY EMISSION TRAJECTORIES

Company	2019 (MT)	2024 (MT)	Change	Net-Zero Target
Google	4.9M	7.2M	+47%	2030
Microsoft	11.6M	16.8M	+45%	2030
Meta	1.6M	2.8M	+75%	2030

All companies show increasing emissions contradicting 2030 net-zero commitments, requiring unprecedented 20-30% annual reductions.

3) Water Resource Consumption

U.S. data centers consumed 21.3 billion gallons in 2024 (95% CI: 18.5-24.1). Large facilities consume up to 5 million gallons daily (1.8 billion annually)—equivalent to towns of 10,000-50,000 residents.

Evaporative cooling systems lose 70-80% of withdrawn water permanently through evaporation. Google facilities discharge only 20% to treatment plants, with 80% lost. Critically, nearly 80% comes from potable municipal water supplies, directly competing with residential use.

Indirect Consumption:

Thermoelectric power generation consumes substantial water. U.S. electric power sector withdraws 11.6 gallons per kWh. For 650 TWh annual data center consumption:

W_indirect = 650 × 10^6 MWh × 320 gal/MWh = 208 billion gallons

This indirect footprint exceeds direct consumption by 10×.

Spatial Distribution:

TABLE IV: DATA CENTER SITING BY WATER STRESS LEVEL

Stress Level	WRI Score	Facilities (2022-2025)	% Total
Low	0.0-1.0	35	21%
Medium-High	2.0-3.0	42	25%
High	3.0-4.0	52	31%
Extremely High	4.0-5.0	17	10%
High+ Total	**≥3.0**	**69**	**41%**

Moran's I spatial autocorrelation statistic ($I = 0.34, p < 0.01$) confirms significant clustering in water-stressed areas.

4) Environmental Justice Analysis

Demographic Disparities:

GIS analysis of 247 data center sites reveals systematic inequities:

TABLE V: DEMOGRAPHICS NEAR DATA CENTERS (2-MILE RADIUS)

Metric	Near Facilities	State Average	Ratio	p-value
% Minority	52.3%	38.2%	1.37	<0.001
% African American	22.7%	13.1%	1.73	<0.001
Median Income	$47,300	$64,900	0.73	<0.001
% Below Poverty	18.4%	12.8%	1.44	<0.001

All comparisons statistically significant ($\alpha = 0.01$).

Regression Analysis:

Probit regression examining siting determinants (n = 247 sites + 988 controls):

TABLE VI: SITING REGRESSION RESULTS

Variable	Coefficient	Std. Error	z-value	p-value
Minority %	0.025	0.008	3.13	0.002
Median Income	-0.082	0.021	-3.90	<0.001
Pop. Density	0.0003	0.0001	3.00	0.003
Grid Capacity	0.041	0.009	4.56	<0.001

Pseudo-$R^2 = 0.31$. After controlling for technical factors, demographic variables remain significant predictors.

Pollution Exposure:

EPA AERMOD dispersion modeling estimates increased PM2.5 concentrations of 1.8 μg/m^3 within 2 miles. Health impacts using EPA BenMAP: 120-180 additional asthma cases and 2-4 premature deaths per facility annually, with economic damages of $18-32M/year.

5) The Optimization Paradox

Potential Climate Benefits:

AI applications demonstrate substantial emissions reduction potential:

- Buildings optimization: 262 Mt CO_2/year (10% energy savings)
- Transportation efficiency: 111 Mt/year (6% aggregate improvement)
- Industrial optimization: 67.5 Mt/year (5% efficiency gains)
- Grid optimization: 45 Mt/year (renewable integration)

Total global potential: 1,400 Mt/year (IEA conservative estimate)

Compared to projected AI footprint of 400 Mt/year, this suggests 3.5× benefit-cost ratio.

Critical Limitations:

Monte Carlo simulation (10,000 iterations) incorporating adoption barriers, rebound effects, and implementation timelines yields:

- Median realized benefits (2035): 385 Mt CO_2/year (95% CI: 145-820)
- Probability benefits exceed AI footprint: 48%
- Probability strongly positive outcome (>1,000 Mt): 12%

Mean rebound effect: 22% (95% CI: 10-35%), with transportation showing highest rebound (30%).

Net Impact Assessment:

Current business-as-usual trajectory shows net negative impact through 2032, with crossover point at 2033 ± 2 years. However, atmospheric CO_2 persistence means near-term emissions cause irreversible warming even if eventual benefits materialize.

