

Beyond Climate Change:

A Holistic Framework for Evaluating the Environmental Impact of Computing Systems

by Yi Ding, Yanran Wu, Tianyao Shi, and Inez Hua

An overview of three recent studies designed to evaluate large-scale computing's impacts on climate change, water, and biodiversity.

The rapid expansion of artificial intelligence (AI) is reshaping global infrastructure through significant investment in computing hardware and software. For example, Nvidia's \$100 billion investment in OpenAI aims to sustain the AI boom and drive demand for its chips.¹ Amazon is building a 1,200 acre, 2.2 gigawatt (GW) AI data center complex in Indiana to support Anthropic, equipped with hundreds of thousands of Trainium 2 chips. This supersized infrastructure reflects a broader industry trend, with companies like Meta and OpenAI also building GW scale data centers to meet the computational demands of advanced AI models.²

Yet, these developments come with escalating environmental impacts. The U.S. Department of Energy projects that data centers could consume up to 12% of national electricity by 2028, with large language models (LLMs) driving much of this growth.³ Studies show that larger AI models emit increasingly more climate changing gases per query,⁴ underscoring the urgent need for sustainability frameworks that address AI's environmental impacts.

Consider the 2.2 GW AI data center in Indiana: its carbon emission stems from both operational energy use (electricity for compute, cooling, and networking) and embodied emissions from manufacturing chips and construction materials. Reducing carbon through renewable energy or efficiency gains may, however, may inadvertently increase water withdrawals for cooling or hydropower. Likewise, land use change and infrastructure expansion can disrupt local ecosystems, contributing to biodiversity loss. Hence, any sustainable computing strategy must jointly assess these dimensions rather than treat them in isolation.

This article presents an overview of three recent studies: FUEL,⁵ SCARF,⁶ and FABRIC,⁷ that collectively enable holistic evaluation of computing's impacts on climate change, water, and biodiversity. These frameworks offer a multi-layered lens for environmental managers, policymakers, and technologists seeking to mitigate the coupled environmental impact of climate change, water scarcity and biodiversity loss at the planetary scale.⁸

Carbon Efficiency: The FUEL Framework

Carbon emissions remain a central concern in computing sustainability. The FUEL (Functional Unit-based Evaluation for LLMs) framework introduces Functional Units (FUs)—tokens generated under specific performance and quality constraints—as a standardized basis for comparing carbon emissions across system configurations (see Figure 1).

Components of Carbon Emissions


- **Operational Carbon:** Energy consumed during inference, multiplied by grid carbon intensity.
- **Embodied Carbon:** Manufacturing and packaging emissions amortized over hardware lifetime.

Case Studies

We conducted three case studies that systematically compared LLM serving systems across three dimensions: model size, quantization strategy, and hardware choice. Each case study offers actionable insights into designing more sustainable AI infrastructures. For example, in evaluating hardware choices, we observed that newer accelerators such as the Nvidia H100 GPU can deliver lower latency,⁹ yet these gains do not automatically translate into lower carbon emissions. The embodied carbon associated with manufacturing advanced GPUs and their elevated operational power draw can offset efficiency improvements. By contrast, when output quality and latency requirements are satisfied, deploying LLMs on older GPUs such as the Nvidia L40¹⁰ can yield lower carbon emissions per functional unit. Our finding highlights an important design trade-off: maximizing performance may not align with minimizing total carbon emission.

Water Stress and Computing: The SCARF Framework

Water consumption in computing is often overlooked, yet it is a critical sustainability metric. Data centers require large amounts of water for cooling, and semiconductor fabrication plants use ultra-pure water in manufacturing processes. The SCARF (Stress-Corrected Assessment of Water Resource Footprint) framework addresses this gap by introducing a



spatially and temporally aware method for evaluating water impact.

Key Innovations

SCARF calculates an Adjusted Water Impact (AWI) by multiplying raw water consumption with a Water Stress Factor (WSF), which accounts for both geographic location and temporal (including seasonal) water availability. This approach enables more accurate assessments of environmental burden, especially in regions with fluctuating water stress (see Figure 2).

Case Studies

We conducted three case studies to evaluate the environmental impact of water consumption for LLM serving, data centers, and semiconductor fabs. Each case study highlights a distinct dimension of water impact, from model-level efficiency to site-level resource management and upstream manufacturing dependencies, illustrating how water sustainability challenges propagate across the computing stack. In our analysis of Google's U.S. data centers,¹¹ we quantified their annual water consumption and local water stress exposure using geospatial and operational datasets. We found that most of Google's facilities are strategically located in

low- to medium-water-stress regions, suggesting that siting decisions already consider hydrological constraints. However, SCARF further reveals that neither total water consumption nor regional water stress alone provides an accurate picture of water impact. In some cases, high water consumption in a medium stressed region can impose greater long-term impacts than moderate consumption in a high stressed region.



