



AI's Thirst, AI's Heat, AI's Waste: Exposing the Hidden Environmental Impact of Every Artificial Intelligence Interaction

EDITORIAL

ARAS BOZKURT 



ABSTRACT

Generative Artificial Intelligence (GenAI) is being rapidly adopted by academic and policy leaders, often with a focus on its revolutionary benefits while overlooking its significant, undisclosed environmental costs. This paper deconstructs the physical footprint of GenAI, moving beyond the abstract cloud to analyze its resource-intensive infrastructure. It synthesizes current research into four key areas of concern. First, it reframes the energy debate, demonstrating that the operational inference (use) phase is the dominant long-term cost, with one study estimating its carbon footprint to be 25 times higher than the one-time training cost. Second, it reveals the hidden water footprint, which includes both direct Scope 1 consumption for cooling and the larger Scope 2 consumption from offsite electricity generation, a critical issue as two-thirds of new US data centers are in water-stressed areas. Third, it examines the systemic and material impacts, including the full Life Cycle Assessment (LCA) of hardware (e-waste, raw material extraction) and the rebound effect (Jevons' Paradox), where efficiency gains increase overall demand. A new side effect includes AI tools propagating non-green code. Finally, the paper highlights the critical lack of corporate transparency, a black box of proprietary data that prevents effective governance. This paper argues that a full-cost accounting is necessary and concludes by proposing actionable recommendations from the Green AI and Green Lean movements. Solutions include mandating public reporting, prioritizing smaller models, and implementing sustainable techniques like model pruning and quantization.

CORRESPONDING AUTHOR: Aras Bozkurt

Anadolu University, Türkiye
arasbozkurt@gmail.com

KEYWORDS:

Generative AI; Green AI; AI Environmental Impact; AI Energy Consumption; AI Carbon Footprint; AI Water Footprint; Data Center Sustainability; Sustainable AI; AI Transparency; AI Policy; Higher Education; Teaching and Learning

TO CITE THIS ARTICLE:

Bozkurt, A. (2025). AI's Thirst, AI's Heat, AI's Waste: Exposing the Hidden Environmental Impact of Every Artificial Intelligence Interaction. *Open Praxis*, 17(4), pp. 638–647. DOI: <https://doi.org/10.55982/openpraxis.17.4.1058>

In *The Animatrix*, human leaders launched “*Operation Dark Storm*,” a final act of war to permanently scorch the sky. This was “*a final solution... to deny the machines the energy of the sun, upon which they had come to depend.*”

— The Animatrix: The Second Renaissance, Part I (Maeda, 2003)

What is the environmental CO₂st of digital pO₂liteness and wH₂O pays it? When a user types *please* or *thank you* into generative AI (GenAI) services, it seems like a harmless, humanizing gesture. Yet, every additional word processed by a large language models (LLMs) increases the computational load, which requires energy and costs money (Deb, 2025). OpenAI’s CEO, Sam Altman, recently estimated the cost of this simple courtesy at “tens of millions of dollars” (Altman, 2025). This single, startling fact exposes a much larger, uncomfortable truth: every interaction with GenAI, no matter how trivial, has a physical, resource-intensive cost.

This revelation forces a series of critical questions for academic leaders, policymakers and indeed for everyone. As we rush to integrate GenAI into our classrooms, research labs, and government agencies, are we overlooking its true environmental cost? The dominant narrative of AI as a disembodied, in the cloud technology (Bashir et al., 2024) is a dangerous misconception. The reality is that GenAI is not magic; it is a physical process running in warehouse-scale data centers (Zewe, 2025) that consume massive amounts of energy (Chen, 2025), consume water (Gordon, 2024; Ren, 2023), require the extraction of raw materials (Ergönül, 2025; Zewe, 2025), generate significant electronic waste (Ergönül, 2025; Walther, 2024) and causing other environmental issues (Bozkurt et al., 2024; Bozkurt & Sharma, 2025).