DISCUSSION

1. Implications for Sustainability

These results prove that we're dealing with what's called "Jevons Paradox." Essentially, even though computers got about four and a half times more efficient between 2020 and 2030, the total amount of energy used actually went up by 42%. It turns out that when you make it cheaper and easier to use AI, people don't save energy—they just build and use more AI. The efficiency gains are basically being swallowed up by the massive growth of the industry.

2. Policy Requirements

Current regulatory vacuum enables unsustainable growth. Critical interventions include:

1. Mandatory Disclosure: Annual reporting of energy by source, water by source/destination, and Scope 1-3 carbon emissions with facility-level granularity
2. Performance Standards: PUE <1.3, WUE <0.5 L/kWh for new facilities with retrofit timelines for existing infrastructure
3. Renewable Mandates: 80% renewable energy by 2030, 95% by 2035
4. Siting Restrictions: Prohibiting construction in regions with WRI water stress ≥3.0
5. Economic Instruments: Carbon pricing at $50/ton (2025) escalating to $150/ton (2030) without sectoral exemptions

3. Technical Solutions

Multi-objective optimization identifies trade-offs requiring balanced approaches:

- Algorithm efficiency: 60% computational reduction potential through architecture optimization
- Hardware advances: 2-5× improvement through neuromorphic/photonic computing
- Cooling transformation: 90% water reduction via immersion/direct-to-chip systems
- Renewable integration: 80-95% clean energy through grid modernization

Representative optimal solution achieves 65% energy reduction, 60% carbon reduction, and 65% water reduction versus baseline but requires $50-80M capital investment and 18-36 month implementation timeline.

FUTURE RESEARCH DIRECTIONS

1) Implications for Sustainability

New Tech on the Horizon We're looking at some big shifts in how computers actually work. "Neuromorphic" chips—which are designed to mimic the human brain—could be 100 to 1,000 times more efficient than what we have now, and we expect them to hit the market between 2028 and 2032. There's also "photonic" computing, which uses light instead of electricity to move data; it's about 5 times better than current chips and should start appearing by the mid-2030s. We also have new ways to cool these machines, like dunking them in special fluids, which could basically stop water waste entirely—though these systems currently cost about 50% more to build.

2) Methodological Advances

Better Ways to Measure Impact To really get a handle on this, we need better tools. Specifically, we need:

Systems that track carbon in real-time so we can schedule big AI tasks for when the sun is shining or the wind is blowing.

Better models for water use that account for how data centers might pollute or degrade the water they do send back.

Fairness scores that look at the "total" impact on a neighborhood over time, not just one building at a time.

Long-term studies to see if making AI more efficient actually saves energy or just encourages people to use it more.

3) Policy Research

Fixing the Rules We need to compare different ways of governing this industry. Should we set "performance goals" (like a miles-per-gallon rating for AI) or just mandate specific technologies? We also need to find a balance where companies are forced to be honest with the public about their resource use without giving away their trade secrets. Finally, we need to figure out how to give local communities a real say in whether a data center gets built in their backyard.

CONCLUSION

What we've found here is that AI's environmental impact is a massive problem that we can't afford to ignore anymore. To put it simply, training just one of these big models can burn through as much energy as a small town uses in a year, and we're on track for data centers to gulp down nearly a trillion kilowatt-hours by 2030. It's the same story with water—we're already at 21 billion gallons a year, and that's likely to quadruple soon. Even worse, most of these new buildings are being put in places that are already running out of water and in neighborhoods where minority residents are far more likely to deal with the fallout.

There is a bit of a twist, though. AI could actually help us save the planet by making things like power grids or shipping routes much more efficient—potentially cutting carbon by way more than the AI itself creates. But don't bank on that happening automatically. Our math shows there's only a 48% chance that AI actually does more good than harm if we just leave it to the "free market." Without rules, we're more likely to see the costs outweigh the benefits.

To fix this, we need to move on five fronts:

- Punt on better tech that uses 60% to 90% fewer resources.
- Pass real laws that force companies to tell the truth about their footprint and meet efficiency standards.
- Hold companies accountable so they actually cut emissions instead of just talking about it.
- Protect communities so that poor and minority neighborhoods don't get stuck with all the pollution.
- Fix the grid so that all this energy comes from wind and sun rather than coal.

The next five years are the "make or break" window. Right now, things are heading in the wrong direction—more pollution, less water, and more inequality. We know it's technically possible to build sustainable AI; that's not the issue. The real question is whether we're going to actually step up and make the hard choices before it's too late. In the end, AI shouldn't be judged by how smart it is, but by whether we were smart enough to use it without destroying our own future.

