"Is the AI Water Issue Fake:" Discourse of GenAI's Environmental Footprint

Nino Migineishvili*
ninom@cs.washington.edu
University of Washington
Seattle, Washington, USA

Amelia Lee Doğan*
dogan@uw.edu
University of Washington
Seattle, Washington, USA

Anoolia Gakhokidze
annygakh@uw.edu
University of Washington
Seattle, Washington, USA

Chandana Robba
crobba@cs.washington.edu
University of Washington
Seattle, Washington, USA

Abstract

Generative artificial intelligence's (GenAI) environmental footprint has been an increasing point of concern, but information about it is obscured, scant, and difficult to access and interpret. In this uncertain landscape, how is meaning around GenAI's environmental impact created, legitimized, and reproduced? Our ongoing efforts to investigate discourse surrounding GenAI and its effects on the environment include the collection and analysis of data across online platforms, news articles, and scientific literature. We believe our work could support HCI opportunities and challenges in this space; namely eco-feedback tool development or on-the-ground engagement of local communities. We welcome further opportunities to ideate around how our research could support broader missions in HCI.

Keywords

AI, Environment, Sustainability, Water Usage, Energy Usage, Carbon Emissions

ACM Reference Format:
Nino Migineishvili, Amelia Lee Doğan, Anoolia Gakhokidze, and Chandana Robba. 2026. "Is the AI Water Issue Fake:" Discourse of GenAI's Environmental Footprint. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (CHI '26).* ACM, New York, NY, USA, 3 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

The matrioshka brain [5], a concept introduced by Robert Bradbury, is a fictional computational megastructure that harnesses the full energy output of a star as fuel. Though purely hypothetical, this idea seems increasingly relevant as major technology companies lobby for the development of such megastructures in the form of hyperscale datacenters in the push for AI [15]. OpenAI, for instance, plans to deploy a $100 million dollar data center with reports estimating that it would require 5GW of power, equivalent to 5 nuclear reactors, to operate [6, 20, 23]. Meanwhile, Meta is targeting an installed base of 650,000 H100-equivalent GPUs [17]. Additionally, GPU cloud provider CoreWeave is making significant investments, including a $1.6 billion facility in Plano, Texas, which will consume up to 50MW of critical IT power, with plans to scale up to a 250MW data center [16].

The rapid growth in demand for AI computing clusters and an increased focus on data center capacity has placed significant pressure on energy production, power grids, and the environment, increasing carbon emissions [15] and straining water resources [19]. These concerns are not new, with the U.S. government flagging escalating energy demands as a pressing concern as early as 2008 [26]. Concurrently, Greenpeace has actively campaigned to address the issue [2], highlighting how growth in the AI and data center sector might be at odds with environmental goals [21]. But the topic of **AI's environmental footprint remains obscured, contested, and ever-evolving**.

One post on Reddit, for instance, reads that a *"100 word email written by AI takes about one bottle of water."* A responding comment disagrees, stating that *"this number comes from a misundertanding of the report."* The commenter presents a different heuristic: One bottle of water for every 300 queries. Similarly, a New York Times podcast [24] asks, "Is the A.I. Water Issue Fake?"

Other news articles link the Great Lakes' falling water levels to new planned data centers [25]. Even Hank Green, a noted YouTuber, has waded into the discussion about AI water usage[1] along with other outlets [3, 18]. In this information disorder ecosystem [14, 29], calculation errors are easy to make, as Karen Hao did in Empire of AI [11], since many of AI's environmental impacts are purposefully obfuscated by corporations and difficult to estimate [4, 7].

This leads us to ask the research questions:

- How are discourses about the environmental impacts of AI systems created, legitimized, and contested?
- What heuristics are users employing, and in what contexts?
- Ultimately, to what extent does knowledge of these AI's environmental footprint support reducing AI's environmental footprint?

*Authors contributed equally to this submission.

