

The Hidden Footprint of AI

Climate, Water, and Justice Costs

Created by Hastings Initiative Members

Last Updated 7/21/2025

Summary

AI tools are becoming a daily fixture of academic and public life. But their environmental footprint – carbon emissions, water usage, and global labor inequality – remains largely invisible. This brief draws from recent research to reveal how AI systems extract energy, water, minerals, and labor from vulnerable communities. It calls on students, faculty, and institutions to use AI more thoughtfully and advocate for transparency, sustainability, and justice in how these tools are developed and deployed.

Introduction

Artificial Intelligence (AI) is often framed as a transformative force, from personalized tutoring and medical diagnostics to scientific discovery. But behind these promises lies a vast, often invisible infrastructure that consumes enormous amounts of electricity, water, and physical hardware. Far from immaterial, AI systems are powered by data centers running on fossil-fueled grids, cooled by freshwater, and built from rare-earth minerals mined under extractive conditions.

These environmental and social costs are not evenly shared. Communities in the Global South and marginalized areas of the Global North often bear the brunt of carbon emissions, water withdrawals, e-waste, and hidden labor that make AI possible (without sharing in its benefits).

As Kate Crawford writes in *The Atlas of AI*, artificial intelligence is not just code or algorithms; it is a planetary industrial system built on layers of extraction: of energy, labor, minerals, and data (2021).

This brief offers a window into AI's hidden footprint. Drawing on recent academic research and investigative journalism, it examines three key concerns: carbon emissions and energy use, water consumption, and global inequalities. Its goal is to help our campus community understand the full cost of AI technologies and to spark conversation about how students, faculty, and institutions might respond with care, equity, and accountability.

Carbon Emissions and Energy Use

AI's rising energy demands are often illustrated by the staggering costs of training large models like GPT-3 or GPT-4. But while training garners headlines, the real long-term burden comes from inference, the everyday use of these systems by billions of users.

According to *MIT Technology Review*, inference now accounts for 80-90% of total AI computing power, and that number is climbing. By 2028, AI alone could consume as much electricity as 22% of U.S. households, a staggering figure given technology's growth (O'Donnell & Crownhart, 2025, p. 4). Even a single ChatGPT prompt uses over 1,000 joules. With an estimated one billion queries daily, that adds up to 109 gigawatt-hours annually, or enough to power more than 10,000 homes (p. 24).

These energy demands are met by sprawling data centers, many of which run on fossil fuel-heavy grids. In Virginia, where electricity is primarily coal-based, AI's carbon footprint is especially high. Data center electricity there is 48% more carbon-intensive than the national

average (p. 21). As Kate Crawford writes, “AI systems are not in the cloud – they are on Earth, in data centers, run on coal” (*The Atlas of AI*, 2021, p. 28).

Exacerbating this impact is a profound lack of transparency. Most tech companies don’t disclose the energy or emissions associated with their AI tools. As O’Donnell and Crownhart put it, “Most firms serve up a total black box” (p. 12), leaving when it comes to energy emissions researchers, regulators, and the public unable to verify environmental claims.

While inference now dominates AI’s emissions profile, training large models still carries a significant carbon cost. Bender et al. (2021) estimate that training a single large transformer model emitted over 284 metric tons of CO₂, which is about five decades of emissions from an average person (p. 612).

And the trend is accelerating. A 2024 Capgemini Invent study projects that under a high-growth scenario, AI’s electricity use could increase 24.4-fold by 2030. Generative AI (GenAI) models are especially power-hungry. A single inference from Meta’s LLaMA 405B model can consume 17 watt-hours, roughly the energy needed to toast a slice of bread (Desroches et al., 2024, p. 14). Though GenAI accounts for less than a third of corporate AI use cases, it drives 99.9% of total energy consumption (p. 18).

AI also depends on vast amounts of physical hardware: GPUs, custom chips, and server clusters built from rare earth minerals. The emissions from mining and manufacturing (especially cobalt and lithium) are rarely counted in carbon reports. As Crawford (2021) notes, the AI supply chain begins in mines and ends in landfills. These extractive processes disproportionately impact the Global South, where environmental protection is often weaker and labor conditions more exploitative. Frequent hardware upgrades also contribute to a growing e-waste crisis, with much of it exported to countries unequipped to safely process toxic components.