Biodiversity Loss: The FABRIC Framework

Biodiversity is a critical planetary boundary, yet its connection to computing has remained largely unexplored. The FABRIC (Fabrication-to-Grave Biodiversity Impact Calculator) framework introduces two new metrics—Embodied Biodiversity Index (EBI) and Operational Biodiversity Index (OBI)—to quantify biodiversity impact across the lifecycle of computing systems (see Figure 3).

Midpoint and Endpoint Metrics

FABRIC uses life cycle assessment (LCA) principles to evaluate three midpoint impacts:

- **Acidification Potential (AP):** Emissions like sulfur dioxide (SO₂) and nitrogen oxides (NO_x) lower the pH of precipitation and surface water.



- **Eutrophication Potential (EP):** Excess nutrients lead to oxygen-depleting algal blooms.
- **Freshwater Ecotoxicity Potential (FETP):** Toxic chemicals harm freshwater aquatic ecosystems.

These are converted into endpoint metrics expressed in species • yr, representing the statistical loss of species per year.

Key Findings

We evaluated FABRIC on seven high performance computing (HPC) workloads across three computing platforms with diverse hardware infrastructure. We discovered:

- **Manufacturing Dominates:** Manufacturing contributes 78–92% of midpoint impacts and 55–75% of endpoint biodiversity loss.
- **System-Level Analysis:** Perlmutter, a petascale supercomputer, has an annual OBI of 2.51×10^{-3} species•yr, which ~60 greater than its annualized EBI.
- **Energy Mix Matters:** Relocating Perlmutter to hydro-rich Québec reduces its OBI by two orders of magnitude compared to fossil-heavy grids.

Implications

FABRIC highlights the need to consider biodiversity alongside carbon and water in sustainable computing. It also underscores the importance of clean energy and efficient hardware reuse in minimizing ecological damage.

Integrated Insights for Environmental Management

The FUEL, SCARF, and FABRIC frameworks collectively offer a multi-dimensional approach to evaluating the environmental impact of computing systems. Key takeaways include:


1. **Context Matters:** Environmental impact varies significantly by location, time, and workload intensity. Spatial and temporal awareness is essential.

2. **Lifecycle Thinking:** Embodied impacts from manufacturing and end-of-life stages must be considered alongside operational emissions.
3. **Beyond Climate Change:** Water stress and biodiversity loss are equally critical and often more regionally sensitive than carbon emissions.
4. **Policy Levers:** Discount rates, energy mix, and infrastructure siting policies can dramatically alter sustainability outcomes.
5. **Technology Choices:** Quantization, model size, and hardware reuse offer practical pathways to reduce environmental burden.

Recommendations for Practitioners

Environmental managers, sustainability officers, and system architects in the computing sector can leverage the FUEL (carbon), SCARF (water), and FABRIC (biodiversity) frameworks in tandem to guide more comprehensive sustainability strategies. When integrated, they enable multi-impact decision-making across the software–hardware–infrastructure lifecycle. To translate these insights into practice:

- **Adopt Multi-Metric Evaluation:** Use AWI, EBI, and OBI to assess environmental impact holistically.
- **Incorporate Regional Data:** Leverage tools like Aqueduct 4.0^{12,13} for water risk and eGRID¹⁴ for regional carbon intensity to inform site selection and workload scheduling.
- **Promote Hardware Longevity:** Extend the lifecycle of computing devices to reduce embodied carbon and biodiversity impact.
- **Advocate for Clean Energy:** Transition to renewable-heavy grids to minimize operational emissions and ecological harm.
- **Support Open Data and Standards:** Encourage transparency in sustainability reporting to enable robust environmental assessments.



Be sure to visit www.awma.org regularly for the latest important information from A&WMA.

AIR & WASTE MANAGEMENT ASSOCIATION

Expand your knowledge

A&WMA offers many opportunities to advance your knowledge and career. Take a course, or a Webinar. Attend ACE or one of our targeted Specialty Conferences. Network at a Section or Chapter meeting.

Discover what you can learn through A&WMA involvement.

Information **Knowledge** **Networking** **Solutions**

Conclusion

As computing systems become integral to modern life, their environmental footprint must be managed with the same rigor applied to traditional industrial sectors. The FUEL, SCARF, and FABRIC frameworks represent a new frontier in

environmental management—one that integrates considerations about climate change, water, and biodiversity into a unified sustainability strategy. By adopting these tools and insights, stakeholders can build a more resilient, equitable, and ecologically sound digital future. **em**

Yi Ding is an Assistant Professor in the Elmore School of Electrical and Computer Engineering (ECE) and **Inez Hua** is a Professor in the School of Sustainability Engineering and Environmental Engineering and the Lyles School of Civil and Construction Engineering at Purdue University, West Lafayette, IN. Dr. Hua is also a sustainability advisor to the technology company Western Digital. **Yanran Wu** is a Ph.D. student in Computer Science and **Tianyao Shi** is a Ph.D. student in ECE. All authors were directly involved in the research efforts reported in this article.

