



Not All Water Consumption Is Equal: A Water Stress Weighted Metric for Sustainable Computing

YANRAN WU, Purdue University, USA

INEZ HUA, Purdue University, USA

YI DING, Purdue University, USA

Water consumption is an increasingly critical dimension of computing sustainability, especially as AI workloads rapidly scale. However, current water impact assessment often overlooks where and when water stress is more severe. To fill in this gap, we present SCARF, the first general framework that evaluates water impact of computing by factoring in both spatial and temporal variations in water stress. SCARF calculates an Adjusted Water Impact (AWI) metric that considers both consumption volume and local water stress over time. Through three case studies on LLM serving, datacenters, and semiconductor fabrication plants, we show the hidden opportunities for reducing water impact by optimizing location and time choices, paving the way for water-sustainable computing. The code is available at <https://github.com/jojacola/SCARF>.

CCS Concepts: • **Social and professional topics** → **Sustainability**; • **Hardware** → **Impact on the environment**; • **Computing methodologies** → **Machine learning**.

Additional Key Words and Phrases: Sustainability, Water Consumption, Large Language Model, Semiconductor Manufacturing, Datacenter

1 Introduction

The rapid growth of computing demand, particularly in AI workloads, has significantly increased water consumption across the whole computing stack, intensifying water-related sustainability concerns [27, 48]. Recent studies show that serving 10 to 50 medium-length queries from a large language model (LLM) such as GPT-3 can consume 500 mL of water [27]. Datacenters, which power AI services, consume vast amounts of water for cooling. One Google datacenter in Iowa consumed about 3.7 million cubic meters of water in a year [12], and according to Bloomberg, nearly two-thirds of new U.S. datacenters since 2022 are located in high water-stress regions [34]. Semiconductor manufacturing is also water-intensive, as large volumes of ultra-pure water are used to clean silicon wafers. For instance, Intel [16] reported its annual global water consumption of 9.6 million cubic meters. Globally, water consumption by all sectors has increased exponentially over the past century [22]. This has led to conflicts between stakeholders, including the computing industry. For example, in Taiwan, TSMC has competed with agriculture for water, especially during drought years [2, 49].

Water consumption refers to the volume of water that is removed from the immediate environment, including water that is evaporated, transpired, or incorporated into products or processes [29, 36]. This metric differs from *water withdrawal*, which includes all water extracted from the sources, including that discharged back after use. Due to this removal from local watershed, *water consumption* is more indicative of assessing regional water impact. Following prior

work [27, 35], in this paper, we focus on *Scope 1* water consumption (i.e., water consumed directly on-site, such as in cooling systems) and *Scope 2* (i.e., off-site water consumption associated with electricity generation). *Scope 3* (i.e., water consumed in upstream supply chain) is not discussed due to the lack of public data.

The environmental impact of water consumption depends heavily on spatial and temporal context. Spatially, the same amount of water consumption can lead to greater environment burden in arid regions than in water-abundant regions. For instance, both Central California and the Mississippi River Basin are major farming regions, but California's dry climate has caused more rapid groundwater level decline and thousands of dry wells [19, 20, 47]. Temporally, water supply fluctuates over time. For example, California experiences distinct wet and dry seasons, and climate projections indicate these fluctuations will become more pronounced in the future [30]. Therefore, we must account for both the geographic location and the temporal trajectory when analyzing water impact.

Prior work has explored water impact in computing, estimating water consumption of AI workloads such as LLM training and inference [8, 27, 33, 50], as well as datacenters [11, 13, 26, 39]. These efforts primarily focus on modeling and profiling the amount of water consumed by computational tasks or facility operations, providing valuable insights of water demand across different computing layers. In parallel, others have begun to incorporate regional water conditions into system design, including infrastructure deployments [6, 40] and workload scheduling [17, 18, 21, 28]. These studies aim to reduce water impact by integrating water consumption as a factor in facility siting or computing workload management.