AI's energy infrastructure is vast, accelerating, and deeply unequal. Any serious conversation about sustainable technology must reckon with this hidden footprint.

Water Use and Resource Scarcity

While AI's carbon emissions have drawing increasing attention, its water footprint remains largely hidden. From cooling servers to powering data centers to manufacturing chips, AI consumes enormous volumes of water. Oftentimes it does so in regions already facing drought and water stress. Researchers typically classify AI's water use into three categories:

1. **Scope-1:** On-site cooling to prevent servers from overheating
2. **Scope-2:** Off-site electricity generation, particularly from water-cooled power plants
3. **Scope-3:** Supply chain production, including chip and infrastructure manufacturing

Together, these sources create a massive and geographically uneven demand for freshwater, with urgent implications for environmental justice.

In one striking example, Li et al. (2025) found that training GPT-3 at a Microsoft's U.S. data center consumed more than 5.4 million liters of water, accounting for both cooling and electricity-related use. Scaled globally, the impact is staggering. The authors project that by 2027, AI could be responsible for 4.2 to 6.6 billion cubic meters of water withdrawals, which is more than the entire annual water use of Denmark or half the United Kingdom (p. 1).

Even everyday AI use adds up. A single GPT-3 query can “drink” a 500 mL water bottle every 10 to 50 prompts, depending on where it's hosted. In Arizona, a drought-prone state, each request uses about 30 mL of water, compared to 7.6 mL in Texas (Li et al., 2025, Table 1, p. 5).

These disparities highlight a deeper equity issue: location matters. Siddik et al. (2021) show that while only 20% of U.S. data centers are located in the arid West and Southwest, these

regions account for 70% of the industry's water scarcity-related impact (p. 7). In places like Arizona, water use per kilowatt-hour can be nine times higher than in cooler climates (Li et al., 2024). By placing massive water-cooled infrastructure in drought-prone areas, tech companies shift the environmental burdens onto communities already facing water insecurity.

As Kate Crawford writes in *The Atlas of AI*, "What looks like machine intelligence is often just water, power, and labor in disguise" (paraphrased from Crawford, 2021, p. 8). The placement of AI infrastructure, whether near overdrawn hydroelectric dams or in dry regions, reveals how environmental risks are unevenly distributed in service of digital convenience.

Ironically, efforts to reduce carbon emissions can increase water use. "Follow the sun" scheduling, which shifts workloads to solar-powered data centers, can raise water demand due to higher temperatures and evaporation rates. As Li et al. caution, "minimizing one footprint might increase the other...water and carbon footprints are not substitutable" (2025, p. 6). This tradeoff underscores the need for holistic sustainability metrics that address multiple resource impacts, not just carbon.

Yet despite these realities, most AI developers do not disclose water use. While some model cards include carbon estimates, water consumption (especially Scope-2 and Scope-3) is almost never reported. Without comprehensive tracking, even engineers struggle to assess the systems they build. And without public reporting, communities have no way to understand the true cost of the technologies reshaping their local water systems.

Geographic and Structural Inequities

AI's environmental impacts are not only massive, but they are also deeply uneven. From carbon-intensive data centers in coal-dependent regions of the United States to water-stressed

facilities in the Global South, the burdens of AI's infrastructure fall hardest on communities with the least power to resist them or benefit from them. These patterns echo long-standing legacies of environmental injustice, resource extraction, and digital colonialism.

Towards Environmentally Equitable AI via Geographical Load Balancing, a study on environmentally equitable AI deployment (Li et al., 2024), reveals how optimization strategies (e.g. routing AI tasks to regions with the cheapest electricity) can unintentionally magnify harm. These systems often ignore local water scarcity or grid pollution. For example, in peak summer months, data centers in Arizona may consume up to 9 liters of water per kilowatt-hour for cooling. This is nine times more than equivalent centers in cooler climates (Li et al., 2024, pg. 3).

The same study highlights stark carbon disparities. In 2020, Google's Finland data center ran on 94% carbon-free energy, while its Singapore facility operated with just 4%. That's a 23-fold difference in emissions per unit of computation (pg. 2). Systems optimized for global efficiency can obscure these localized costs, concentrating environmental and health harms in specific communities.