References

1. How Nvidia Is Backstopping America's AI Boom; *WSJ*; [https://www.nytimes.com/2025/06/24/technology/amazon-ai-data-centers.html](https://www.wsj.com/tech/ai/how-nvidia-is-backstopping-americas-ai-boom-875cl3467gaa_at=eafs&gaa_n=ASWzDAg7D6Q05NEzZptDAAfklobP1dxDS7cXLPY5ni42KWLxWkT1J3vGSDKC4rPsw8963D&gaa_ts=68d5c7b3&gaa_sig=hqGeKOKNWVq4o6bO2Yqtspl3NGwLW0xUeLgpZJ8Qap_n7JpFfMiFMWrgfahC4V6jCbxnfVj4VlzfIzDoGm0A%3D%3D (accessed 2025-09-24).
2. At Amazon's Biggest Data Center, Everything Is Supersized for A.I.; <i>The New York Times</i>; <a href=) (accessed 2025-09-24).
3. Shehabi, A.; Newkirk, A.; Smith, S. J.; Hubbard, A.; Lei, N.; Siddik, M.A.B.; Holecek, B.; Koomey, J.; Masanet, E.; Sartor, D. Lawrence Berkeley National Laboratory LBL Publications Title 2024 United States Data Center Energy Usage Report Permalink <https://escholarship.org/uc/item/32d6m0d1> Publication Date. 2024; <https://doi.org/10.71468/P1WC7Q>.
4. Your A.I. Queries Come With a Climate Cost; *The New York Times*; https://www.nytimes.com/2025/06/19/climate/ai-emissions-chatbot-accuracy.html?unlocked_article_code=1.ak8.5Yoa.lkMybe21ulzk&smid=url-share (accessed 2025-09-24).
5. Wu, Y.; Hua, I.; Ding, Y. Unveiling Environmental Impacts of Large Language Model Serving: A Functional Unit View. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; Association for Computational Linguistics: Vienna, Austria, 2025; pp 10560-10576.
6. Wu, Y.; Hua, I.; Ding, Y. Not All Water Consumption Is Equal: A Water Stress Weighted Metric for Sustainable Computing; *ACM SIGENERGY Energy Informatics Review* 2025, 5 (2), 84-90.
7. Shi, T.; Kumar, R.; Hua, I.; Ding, Y. When Servers Meet Species: A Fab-to-Grave Lens on Computing's Biodiversity Impact; *ACM SIGENERGY Energy Informatics Review* 2025, 5 (2), 34-40.
8. Richardson, K.; Steffen, W.; Lucht, W.; Bendtsen, J.; Cornell, S. E.; Donges, J. F.; Drücke, M.; Fetzer, I.; Bala, G.; Von Bloh, W.; Feulner, G.; Fiedler, S.; Gerten, D.; Gleeson, T.; Hofmann, M.; Huiskamp, W.; Kummel, M.; Mohan, C.; Nogués-Bravo, D.; Petri, S.; Porkka, M.; Rahmstorf, S.; Schaphoff, S.; Thonicke, K.; Tobian, A.; Virkki, V.; Wang-Erlandsson, L.; Weber, L.; Rockström, J. Earth beyond Six of Nine Planetary Boundaries, *Science*, 2023; <https://www.science.org>.
9. H100 Tensor Core GPU, NVIDIA. See <https://www.nvidia.com/en-us/data-center/h100/> (accessed 2025-10-19).
10. NVIDIA L40 GPU for Data Center, NVIDIA. See <https://www.nvidia.com/en-us/data-center/l40/> (accessed 2025-10-19).
11. Google Data Centers. See <https://datacenters.google/> (accessed 2025-10-19).
12. Kuzma, S.; Bierkens, M.F.P.; Lakshman, S.; Luo, T.; Saccoccia, L.; Sutanudja, E.H.; Beek, R. Van. Aqueduct 4.0: Updated Decision-Relevant Global Water Risk Indicators; World Resources Institute, 2023; <https://doi.org/10.46830/WRITN.23.00061>.
13. Aqueduct 4.0: Updated Decision-Relevant Global Water Risk Indicators; World Resources Institute; <https://www.wri.org/research/aqueduct-40-updated-decision-relevant-global-water-risk-indicators> (accessed 2025-10-19).
14. Emissions & Generation Resource Integrated Database (eGRID); U.S. Environmental Protection Agency; <https://www.epa.gov/egrid> (accessed 2025-10-19).



A&WMA Buyers Guide

Tap into the incredible network of the Air & Waste Management Association with the A&WMA Buyers Guide. Powered by MultiView, the Guide is the premier search tool for environmental professionals. Find the suppliers you need, within the network of the association you trust.

Start your search today at awma.org.



AIR & WASTE MANAGEMENT
ASSOCIATION



MULTIVIEW