CHI '26, Woodstock, NY

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
https://doi.org/XXXXXXX.XXXXXXX

[1] Why is Everyone So Wrong About AI Water Use?

In the following sections, we describe initial steps undertaken to investigate discourse of AI's environmental footprint (Section 2), as well as minor preliminary findings (Section 3). We then connect our findings and experiences to the grand challenges identified by the HCI-TERRA workshop organizers (Section 4), and conclude with remaining open questions (Section 5).

2 Methods

2.1 Data Collection

To understand discourse around the environmental impacts of AI, we are collecting data from a variety of sources.

- **Reddit Data.** We are collecting Reddit data related to GenAI and its environmental footprint. Our initial approach has been to scrape all posts and associated comments that meet this criteria, irrespective of specific subreddits. Using keywords related to GenAI (e.g., "generative AI," "ChatGPT," "Large Language Models") and environmental impacts (e.g., "water use," "carbon emissions," "ewaste"), we have collected over 500 posts made since Novemeber 30th, 2022 (the release of ChatGPT) through the February 2026.
- **News Reports and Scholarly Articles.** We are also collecting newspaper articles and scholarly articles related to GenAI and its environmental footprint. As of now, we have identified 83 relevant sources, but our search is ongoing.

2.2 Data Analysis

We are employing both a deductive and an inductive coding approach to analyze our data. Our deductive coding is guided the research questions around discourse creation. These codes correspond to, for instance, the tone and metaphors used in Reddit posts and news articles. We document the types of comparisons invoked to create heuristics to contextualize metrics of AI's environmental footprint. Inductive coding is meant to uncover underlying themes we could not anticipate.

We are in the process of deriving a codebook whereby three authors are reviewing the data in parallel and negotiating commonalities and differences in their observations.

3 Preliminary Findings

The following represent initial observations we have made through our data collection and analysis process. They by no means constitute rigorous conclusions reached from our data. Nevertheless, we believe that preliminary findings can be informative to workshop discussions, and in shaping our thoughts around HCI's role in addressing AI's environmental impacts.

Expressions of Care. Individuals online have engaged in discourse around *care*. One Reddit user, for instance, asked why *"people [aren't] more concerned about the environmental impact of AI?...Do people really not care?"* Several users felt isolated in their concerns or wanted to understand why others did not share their worries. This reflects prior work by scholars; caring for the environment is a precursor to change in the ways we design and use technology [22].

Polarized Opinions. Perhaps unsurprisingly, discourse around AI's environmental impact appears to be highly polarized along the anti- versus pro-AI fault line. On the one hand, one user wondered why *"pro-AI people [do] not see the negative environmental impacts"* of the technology, implying their stance as *"anti-AI."* On the other hand, another Reddit user wrote that they would not *"stop using artificial intelligence because [others] pretend to care about the environment,"* citing hypocrisy of those engaged with technology at all.

Inconsistent and Non-Standardized Metrics. As previously noted, estimates and measurements of AI's environmental footprint are difficult to assess, obtain, and contextualize [27]. Across sources, we've encountered drastically different estimates. For instance, a Washington Post article projects that demand for power up from AI will increase 50% by 2030 [12]. A different Reuters article, however, puts that number at 6.7% [13]. While these discrepancies have been documented previously, we are hoping to quantify the extent of these contradictions.

4 HCI Opportunities and Challenges

Our research presents several directions for HCI challenges and opportunities. We welcome further discussion of these directions during the workshop, and hope to contribute in reducing GenAI's climate impacts.

4.1 Tools for Eco-Feedback

Our work could inform the design of an eco-feedback tool for GenAI users. Through analyzing discourse, we will be able to extract metrics of AI's environmental footprint and collate them from multiple sources. This can inform the creation of metrics presented to users not as absolute numbers, but rather as ranges incorporating uncertainty. Moreover, our exploration of the metaphors invoked and comparisons made in the conceptualization and sense-making of environmental impacts could dictate how an eco-feedback tool should be framed, and what response this framing might evoke.