But these inequities extend beyond water and carbon. In her 2025 paper *Digital Colonialism*, Samavia Zia argues that AI development replicated global power imbalances rooted in colonial history. Corporations based in the Global North extract massive amounts of personal data from the Global South (often without meaningful consent or compensation) and turn it into profit. Just as with past extractive industries, value flows out, and harm stays behind.

Zia also exposes the hidden labor force behind AI. In countries like the Philippines, India, and Kenya, underpaid workers perform critical tasks: labeling training data, flagging toxic content, and moderating platform abuse. This so-called automation still depends on human labor, but it's labor that remains hidden and undervalued. At the same time, AI surveillance tools are

disproportionately deployed in postcolonial and marginalized communities, deepening existing structures of control.

These realities mirror what Kate Crawford describes as the “new geographies of AI,” a global infrastructure that reproduces imperial dynamics under the banner of technological innovation. AI, she writes, is not abstract but “extractive, unequal, and deeply material” (Crawford, 2021, p. 39).

Taken together, these perspectives challenge conventional ideas of sustainability. A system may appear “green” on average while externalizing serious harm to specific people and places. A truly ethical and sustainable AI must not just account for watts, liters and carbon, but also for geography, labor, history, and justice.

Gaps in Transparency and Accountability

Despite AIs vast and growing environmental footprint, there is surprisingly little public visibility into how much energy, water, or material these systems actually consume. The lack of standardized environmental reporting limits public understanding, stalls scientific research, and hinders both corporate accountability and policymaking. In short, the transparency gap is one of the biggest barriers to meaningful climate action in the AI industry.

As the *MIT Technology Review*’s reports, most AI companies do not disclose basic data like electricity use, carbon emissions, or water consumption (O’Donnell & Crownhart, 2025, p. 12). Without this information, researchers, regulators, and clients are left in the dark, unable to verify sustainability claims or make informed decisions.

This secrecy is not accidental. As Kate Crawford argues in *The Atlas of AI*, the industry operates under a “facade of abstraction,” where glossy narratives about models and innovation

conceal the material realities beneath. She writes, “Artificial intelligence is both embodied and material, made from natural resources, fuel, human labor, infrastructures, logistics, histories, and classification” (2021, p. 8). The phrase “artificial intelligence” itself masks the global networks of mining, manufacturing, waste, and labor extraction that make AI possible. “AI extracts far more from us and the planet than is widely known” (Crawford, p. 32).

Some developers have called for greater transparency, but current tools fall short. Most large language models include technical specs (e.g. training time, dataset size) but rarely disclose environmental costs. Water usage is almost entirely never reported in technical papers, model cards, and cloud dashboards. As Li et al. (2025) observe, this omission “makes it difficult for even engineers to assess the true resource costs of the systems they build” (p. 1).

To address this, Li and colleagues call for full-scope environmental accounting, including scope-1(on-site), scope-2(electricity-related), and scope-3 (hardware and supply chain) water usage. This mirrors the more mature carbon accounting frameworks now common in climate science. Without such tracking, sustainability efforts will remain reactive and incomplete.

Some institutions are proposing more concrete tools. The *Capgemini Invent* report (Desroches et al., 2024) recommends the standardized environmental “eco-labels” that rate AI models by energy use per task (e.g. chat, image generation, RAG). These metrics could empower organizations to choose lower-impact tools and incentivize more efficiency model design.

But without enforced disclosure requirements, AI companies continue to expand operations often with public subsidies and limited oversight. In Virginia, for instance, residents may face monthly utility hikes of up to \$37.50 to cover rising data center energy needs (O’Donnell & Crownhart, 2025, p. 30). If the public is expected to foot the bill, they deserve access to the environmental facts.

Transparency is not a luxury, it is a baseline requirement for climate accountability, sustainability, and democratic oversight in the age of AI. But as Crawford reminds us, transparency alone is not enough. We must not only ask what is being optimized, but also for whom and at whose expense.

Conclusion

Artificial intelligence is reshaping the world through both its capabilities and the infrastructure it requires. As this brief has shown, AI systems depend on immense amounts of energy, water, hardware, and labor. These costs are often hidden from public view, yet they fall heaviest on communities already facing climate stress, water scarcity, and systemic inequality.