This paper argues that the current discourse on GenAI adoption is dangerously incomplete without a critical assessment of its environmental externalities (Bashir et al., 2024). In this regard, the paper moves beyond the hype to deconstruct the four pillars of AI’s hidden environmental footprint: (1) its massive energy consumption, driven by inference, not just training (Chien et al., 2023); (2) its hidden water footprint for cooling (Ren & Luers, 2025); (3) its total material footprint, from raw mineral extraction to e-waste (Berthelot et al., 2025); and (4) the corporate black box that makes accountability nearly impossible (Chen, 2025). This paper aligns with the principles of Green AI, which includes both Green-by AI (using AI for environmental solutions) and Green-in AI (making AI itself sustainable) (Bolón-Canedo et al., 2024), and Green Lean (Setyadi et al., 2025) to provide a full-cost accounting and propose a sustainable path forward.

THE SCALE OF ENERGY CONSUMPTION

The most common discussion around AI’s energy use is deeply flawed because it focuses on the wrong part of the problem. Public attention, drawn to headline-grabbing figures, has centered on the massive, one-time energy cost of training a model, which is similar to the cost of building a factory. However, this perspective completely misses the far larger, cumulative, and persistent environmental cost of inference, that is, the energy required to run that factory every single day to handle what may be billions of daily user queries.

THE TRAINING ICEBERG TIP

Training is the one-time, energy-intensive process of creating a large-scale model. A study, for example, estimated that training GPT-3 alone consumed 1,287 megawatt-hours (Bolón-Canedo et al., 2024; Tabbakh et al., 2024; Zewe, 2025). One study further noted that training a single AI model can emit as much carbon as five cars over their lifetimes (Hao, 2019). Historically, this focus on training was understandable, as models were viewed as research projects rather than mass-market products (Berthelot et al., 2025).

THE INFERENCE SUBMERGED MASS

Inference is the energy cost of *using* the trained model. It is the *operational expenditure* of energy, repeated every single time a user hits *enter* to handle what may be billions of daily queries (Chen, 2025). The focus on training obscures the dominant, long-term cost.

- Carbon Cost: Researchers from the University of Chicago calculated that for a ChatGPT-like service, one year of inference can produce 25 times the carbon emissions of the model's initial training (Chien et al., 2023).
- Energy Cost: Recent studies suggest that inference, not training, will account for 60% to 90% of a model's total lifecycle energy consumption (Heikkilä, 2023; O'Donnell & Crownhart, 2025). Each query, no matter how small, has a cost.

The most effective Green AI strategy is often to avoid using a massive, all-purpose model for a simple task. This brute-force approach is sometimes called Red AI, buying better results with massive computational power, in contrast to the efficiency of Green AI (Bolón-Canedo et al., 2024).

- Smaller Models: Research shows that using smaller, specialized models tailored to a specific task can reduce energy use by up to 90% (UNESCO, n.d.).
- Knowledge Distillation: The DistilBERT model, for example, was distilled from the larger BERT model. It is 40% smaller and 60% faster while retaining 97% of BERT's performance (Sanh et al., 2019).
- Energy Disparity: The energy costs are not linear. Training the massive GPT-3 model required an estimated 1,290,000 kWh. By contrast, training a high-performing green learning model for a specific task, such as GANSynth, required only 32.4 kWh (Kusummaraju et al., 2024).

As institutions integrate GenAI, indeed, they are not just buying software; they are subscribing to a massive, ongoing, and cumulative energy liability. Opting for smaller, task-specific models over a one-size-fits-all large model is a critical, high-impact sustainability decision. So, the *critical decision* for leaders is to avoid the trap of *bigger is better* and instead select the smallest, most efficient model that is appropriate for the task.

THE HIDDEN WATER FOOTPRINT: COOLING THE DIGITAL BRAIN

The energy consumption of generative AI creates a direct, and often invisible, second-order crisis: extreme heat (Barnett, 2025; Ergönül, 2025). The high-performance Graphics Processing Units (GPUs) that power AI models generate heat so intense that data centers must rely on massive cooling systems, which have become voracious consumers of water (Ren, 2023; Zewe, 2025).

DIRECT VS. INDIRECT WATER USE

AI's water footprint is twofold, and to grasp the full environmental cost, it is crucial to distinguish between the water consumed directly by the data center and the much larger amount of *hidden* water consumed to power it (Barnett, 2025; Ren, 2023; Ren & Luers, 2025).