However, existing efforts face two key limitations. First, many focus solely on the quantity of water consumed, ignoring spatial variability in water stress [11, 13, 17, 18, 26–28, 33, 39, 50]. Second, studies that incorporate regional water data often rely on static or short-term information [6, 8, 21, 40], failing to account for how water stress evolves over time due to climate change or long-term resource depletion. These two limitations ignore where and when water stress occurs, leading to an incomplete picture of water impact. Thus, there is an urgent need for a spatial- and temporal-aware framework to evaluate water impact in computing.

To bridge this gap, we present SCARF (Stress-Corrected Assessment of Water Resource Footprint), a general framework for evaluating water impact from computing. The key insight is to integrate spatial and temporal variation in *water stress*, which is the ratio of local water demand to supply, reflecting regional scarcity. SCARF follows four steps: (1) model water consumption volume, including both on-site (Scope 1) and off-site (Scope 2) consumption; (2) map each site to its corresponding hydrological basin; (3) calculate a Water Stress Factor (WSF) using cumulative projections over time; and (4)

Authors' Contact Information: Yanran Wu, Purdue University, West Lafayette, IN, USA, wu2187@purdue.edu; Inez Hua, Purdue University, West Lafayette, IN, USA, hua@purdue.edu; Yi Ding, Purdue University, West Lafayette, IN, USA, yiding@purdue.edu.




Fig. 1. Overview of the SCARF Framework.

compute the Adjusted Water Impact (AWI) by multiplying total water consumption with the WSF, yielding a unified metric for water sustainability. To demonstrate the effectiveness and generalizability of SCARF, we present three case studies covering LLM serving, datacenters, and semiconductor fabrication plants. Our key takeaways for sustainable computing deployment include:

- *LLM serving*: The adjusted water impact of deploying LLMs is highly location-sensitive and can vary significantly due to seasonal water supply changes.
- *Datacenters*: The adjusted water impact of datacenters depends on regional water stress, consumption volume, and weighting of long-term versus short-term water stress.
- *Semiconductor fabrication plants*: Fab plants located in high water stress regions can impose a much higher adjusted water impact than those in low-stress regions.

We summarize the contributions as follows:

- Introducing SCARF, the first systematic and general framework for evaluating the water impact with water stress in computing.
- Designing a unified metric that incorporates both spatial and regional variations in water stress, enabling sustainable planning of the computing infrastructure.
- Conducting case studies in LLM serving, datacenters, and semiconductor fab plants, enabling comparison for water impact evaluation across the full stack of computing.

2 The SCARF Framework

Figure 1 shows the four steps of SCARF. The key insight is to assess the environmental impact of water consumption by factoring in where and when water stress occurs. SCARF models total water consumption (on-site and off-site), maps site locations to hydrological basins, calculates a Water Stress Factor (WSF) by aggregating water stress over time, and computes the Adjusted Water Impact (AWI) by applying the WSF. We detail each step next.

2.1 Step 1: Raw Water Consumption Modeling

We first model the raw operational water consumption, which includes both on-site (Scope 1) and off-site (Scope 2) consumption, following the same methodology from prior work [27, 33, 50].

On-site water consumption refers to water that is directly consumed at the facility, including cooling, humidification, and other operational needs. It can be quantified using the *on-site Water Usage Effectiveness* (WUE_{on}) metric in the unit of liters/kWh:

$$W_{on} = P \cdot t \cdot WUE_{on}, \quad (1)$$

where W_{on} is the on-site water consumption (in liters), P is the average power draw (in kilowatts), and t is the runtime (in hours).

Off-site water consumption refers to the indirect water consumption associated with electricity generation at power plants. This can be modeled by first computing the total energy used and then applying the water intensity of upstream electricity sources, denoted as the *off-site Water Usage Effectiveness* (WUE_{off}) metric in the unit of liters/kWh:

$$W_{off} = P \cdot t \cdot PUE \cdot WUE_{off}, \quad (2)$$

where W_{off} is the off-site water consumption (in liters), and PUE is the *Power Usage Effectiveness* calculated by dividing the total energy used by the facility by the energy used by its IT equipment.