REFERENCE

[1] E. Strubell et al., "Energy and Policy Considerations for Deep Learning in NLP," in Proc. 57th Annual Meeting ACL, 2019, pp. 3645-3650.

[2] D. Patterson et al., "Carbon Emissions and Large Neural Network Training," arXiv:2104.10350, 2021.

[3] International Energy Agency, "World Energy Outlook 2025," IEA Publications, Paris, 2025.

[4] S. Ren et al., "Making AI Less 'Thirsty': Uncovering and Addressing the Secret Water Footprint of AI Models," arXiv:2304.03271, 2023.

[5] A. Shehabi et al., "United States Data Center Energy and Water Usage Report - 2024 Update," LBNL Technical Report, 2024.

[6] World Resources Institute, "Aqueduct 3.0: Updated Decision-Relevant Global Water Risk Indicators," WRI Technical Note, 2019.

[7] Goldman Sachs Research, "AI and Energy: Data Center Electricity Demand Forecasts," Aug. 2025.

[8] A. Lacoste et al., "Quantifying the Carbon Emissions of Machine Learning," arXiv:1910.09700, 2019.

[9] D. Rolnick et al., "Tackling Climate Change with Machine Learning," ACM Computing Surveys, vol. 55, no. 2, pp. 1-96, 2022.

[10] R. Vinuesa et al., "The role of artificial intelligence in achieving the Sustainable Development Goals," Nature Communications, vol. 11, art. 233, 2020.

[11] E. Masanet et al., "Recalibrating global data center energy-use estimates," Science, vol. 367, no. 6481, pp. 984-986, 2020.

[12] M. Li et al., "Quantifying the energy and carbon footprint of AI computing," Energy and AI, vol. 12, pp. 100235, 2023.

[13] U. Gupta et al., "Chasing Carbon: The Elusive Environmental Footprint of Computing," in Proc. IEEE HPCA, 2021, pp. 854-867.

[14] K. Crawford, "Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence," Yale University Press, 2021.

[15] Southern Environmental Law Center, "Legal Challenge to xAI Memphis Facility for CAA Violations," SELC Press Release, Mar. 2025.

[16] D. E. Taylor, "The Rise of the Environmental Justice Paradigm," American Behavioral Scientist, vol. 43, no. 4, pp. 508-580, 2000.

[17] European Parliament, "Regulation on Artificial Intelligence (AI Act)," Regulation (EU) 2024/1689, 2024.

[18] U.S. Senate, "AI Environmental Impact Disclosure Act of 2025," S. 1847, 119th Congress, 2025.

[19] T. Le Scao et al., "BLOOM: A 176B-Parameter Open-Access Multilingual Language Model," arXiv:2211.05100, 2022.

[20] A. Chowdhery et al., "PaLM: Scaling Language Modeling with Pathways," arXiv:2204.02311, 2022.

[21] U.S. EPA, "Emissions & Generation Resource Integrated Database (eGRID) 2024," EPA, 2024.

[22] M. A. Maupin et al., "Estimated use of water in the United States in 2020," USGS Circular 1441, 2022.

[23] Bloomberg News, "The AI Boom Is Draining Water From Areas That Need It Most," May 2025.

[24] K. Deb et al., "A fast and elitist multiobjective genetic algorithm: NSGA-II," IEEE Trans. Evolutionary Computation, vol. 6, no. 2, pp. 182-197, 2002.

[25] International Organization for Standardization, "ISO 14040:2006 - Environmental Management - Life Cycle Assessment," Geneva, 2006.

[26] Google, "Environmental Report 2024," Google Sustainability, Mountain View, CA, 2024.

[27] Microsoft, "2024 Environmental Sustainability Report," Microsoft CSR, Redmond, WA, 2024.

[28] Meta, "Sustainability Report 2024," Meta Platforms Inc., Menlo Park, CA, 2024.

[29] R. Nanjundappa, Gajendra N, Lakshya et al. "AWAF: AI Enabled Web Contents Authoring Framework", IEEE 17th India Council International Conference (INDICON), 2020, pp. 1-5, doi:10.1109/INDICON49873.2020.9342385. https://ieeexplore.ieee.org/document/9342385.

[30] Lakshya, S. Kota, Mallikarujun Volleti and Shivraj Singh, "Compressed Domain Consistent Motion based Frame Scoring for IoT Edge Surveillance Videos", International Symposium on Visual Computing, 2021, pp. 534-545. https://link.springer.com/chapter/10.1007/978-3-030- 90439-5_42.