- Scope 1 (Direct Use): This is the onsite freshwater withdrawn from local sources (like rivers and municipal supplies) and consumed by data center cooling towers (Barnett, 2025; Ren, 2023; Ren & Luers, 2025). The water evaporates to dissipate heat, permanently removing it from the local watershed (Ren & Luers, 2025).
- Scope 2 (Indirect Use): This is the offsite water used by power plants (e.g., fossil fuel or nuclear) to generate the electricity the data center consumes (Barnett, 2025; Ren, 2023; Ren & Luers, 2025).

For most U.S. data centers, the Scope 2 indirect water use accounts for 80% or more of the total water footprint (Ren & Luers, 2025). For example, one study estimated a single GPT-3 query consumed 16.9 ml of water, but only 2.2 ml was for direct cooling (Scope 1) while 14.7 ml (87%) was for the off-site electricity generation (Scope 2) (Ren & Luers, 2025).

Due to a significant lack of corporate transparency, the precise environmental costs of generative AI are often treated as proprietary trade secrets. This opacity has forced researchers to create their own models to calculate the footprint, and the resulting estimates reveal an alarming scale of resource consumption.

The scale of this water consumption is evident at every level: the one-time process of training GPT-3 was estimated to have directly consumed 700,000 liters (185,000 gallons) of clean freshwater (Barnett, 2025; Ren, 2023), while a simple ongoing conversation with ChatGPT (10–50 queries) could *drink* a 500ml (16oz) bottle of water (Barnett, 2025; Ren, 2023). These individual costs scale up to a massive industrial footprint, with U.S. data centers' direct (Scope 1) water consumption estimated at 17.5 billion gallons in 2023 (Ren & Luers, 2025)—a figure projected to double or even quadruple by 2028 (Ren & Luers, 2025). This surge is not theoretical; it is confirmed by corporate reports, such as Microsoft's 34% increase in water consumption from 2021 to 2022, largely attributed to its AI operations (Barnett, 2025; Ren, 2023), and Google's 20% increase in the same period (Ren, 2023).

ENVIRONMENTAL JUSTICE AND THE WATER-ENERGY TRADE-OFF

This water consumption is not happening in a vacuum. It becomes a critical policy and environmental justice issue when *where* this is happening is considered (Ren & Wierman, 2024). Data centers are frequently built in hot, arid regions (Ren, 2023; Zewe, 2025), with an estimated two-thirds of new U.S. data centers located in high water-stress areas (Ren & Luers, 2025). This strategy places the tech industry in direct competition with local communities, Indigenous nations, and agriculture for a diminishing resource (Ren, 2023).

This creates a complex energy-water trade-off. To reduce their *direct* (Scope 1) water use, data centers can use zero water air-cooling systems. However, these systems are less efficient and can increase electricity demand by ~10% (Ren & Luers, 2025). This, in turn, increases the *indirect* (Scope 2) water footprint from power plants, which, as noted, is already 80% of the total burden. In water-stressed regions, operators must navigate tough trade-offs between global climate goals (lower energy) and local water needs (lower direct water) (Ren & Luers, 2025).

THE CARBON, MATERIAL, AND REBOUND EFFECT FOOTPRINT

Beyond the direct consumption of electricity and water, the *source* of that power, the *materials* required to build the hardware, and the *way* we use the technology create a compounding environmental crisis.

THE GRID'S CARBON INTENSITY

Tech companies' net-zero pledges often mask a more complex reality. An AI model's true carbon footprint is determined by the carbon intensity of the specific power grid it relies on (Berthelot et al., 2025; Ding et al., 2025).

Data centers require massive, 24/7, reliable power (O'Donnell & Crownhart, 2025). In many regions, this baseload power is provided by fossil fuels, primarily natural gas (Zewe, 2025). This dependency is critical as global data center electricity demand is projected to soar from 460 TWh in 2022 to 1,050 TWh by 2026 (Zewe, 2025).

These location disparities create significant carbon liabilities. Data center electricity is estimated to be 48% more carbon-intensive than the US average, precisely because they are "clustered in places that have dirtier grids" (O'Donnell & Crownhart, 2025), with the majority of current GenAI models deployed in high-carbon regions like the US and China (Ding et al., 2025). This is starkly illustrated by the fact that while these two countries are responsible for over 99% of global GAI emissions, China has lower energy consumption but higher emissions (contributing 54.4%) due to its grid's higher carbon intensity (Ding et al., 2025). This disparity is even visible within a single company: in 2022, Google's data center in Finland ran on 97% carbon-free energy, while its data centers in Asia ran on only 4–18% carbon-free energy (Ren & Wierman, 2024).