4.2 Supporting Communities with Resistance On-the-Ground

An eco-feedback tool [9] could be a low-hanging continuation of our work and is an established research area in sustainable HCI. However, we believe that a much more effective approach involves supporting communities fighting efforts to expand data centers and GenAI facilities [8].

A Data Watch report [1] has identified at least 142 activist groups across 24 states opposing or delaying the construction of new data centers. Cases in California [28], Oregon [1], Virginia [1], or Georgia [10] show that local communities are fighting back, often motivated by water consumption, noise, and green space preservation. The environmental concerns are shaping direct actions of activists.

Our research team is considering more directly engaging activist groups. One method of doing this would be through more intentional or purposive data collection of their perspectives, information, or technology needs. During a recent panel discussions the authors attended around data centers in California, speakers consistently noted that accurate information of AI's environmental footprint was unavailable. These obscured metrics have impacts on how communities on-the-ground can organize.

More compelling, therefore, would be to involve them in shaping our research design. Participatory efforts could inform what type of discourses around environmental impacts should be collected in the first place and offer a different lens of data analysis.

5 Open Questions

There are many challenges we foresee in this work. Primarily, GenAI's environmental impact is polarized and takes place in isolated silos. Can HCI methods support breaking this paradigm, and/or should they? Second, how can we as researchers wrestle with and understand GenAI's environmental impacts, given the degree of contestations and uncertainty?

References

[1] 2025. $64 billion of data center projects have been blocked or delayed amid local opposition. https://static1.squarespace.com/static/67819031da098341c45ac84a/t/6849bcfe640a951f79e00715/1749662975141/Data+Center+Watch+Report+.pdf

[2] Gary Cook and Elizabeth Jardim. 2019. *Clicking Clean Virginia: The Dirty Energy Powering Data Center Alley.* Technical Report. Greenpeace Reports, Washington, DC. https://storage.googleapis.com/planet4-usa-stateless/2024/11/315295ef-greenpeace-click-clean-virginia-2019.pdf

[3] Casey Crownhart. 2025. These four charts sum up the state of AI and energy. https://www.technologyreview.com/2025/04/17/1115320/four-charts-ai-energy/

[4] Alex de Vries-Gao. 2026. The carbon and water footprints of data centers and what this could mean for artificial intelligence. *Patterns* 7, 1 (2026).

[5] Jacob Eddison, Joe Marsden, Guy Levin, and Darshan Vigneswara. 2017. P6_5 Matrioshka Brain. *Physics Special Topics* 16, 1 (2017).

[6] David Gelles. 2024. A.I.'s insatiable appetite for Energy. https://www.nytimes.com/2024/07/11/climate/artificial-intelligence-energy-usage.html?unlocked_article_code=1.pE4.5bCc.NeyorV1DWfnH&smid=url-share&login=smartlock&auth=login-smartlock

[7] A Shaji George, AS Hovan George, and AS Gabrio Martin. 2023. The environmental impact of AI: a case study of water consumption by chat GPT. *Partners Universal International Innovation Journal* 1, 2 (2023), 97–104.

[8] Sinem Görücü, Luiz A. Morais, and Georgia Panagiotidou. 2025. A Critical Analysis of Machine Learning Eco-feedback Tools through the Lens of Sustainable HCI. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25).* Association for Computing Machinery, New York, NY, USA, Article 885, 18 pages. doi:10.1145/3706598.3713198

[9] Nora Graves, Vitus Larrieu, Yingyue Trace Zhang, Joanne Peng, Varun Nagaraj Rao, Yuhan Liu, and Andrés Monroy-Hernández. 2025. GPTFootprint: Increasing Consumer Awareness of the Environmental Impacts of LLMs. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25).* Association for Computing Machinery, New York, NY, USA, Article 296, 16 pages. doi:10.1145/3706599.3719708

[10] Kylie Hanson. 2026. Why residents in one small town are fighting data centers.