2.2 Step 2: Spatial Mapping

To account for spatial variation in water stress, we map the physical location of each facility to its corresponding hydrological basin, where water stress is actually measured [23]. Since water stress varies by watershed—not by country or region—the same facility in two different watersheds can face different water challenges (see Figure 2). To perform this mapping, we first use the facility’s geographic name (e.g., city or administrative region) to query the Aqeduct API [14], which returns the corresponding watershed ID. These watershed IDs from the global hydrographic classification [25] are unique for each hydrological basin. We then use this watershed ID to retrieve both current and projected water stress values from the Aqeduct 4.0 dataset [23]. Aqeduct 4.0 is the latest iteration of World Resource Institute’s water risk framework to translate complex hydrological data into water risk indicators, including water stress levels that will be used in the next step.

2.3 Step 3: Temporal Water Stress Aggregation

To account for temporal variation in water stress, we introduce the Water Stress Factor (WSF)—a time-weighted aggregation of basin-level water stress values. To reflect future uncertainty and potential impact, WSF applies a user-defined discount rate γ , following principles from environmental economics [43, 44]. A higher discount rate reduces the weight of future stress; an infinitely large rate ignores future stress entirely. Below, we derive both short-term and long-term WSFs to support different application needs.

Short-Term WSF. For short-term WSF analyses, such as evaluating the immediate water stress of computing tasks like LLM serving, we focus on the current impact only:

$$WSF_b^{short} = WS_{t_0, b}, \quad (3)$$

where $WS_{t_0, b}$ denotes the water stress at basin b at time t_0 .

Long-Term WSF. For long-term facilities like datacenters, which operate for decades, factoring in future water stress provides a

more accurate view of sustainability, as regional water availability may change over time. However, future conditions are uncertain due to factors like climate variability and technological shifts. To account for this and weigh future impacts appropriately, we apply a discounting approach inspired by environmental economics [43, 44]:

$$WSF_b^{\text{long}} = \sum_{t \in T} w_t \cdot WS_{t,b}, \quad (4)$$

where $WS_{t,b}$ denotes the water stress at basin b at time t . The weight w_t is computed using a discount rate γ , reflecting the relative importance of future water stress:

$$w'_t = \frac{1}{(1+\gamma)^{t-t_0}}, \quad w_t = \frac{w'_t}{\sum_{t' \in T} w'_{t'}} \quad (5)$$

where t_0 is the baseline year, and T is the set of years considered. According to Aqueduct 4.0 [23], the baseline year is 2019, representing current conditions based on over 40 years of hydrological data (1979–2019), and $T = \{\text{baseline}, 2030, 2050, 2080\}$. A lower discount rate gives more weight to future water stress, while a higher rate prioritizes near-term impacts. Note that when $\gamma \rightarrow \infty$ where near-short impacts are maximized, all w'_t tend to zero except for $t = t_0$, where $w'_{t_0} = 1$. In this case, the equation will be reduced to the short-term WSF in Equation (3).

2.4 Step 4: Adjusted Water Impact Calculation

After aggregating water stress at the basin level, we compute the **Adjusted Water Impact** (AWI) by multiplying the total raw water consumption by the corresponding Water Stress Factor (WSF). AWI captures both onsite and offsite water consumption associated with water stress levels in basin b . Depending on the time horizon, we apply either the short-term or long-term WSF.

$$AWI = (W_{\text{on}} + W_{\text{off}}) \times WSF_b \quad (6)$$

3 Case Study I: LLM Serving

In this section, we use SCARF to evaluate the short-term water impact of LLM serving. LLM serving is flexible and dynamic; cloud providers can shift regions, scale deployments, or update models within months [45, 46]. Because of this short lifecycle, the adjusted water impact depends mainly on the immediate time and location of execution, not long-term trends.

3.1 Evaluation Methodology

We run three Qwen2.5 models (7B, 14B, 32B) [46] on a server with an NVIDIA H100 GPU and Intel Xeon 8480+ CPU. We assume the workload runs in Microsoft datacenters across different regions, where on-site WUEs and PUEs are publicly reported [31]. Location-specific off-site WUEs are obtained from the Water Impact Tool [41].