This evidence points to a powerful and practical Green AI strategy which is geographical load balancing. Instead of running a task in the nearest data center, which may be on a 'dirty' grid, this approach involves intentionally shifting training and inference workloads to data centers in regions with low carbon intensity, such as those powered by high renewable output (e.g., Sweden, UK) (Ding et al., 2025; Ren & Wierman, 2024). As the Google case study demonstrates, this single decision can dramatically reduce the carbon footprint of the exact same computational task.

Focusing exclusively on the energy AI consumes during active operation provides a dangerously incomplete picture of its environmental impact. To understand the true environmental toll, a full Life Cycle Assessment (LCA) is necessary, as this method uncovers the massive embodied costs hidden in the hardware's manufacturing, transportation, and eventual disposal (Berthelot et al., 2025; Ergönül, 2025).

Generative AI is driving an exponential demand for specialized, high-performance GPUs and TPUs (Zewe, 2025). This hardware often has a very short shelf-life and becomes obsolete quickly, which contributes to sky-high amounts of e-waste (Ergönül, 2025; Walther, 2024). This hardware also requires the extraction of raw materials like cobalt, lithium, and rare earth metals, often through dirty mining procedures (Zewe, 2025; Ergönül, 2025). The impact is not limited to the data center; a full Life Cycle Assessment (LCA) of a GenAI service (like Stable Diffusion) found that the end-user terminals and networks—the devices we all use—accounted for 85% of the Abiotic Depletion Potential (metal and mineral consumption) and 45% of the carbon footprint (Berthelot et al., 2025).

THE REBOUND EFFECT (JEVONS' PARADOX)

The most insidious threat to sustainable AI is a 19th-century economic theory known as Jevons' Paradox, or the rebound effect (Chen, 2025; Berthelot et al., 2025; Bashir et al., 2024). The paradox states that as a technology becomes more efficient (and thus cheaper to use), demand for it explodes, leading to a net *increase* in the total consumption of that resource (Chen, 2025). This relentless demand means efficiency gains are often used as an argument for limited regulation and instead result in increased adoption without fundamentally considering the vast sustainability implications (Bashir et al., 2024).

THE NON-GREEN CODE SIDE EFFECT

A new and concerning finding is that GenAI tools are actively propagating unsustainable practices. Developers increasingly rely on tools like ChatGPT and GitHub Copilot to write code (Sikand et al., 2024). However, an investigation found these tools exhibit a default non-green behavior (Sikand et al., 2024). For example, when asked to perform a basic file-reading task in Java, they default to energy-inefficient I/O functions instead of more sustainable, efficient alternatives. This means AI is baking in energy-hungry code, scaling the environmental problem at the software level (Sikand et al., 2024).

THE BLACK BOX OF ACCOUNTABILITY: THE TRANSPARENCY PROBLEM

The most significant obstacle to managing the environmental footprint of generative AI isn't a technical limitation; it is, at its core, a political one. This barrier consists of a deliberate and strategic lack of transparency maintained by the small group of corporations that monopolize the development and deployment of large-scale models.

A CULTURE OF CORPORATE SECRECY

The true energy consumption, water usage, and carbon footprints of models like GPT-4, Gemini, or Claude are treated as proprietary trade secrets (Zewe, 2025; O'Donnell & Crownhart, 2025). This lack of data is the single greatest barrier to governance (Ren, 2023). Researchers in the field are vocal about the primary challenge they face: operating with very little detailed data (Koomey, as cited in Chen, 2025). This pervasive lack of information is a source of intense frustration, with experts noting that researchers are going crazy because they cannot get the data they need (de Vries, as cited in Chen, 2025).

MISINFORMATION BY OMISSION

This corporate opacity is a form of misinformation by omission. This is particularly glaring for water consumption. Model cards (fact sheets) that may report carbon data contain almost no information about water (Barnett, 2025; Ren, 2023). One researcher compared this to "excluding

the calorie information from the nutrition facts label of a food product” (Ren, 2023). This forces researchers to use bottom-up estimates on open-source models, which are acknowledged as a *lower bound*, as they cannot access proprietary chips or accurate data on cooling, which is a major energy consumer (Chen, 2025; O'Donnell & Crownhart, 2025).