The environmental impacts of AI are not evenly shared. Data centers are sited in water-stressed regions. Power comes from fossil-fueled grids. Supply chains rely on rare-earth minerals mined under extractive conditions. The benefits of AI tend to concentrate in wealthy institutions and corporations, while the burdens (carbon emissions, water use, e-waste, and invisible labor) are pushed onto the Global South and marginalized communities in the Global North.

As a college community, we may not be able to redirect the entire trajectory of AI's development. We can demand greater transparency from tech companies. We can advocate for sustainable research practices and more ethical standards for digital infrastructure. More importantly, we can refuse the myth that AI is somehow immaterial or neutral. We can understand it as a powerful system whose impacts on water, carbon, and justice must be confronted and reimagined.

Recommendations: Acting with Care in an Age of AI

No single college can reverse the planetary-scale systems that underlie artificial intelligence. But institutions of higher education do shape how technologies are adopted, discussed, and taught. As students, educators, and researchers, we have the opportunity and responsibility to use AI thoughtfully, ask better questions, and advocate for more just technological futures. Below are a few ways our community can start:

1. **Use AI with awareness.** Every query carries hidden environmental and social costs. Just as we learn to conserve electricity or water, we can be mindful of our digital consumption. Avoid unnecessary or repetitive AI prompts, especially for image and video generation (energy-intensive). Treat AI as a tool, not a default for all writing or thinking. Consider the labor and resources behind the screen. Whose data, water and energy make these tools possible.
2. **Include environmental ethics in AI education.** Including materials on AI's environmental footprint and global inequalities. Teaching sustainable computing practices, such as efficient model design or greener deployment.
3. **Reframe how we talk about intelligence.** We can challenge the idea that artificial intelligence is abstract, weightless, or inevitable. We can ask the important questions in our dialogues. Who builds it? Who benefits? Who bears the cost? What kinds of intelligence do we overlook when we overvalue automation? How might we build systems grounded not just in performance, but in equity, sustainability, and care?

References

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). *On the dangers of stochastic parrots: Can language models be too big?* In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, (pp. 610–623).

<https://doi.org/10.1145/3442188.3445922>

- A foundational critique of large-scale language models, highlighting their environmental impact and the ethical harms of excluding marginalized groups.

Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

- Reframes AI as a global industrial system built on extraction, from energy and minerals to labor and data, rather than as a purely digital technology.

Desroches, C., Chauvin, M., Ladan, L., Vateau, C., Gosset, S., & Cordier, P. (2024). *Exploring the sustainable scaling of AI dilemma: A projective study of corporations' AI environmental impacts*. Capgemini Invent.

<https://www.capgemini.com/insights/research-library/sustainable-gen-ai/>

- An industry report projecting exponential growth in AI's electricity demands and recommending energy intensity scores for generative models.

Li, P., Yang, J., Islam, M. A., & Ren, S. (2025). *Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models*. arXiv:2304.03271.

<https://doi.org/10.48550/arXiv.2304.03271>

- Breaks new ground by quantifying the water use of GPT-3 across regions and proposing sustainability tradeoff frameworks for AI deployment.

Li, P., Yang, J., Wierman, A., & Ren, S. (2024). *Towards environmentally equitable AI via geographical load balancing*. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems* (pp. 291–307).

<https://doi.org/10.1145/3632775.3661938>

- Explores how current AI workload balancing practices shift environmental harm to vulnerable areas and proposes more equitable alternatives.

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/>

- Investigative report quantifying AI's growing carbon footprint and the lack of transparency around emissions and electricity use.

Siddik, M. A. B., Shehabi, A., & Marston, L. (2021). *The environmental footprint of data centers in the United States*. *Environmental Research Letters*, 16(6), 064017.

<https://doi.org/10.1088/1748-9326/abfbal>

- Finds that water-stressed Western U.S. regions, though home to a minority of data centers, bears most of the water-related environmental burden.

Zia, S. (2025). *Digital Colonialism: Reimagining Power, Identity, and Resistance by Decolonizing AI*. Yayasan Drestanta Pelita Indonesia, 129–147. Retrieved from

<https://publisher.yayasandpi.or.id/index.php/dpipress/article/view/1722>

- A critical perspective on how AI development mirrors colonial power structures through data mining, invisible labor, and surveillance.