Power measurement. We collect real-time power consumption data P for the GPU and CPU during inference. Power readings are sampled every 200ms using NVIDIA’s `pynvml` API for the GPU and Intel’s `psutil` API for the CPU.

Workload and profiling. We conduct experiments on the open-source LLM serving platform vLLM [24] and simulate user queries with requests sampled from the ShareGPT dataset [38], which contains real multi-turn conversations. For each model, we gradually increase the query-per-second (QPS) rate from 1 to 30. We identify




Fig. 2. On-site and Off-site WUEs of Microsoft datacenter locations (top) and their corresponding local water stress values (bottom). ★ refers to the location with lowest value.




Fig. 3. Adjusted Water Impact (AWI) per request for different model sizes in different Microsoft datacenter locations.

the QPS that delivers peak throughput, then record the power P and latency t for each request at that QPS.

Adjusted water impact. Once we have power P , execution time t , PUE, and both on-site and off-site WUEs, SCARF calculates the raw water consumption based on Equation (1) and Equation (2). Then, SCARF computes the final adjusted water impacts using Equation (6) for each model size at different locations.

3.2 Evaluation Results

3.2.1 Adjusted water impact in different locations. We first compare the adjusted water impact across different locations. In Figure 2, the top chart shows on-site and off-site WUEs for datacenters by location. The Texas (TX) datacenter has the lowest WUE, meaning it consumes the least water per unit of electricity. The bottom chart shows water stress levels, with Iowa (IA) having the lowest water stress. These charts reveal large differences in both operational efficiency and environmental conditions across locations.

We then compute adjusted water impact per request by model size using SCARF. As shown in Figure 3, the same model can have drastically different adjusted water impacts depending on where it is deployed. For example, the datacenter in Quincy, Washington (WA)—a high-desert area—faces high water stress and less efficient WUE, leading to over 1000× higher adjusted water impact per request than in low-stress locations. Comparing Illinois (IL) and Wyoming




Fig. 4. Water stress levels across various Microsoft datacenter locations over the 12 months of the baseline year 2019.




Fig. 5. Adjusted water impact (AWI) per request of 32B model in two locations: Wyoming (WY) and Arizona (AZ).

(WY), serving a small 7B model in WY can have a higher adjusted water impact than running a large 32B model in IL.

Takeaway 1: The adjusted water impact of deploying LLMs is highly location-sensitive. Same workloads can have orders-of-magnitude differences in adjusted water impact depending on where they are served.

3.2.2 Adjusted water impact in different months. Water stress at the same location can change month to month due to seasonal shifts in climate and water availability. As shown in Figure 4, water stress fluctuates over time, with locations like Wyoming (WY) and Arizona (AZ) showing particularly large swings.

We therefore select WY and AZ to examine how adjusted water impact changes over time for the same model. Figure 5 shows the monthly adjusted water impact per request for a 32B model. From January to March, AZ has relatively lower water stress, resulting in lower adjusted water impact than WY. But during the arid months of April to June, AZ’s water stress spikes. As a result, its adjusted impact surpasses WY’s, even reaching more than twice as much.

Takeaway 2: Even in the same location, seasonal changes can significantly affect adjusted water impact. When you deploy matters—not just where.

4 Case Study II: Datacenters

In this section, we use SCARF to evaluate the long-term water impact of datacenters. Since datacenters operate for decades [9, 10],

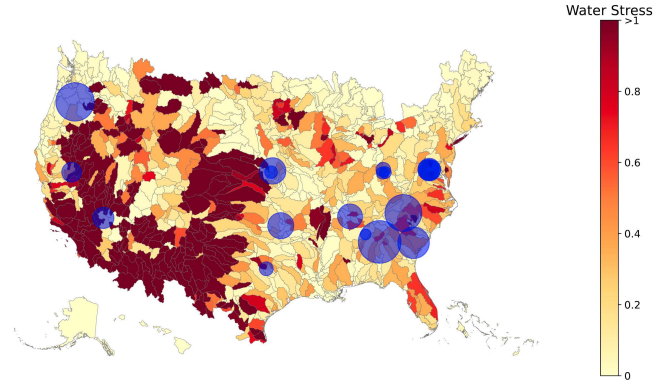


Fig. 6. Geographic distribution of water stress (background) and raw water consumption of Google’s U.S. datacenters (blue circles). Circle size indicates annual water consumption volume.

their adjusted water impact must account for both current stress and future projections.