This black box of accountability, where environmental data is intentionally withheld, makes responsible governance impossible. When leaders and policymakers are forced to operate in the dark, they cannot fulfill their most basic functions: academic leaders are unable to make sustainable procurement decisions, and governments cannot regulate an industry whose impact they are legally prevented from measuring.

- For Academic Leaders: How can a university make a sustainable procurement decision when a vendor refuses to disclose a product's environmental cost?
- For Policymakers: How can a government regulate an industry's environmental impact when it is legally impossible to measure it?

Notably, the EU AI Act (European Union, 2024) does establish obligations for environmental sustainability for high-risk AI systems (Bolón-Canedo et al., 2024). However, similar regulations in the US and China currently lack such specific environmental enforcement, focusing more on security and safety (Bolón-Canedo et al., 2024).

CONCLUSION & RECOMMENDATIONS

The adoption of generative AI is not a simple software upgrade; it is a major institutional investment in a new, resource-intensive industrial infrastructure. The magic of AI is not magic. It is a physical process of computation, heat, water, and carbon, and we must begin to account for it as such. For academic and policy leaders who are committed to sustainability goals, the path forward is not to halt innovation, but to enforce accountability. The current trajectory of rapid, blind adoption is unsustainable.

This paper proposes the following key recommendations based on the principles of Green AI (Tabbakh et al., 2024) and Green Lean operations (Setyadi et al., 2025).

DEMAND TRANSPARENCY BEFORE PROCUREMENT

As leaders of universities, public institutions, and government bodies, you possess immense purchasing power. We must stop signing blank checks. Make any institutional license or adoption of a GenAI tool conditional on the vendor providing a full, audited Environmental Impact Report (EIR). This report must include, at a minimum:

- Energy Consumption per 1,000 queries (inference cost).
- Total Water Footprint (distinguishing Scope 1 direct use and Scope 2 indirect use) (Ren, 2023; Ren & Luers, 2025).
- Material Footprint, including Abiotic Depletion Potential (minerals and metals) (Berthelot et al., 2025).
- Data Center Locations and the Carbon Intensity of their power grids (Ding et al., 2025).

MANDATE PUBLIC REPORTING

Policymakers must move to regulate foundation models as a new class of industrial product. Just as we mandate emissions standards for automobiles and efficiency labels for appliances, we must create a legal framework for mandatory, standardized environmental disclosures for all large-scale AI models (Setyadi et al., 2025; Tabbakh et al., 2024). The EU AI Act, which requires sustainability reporting for high-risk AI, can serve as a global model (Bolón-Canedo et al., 2024).

FUND AND ADOPT GREEN AI STRATEGIES

Leaders must shift the institutional mindset from bigger is better to smarter is better. This involves actively funding and prioritizing known Green AI techniques (Tabbakh et al., 2024).

Model and Algorithm Optimization: This strategy focuses on the software itself, moving away from the brute-force or Red AI mindset. Instead of relying on raw computational power, these Green AI techniques optimize the algorithms to create models that are smaller, faster, and dramatically more energy-efficient. Accordingly, to make AI more efficient, we can prioritize smaller, specialized models for specific tasks instead of using massive, energy-hungry ones. This Green AI approach uses several smart techniques. Knowledge distillation creates a compact student model that learns to mimic a larger teacher model. Model pruning trims away useless parts of a model to make it smaller and faster. Quantization simplifies the model's math by using smaller numbers, which saves memory and energy. Finally, transfer learning saves the most energy by recycling a pre-trained model and just fine-tuning it for a new, similar task, completely avoiding the massive cost of training from scratch.

Sustainable Hardware and Infrastructure: In addition to software-level optimizations, addressing the physical foundation of AI is a critical Green AI strategy. This involves not just what code is run, but where and on what it is run. Accordingly, to make AI's physical infrastructure more sustainable, we can use several key strategies. First, invest in Green Hardware by swapping power-hungry GPUs for more efficient chips like TPUs, FPGAs, or new brain-inspired neuromorphic chips. Second, prioritize efficient data centers that use advanced liquid cooling, optimize their overall power consumption (PUE), and run on renewable energy. Third, use Geographical Load Balancing to strategically shift AI tasks to data centers located in regions with low-carbon (clean) energy grids. Finally, Edge Intelligence avoids the data center entirely for some tasks by running smaller AI models on local devices, which cuts cloud energy costs and improves privacy.