[11] Karen Hao. [n. d.]. Empire of ai - water footprint changes. https://karendhao.com/20251217/empire-water-changes

[12] Evan Harper and Caroline O'Donovan. 2024. As data centers for AI strain the power grid, bills rise for everyday customers.

[13] Leila Kearney and Hampton Liz. 2025. US power stocks plummet as DeepSeek raises data center demand doubts.

[14] Nina Lutz, Stephen Prochaska, Laura Kurek, Marianne Aubin Le Quéré, Jason Greenfield, Joseph S Schafer, Phil Tinn, Daniel Schroeder, Shiva Darian, Sukrit Venkatagiri, et al. 2025. Beyond Information: Online Participatory Culture and Information Disorder. In *Companion Publication of the 2025 Conference on Computer-Supported Cooperative Work and Social Computing.* 120–126.

[15] Dan Milmo. 2024. Google's emissions climb nearly 50% in five years due to AI Energy Demand. https://www.theguardian.com/technology/article/2024/jul/02/google-ai-emissions

[16] Sebastian Moss. 2023. CoreWeave plans $1.6bn Ai Cloud Data Center in Plano, Texas. https://www.datacenterdynamics.com/en/news/coreweave-plans-16bn-ai-cloud-data-center-in-plano-texas/

[17] Sebastian Moss. 2024. Meta to operate "600,000 H100 GPU equivalents of compute" by year-end.

[18] Leonardo Nicoletti, Michelle Ma, and Dina Bass. [n. d.]. The Ai Boom is draining water from the areas that need it most. https://www.bloomberg.com/graphics/2025-ai-impacts-data-centers-water-data/

[19] Eric Olson and Anna Grau. 2024. Data centers draining resources in water-stressed communities. https://utulsa.edu/news/data-centers-draining-resources-in-water-stressed-communities

[20] Dylan Patel and Daniel Nishball. 2024. Microsoft Swallows openai's core team – GPU capacity, incentive structure, intellectual property, openai rump state. https://semianalysis.com/2023/11/20/microsoft-swallows-openais-core-team/

[21] Chris Preist, Daniel Schien, and Eli Blevis. 2016. Understanding and mitigating the effects of device and cloud service design decisions on the environmental footprint of digital infrastructure. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems.* 1324–1337.

[22] Maria Puig De La Bellacasa. 2015. Making time for soil: Technoscientific futurity and the pace of care. *Social Studies of Science* 45, 5 (Oct. 2015), 691–716. doi:10.1177/0306312715599851

[23] Reuters. [n. d.]. Microsoft, OpenAI plan $100 billion data-center project, Media Report says | Reuters. https://www.reuters.com/technology/microsoft-openai-planning-100-billion-data-center-project-information-reports-2024-03-29/

[24] Kevin Roose, Casey Newton, Rachel Cohn, Whitney Jones, Jen Poyant, Katie McMurran, Dan Powell, Marion Lozano, and Rowan Niemisto. 2025. Australia Kicks Kids Off Social Media + Is the A.I. Water Issue Fake? + Hard Fork Wrapped. *The New York Times* (Dec. 2025). https://www.nytimes.com/2025/12/12/podcasts/hardfork-australia-water.html

[25] Stephen Starr. 2025. Water levels across the Great Lakes are falling – just as US data centers move in. https://www.theguardian.com/technology/2025/dec/16/great-lakes-us-data-centers

[26] Global Trends. 2008. A transformed World. *National Intelligence Council* 15 (2008).

[27] Alex de Vries-Gao. 2026. The carbon and water footprints of data centers and what this could mean for artificial intelligence. *Patterns* 7, 1 (Jan. 2026). doi:10.1016/j.patter.2025.101430

[28] Claire Wang. 2025. Rage against the machine: a California community rallied against a datacenter – and won.

[29] Claire Wardle. 2018. The need for smarter definitions and practical, timely empirical research on information disorder. *Digital journalism* 6, 8 (2018), 951–963.

Received 14 February 2026