Adopt Green Lean Operational Strategy: Technical fixes alone are insufficient. To create lasting change, these algorithmic and hardware-level solutions must be embedded within an organizational culture that actively and continuously pursues environmental efficiency. This is achieved through the Green Lean framework, which merges the operational efficiency of Lean management with the Green imperatives of sustainability. This framework involves using practical hybrid tools like Eco-Value Stream Mapping (Eco-VSM), which adds energy and emission metrics to process maps, and Green Kaizen, which empowers employees to find continuous environmental improvements.

Adopting AI without this transparency is not only an environmental risk, but also an abdication of fiscal and ethical responsibility. The time to demand accountability is now, before these systems become so deeply embedded that their hidden costs are irreversible.

This blind adoption forces critical questions: When the full environmental bill for this magic comes due, who will be forced to pay? When data centers drain local aquifers, will we blame the secretive corporations or the public leaders who failed to demand transparency? Are we locking our core infrastructure into a physically unsustainable system? If we fail to force this reckoning, future generations will rightly ask: how could we claim to be building a better future when we refused to even ask the price?

TOWARDS A SPECULATIVE FUTURE

The trajectory of this resource-intensive industry, if left unregulated, points toward a future where the black box of corporate secrecy becomes a physical, sovereign infrastructure. As the computational race accelerates, the current model of relying on public grids and competing for public resources, especially fresh water, becomes a critical bottleneck.

It is therefore not just plausible, but strategically probable, that corporations will seek to secede from public infrastructure entirely. We may witness the strategic migration of data centers *into* the sea, using the ocean as a virtually infinite heat sink to solve their cooling crisis, bypassing all competition for freshwater.

To power these aquatic data-fortresses, a reliance on the public grid is a liability. The logical, proprietary solution will be the pursuit of private, next-generation nuclear power, such as corporate-owned Small Modular Reactors.

This is the ultimate end-point of the accountability-free black box: a future of self-cooling, self-powered corporate entities. Physically decoupled from public utilities and ethically unbound by public governance, they would be free to pursue their core capitalist objectives, influencing

DATA ACCESSIBILITY STATEMENT

Data sharing does not apply to this article as no datasets were generated or analyzed during the current study.

SUSTAINABLE DEVELOPMENT GOALS (SDGs)

This study is linked to the following SDG(s): Clean water and sanitation (SDG 6), Affordable and clean energy (SDG 7), Industry, innovation and infrastructure (SDG 9), Responsible consumption and production (SDG 12), Climate action (SDG 13), Life below water (SDG 14), Life on land (SDG 15).

ACKNOWLEDGEMENTS

The author extends a special acknowledgement to Greta Thunberg. Her bravery in speaking uncomfortable truths to power and her unwavering advocacy for climate accountability have inspired a global generation. This paper, which seeks to bring transparency to the hidden environmental costs of technology, is deeply indebted to the urgent public conversation on green future and sust[ai]nability she started.

FUNDING INFORMATION

This paper is funded by Anadolu University with grant number YTS-2024-2559 and YTS-2025-2776.

COMPETING INTERESTS

The author has no competing interests to declare.

AUTHOR CONTRIBUTIONS (CRediT)

Aras Bozkurt: Conceptualization, methodology, formal analysis, investigation, data curation, writing—original draft preparation, writing—review and editing. The author has read and agreed to the published version of the manuscript.

AUTHOR NOTES

Based on *Academic Integrity and Transparency in AI-assisted Research and Specification Framework* (Bozkurt, 2024), the authors of this paper acknowledge that the paper was reviewed, edited, and refined with the assistance of DeepL and Google's Gemini (Versions as of November 2025), complementing the human editorial process. The human authors critically assessed and validated the content to maintain academic rigor. The authors also assessed and addressed potential biases inherent in the AI-generated content. The final version of the paper is the sole responsibility of the human authors.

AUTHOR AFFILIATIONS

Aras Bozkurt  orcid.org/0000-0002-4520-642X
Anadolu University, Türkiye

REFERENCES

Altman, S. [@sama]. (2025, November 10). tens of millions of dollars well spent--you never know [Tweet]. X. <https://x.com/sama/status/1912646035979239430>

- Barnett, J.** (2025, February 25). *The often overlooked water footprint of AI models*. Generative AI in the Newsroom. <https://generative-ai-newsroom.com/the-often-overlooked-water-footprint-of-ai-models-46991e3094b6>
- Bashir, N., Donti, P., Cuff, J., Sroka, S., Ilic, M., Sze, V., Delimitrou, C., & Olivetti, E.** (2024, March 27). *The climate and sustainability implications of generative AI*. An MIT Exploration of Generative AI From Novel Chemicals to Opera. <https://mit-genai.pubpub.org/pub/8ulgrckc/release/2>
- Berthelot, A., Caron, E., Jay, M., & Lefèvre, L.** (2025). Understanding the environmental impact of generative AI services. *Communications of the ACM*, 68(7), 46–53. <https://doi.org/10.1145/3725984>
- Bolón-Canedo, V., Morán-Fernández, L., Cancela, B., & Alonso-Betanzos, A.** (2024). A review of green artificial intelligence: Towards a more sustainable future. *Neurocomputing*, 599, 128096. <https://doi.org/10.1016/j.neucom.2024.128096>
- Bozkurt, A.** (2024). GenAI et al.: Cocreation, Authorship, Ownership, Academic Ethics and Integrity in a Time of Generative AI. *Open Praxis*, 16(1), 1–10. <https://doi.org/10.55982/openpraxis.16.1.654>
- Bozkurt, A., & Sharma, R. C.** (2025). G for Generative is a Gluttonous Fraud and Shit, G for Green is Genuine Intelligence: Exploring the Unspoken Environmental Cost and the Computational Opulence Paradox in the Artificial Intelligence Landscape. *Asian Journal of Distance Education*, 20(2). <https://doi.org/10.5281/zenodo.17665240>
- Bozkurt, A., Xiao, J., Farrow, R., Bai, J. Y. H., Nerantzi, C., Moore, S., Dron, J., Stracke, C. M., Singh, L., Crompton, H., Koutropoulos, A., Terentev, E., Pazurek, A., Nichols, M., Sidorkin, A. M., Costello, E., Watson, S., Mulligan, D., Honeychurch, S., ... & Asino, T. I.** (2024). The manifesto for teaching and learning in a time of generative AI: A critical collective stance to better navigate the future. *Open Praxis*, 16(4), 487–513. <https://doi.org/10.55982/openpraxis.16.4.777>
- Chen, S.** (2025). How much energy will AI really consume? The good, the bad and the unknown. *Nature*, 639(8053), 22–24. <https://doi.org/10.1038/d41586-025-00616-z>
- Chien, A. A., Lin, L., Nguyen, H., Rao, V., Sharma, T., & Wijayawardana, R.** (2023). *Reducing the Carbon Impact of Generative AI Inference* (today and in 2035). Proceedings of the 2nd Workshop on Sustainable Computer Systems, 1–7. <https://doi.org/10.1145/3604930.3605705>
- Deb, S.** (2025, April 24). Saying ‘thank you’ to ChatGPT is costly. But maybe it’s worth the price. *The New York Times*. <https://www.nytimes.com/2025/04/24/technology/chatgpt-alexia-please-thank-you.html>
- Ding, Z., Wang, J., Song, Y., Zheng, X., He, G., Chen, X., Zhang, T., Lee, W.-J., & Song, J.** (2025). Tracking the carbon footprint of global generative artificial intelligence. *The Innovation*, 6(5), 100866. <https://doi.org/10.1016/j.xinn.2025.100866>
- Ergönül, D.** (2025). *The Digital Forest: Exploring AI’s Ecological Footprint and Sustainability Challenges*. Bachelor Thesis, University of Oulu. <https://doi.org/10.5281/zenodo.17189940>
- European Union.** (2024). *The EU AI Act*. European Commission. <https://artificialintelligenceact.eu/the-act/>
- Gordon, C.** (2024, February 25). AI is accelerating the loss of our scarcest natural resource: Water. *Forbes*. <https://www.forbes.com/sites/cindygordon/2024/02/25/ai-is-accelerating-the-loss-of-our-scarcest-natural-resource-water/>
- Hao, K.** (2019). *Training a single AI model can emit as much carbon as five cars in their lifetimes*. MIT technology Review. <https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/>
- Heikkilä, M.** (2023, December 1). *Making an image with generative AI uses as much energy as charging your phone*. MIT Technology Review. <https://www.technologyreview.com/2023/12/01/1084189/making-an-image-with-generative-ai-uses-as-much-energy-as-charging-your-phone/>
- Kusumaraju, S. H., Suneesh, A., Rana, A., Bodicherla, S., & Tyagi, B.** (2024). *EcoGen: Fusing Generative AI and Edge Intelligence for Sustainable Scalability*. <https://vixra.org/pdf/2402.0083v1.pdf>
- Maeda, M.** (Director). (2003). *The Second Renaissance, Part I*. In M. Arias, A. Jones, L. Wachowski, & L. Wachowski (Producers) (Eds.), *The Animatrix* [Film]. Warner Bros. Pictures.
- O’Donnell, J., & Crownhart, C.** (2025, May 20). *We did the math on AI’s energy footprint. Here’s the story you haven’t heard*. MIT Technology Review. <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>
- Ren, S.** (2023, November 30). *How much water does AI consume? The public deserves to know*. *OECD.AI*. <https://oecd.ai/en/wonk/how-much-water-does-ai-consume>
- Ren, S., & Luers, A.** (2025, September 10). *The real story on AI’s water use—and how to tackle it*. *IEEE Spectrum*. <https://spectrum.ieee.org/ai-water-usage>
- Ren, S., & Wierman, A.** (2024, July 15). *The uneven distribution of AI’s environmental impacts*. *Harvard Business Review*. <https://hbr.org/2024/07/the-uneven-distribution-of-ais-environmental-impacts>
- Sanh, V., Debut, L., Chaumond, J., & Wolf, T.** (2019). *DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter*. *arXiv*. <https://doi.org/10.48550/arXiv.1910.01108>
- Setyadi, A., Pawirosumarto, S., & Damaris, A.** (2025). Sustainable operations strategy in the age of climate change: Integrating green lean practices into operational excellence. *Sustainability*, 17(13), 5954. <https://doi.org/10.3390/su17135954>

- Sikand, S., Mehra, R., Sharma, V. S., Kaulgud, V., Podder, S., & Burden, A. P.** (2024). Do Generative AI Tools Ensure Green Code? An Investigative Study. *Proceedings of the 2nd International Workshop on Responsible AI Engineering*, (pp. 52–55). Lisbon, Portugal. <https://doi.org/10.1145/3643691.3648588>
- Tabbakh, A., Al Amin, L., Islam, M., Mahmud, G. M. I., Chowdhury, I. K., & Mukta, M. S. H.** (2024). Towards sustainable AI: A comprehensive framework for Green AI. *Discover Sustainability*, 5(408). <https://doi.org/10.1007/s43621-024-00641-4>
- UNESCO.** (n.d.). *AI Large Language Models: new report shows small changes can reduce energy use by 90%*. United Nations Educational, Scientific and Cultural Organization. <https://www.unesco.org/en/articles/ai-large-language-models-new-report-shows-small-changes-can-reduce-energy-use-90>
- Walther, C. C.** (2024, November 12). Generative AI's impact on climate change: Benefits and costs. *Forbes*. <https://www.forbes.com/sites/corneliawalther/2024/11/12/generative-ais-impact-on-climate-change-benefits-and-costs/>
- Zewe, A.** (2025, January 17). Explained: Generative AI's environmental impact. *MIT News*. <https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117>

Bozkurt
Open Praxis
DOI: 10.55982/
openpraxis.17.4.1058

647

TO CITE THIS ARTICLE:

Bozkurt, A. (2025). AI's Thirst, AI's Heat, AI's Waste: Exposing the Hidden Environmental Impact of Every Artificial Intelligence Interaction. *Open Praxis*, 17(4), pp. 638–647. DOI: <https://doi.org/10.55982/openpraxis.17.4.1058>

Submitted: 08 November 2025

Accepted: 18 November 2025

Published: 25 November 2025

COPYRIGHT:

© 2025 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See <http://creativecommons.org/licenses/by/4.0/>.

Open Praxis is a peer-reviewed open access journal published by International Council for Open and Distance Education.