4.1 Evaluation Methodology

We analyze the annual water consumption of Google’s U.S. datacenters. We focus on Google because it provides the most detailed public reporting among major cloud providers [4, 12, 32], including site-specific data on annual water consumption and PUE [12].

Energy consumption estimation. Since Google does not disclose total energy consumption per datacenter, we estimate each site’s power capacity using public data from the industry database DataCenters.com [7]. For each datacenter, we identify all nearby sites within a 100-mile radius and take the maximum reported power capacity P (kWh) as a proxy for the Google site’s power capacity. We note that datacenters do not continuously operate at peak power. According to Google’s datacenter power dataset [37], the average power utilization of datacenters is approximately 70%. Assuming 24/7 operation, we compute the annual energy consumption as $E = P \times 24 \times 365 \times 0.7$. We then apply the reported Power Usage Effectiveness (PUE) to estimate the total energy consumption.

Adjusted water impact. Given the decades-long lifespan of datacenters, SCARF evaluates the annual adjusted water impact using both current water stress data and future projections for 2030, 2050, and 2080. Aqueduct 4.0 provides three pathways to represent future scenarios of climate and socioeconomic change [23]: *Business-as-Usual* pathway with moderate emissions and adaptation, *Optimistic* scenario with sustainable development, and *Pessimistic* scenario under high emissions. We adopt business-as-usual pathway in this analysis. We apply a discount rate of 3% in long-term water stress factor calculation, following the U.S. Water Resources Development Act of 1974 (WRDA) [1]. We also analyze how changes in discount rates and future scenarios affect long-term adjusted water impact, helping to understand sustainability trade-offs under uncertainty.




Fig. 7. (a) Annual water consumption of Google's U.S. datacenters, (b) AWI of Google's U.S. datacenters using the 3% discount rate, both grouped by regional water stress levels.

4.2 Evaluation Results

4.2.1 Adjusted water impact in different locations. We begin by mapping water stress and datacenter water consumption. In Figure 6, western areas such as Southern California and Arizona exhibit higher water stress. Many large water-consuming datacenters are located in lower-stress regions, suggesting water availability may influence site choice. Still, some datacenters are placed in high-stress watersheds, raising long-term sustainability concerns.

To dig deeper, Figure 7(a) groups datacenters by their regional water stress levels and compares their annual water consumption. Following Aqueduct's classification [23], we define water stress levels as: **low** (<0.1), **medium** (0.1-0.4), and **high** (>0.4). While most datacenters fall in low- or medium-stress regions, several high-stress regions still show substantial water consumption.

SCARF then calculates the adjusted water impact by integrating both water consumption and water stress factors. As shown in Figure 7(b), datacenters in medium-stress regions exhibit the highest adjusted water impact, more than 2× higher than datacenters in high-stress regions. In particular, sites in GA, NC, and SC show significantly higher adjusted water impact than those in TX or NV2, which face higher stress but consume less water.

Takeaway 3: Water stress alone is not sufficient to capture true water impact. High water consumption in medium-stress regions can cause greater long-term impact than moderate consumption in high-stress regions.

4.2.2 Sensitivity check of discount rate. The discount rate controls how much future impacts are valued compared to present ones. As shown in Figure 8(a), higher discount rates sharply reduce the weight of future years, emphasizing short-term outcomes. We adopted a 3% discount rate in our main analysis above, following the U.S. Water Resources Development Act of 1974 (WRDA 1974) [1].

However, discount rates vary by perspectives on intergenerational responsibility and policy planning. For example, [42] advocates a low discount rate of 1.4% to reflect ethical responsibility to




Fig. 8. Sensitivity Check of Discount Rate. (a) Discounting weights of future years under different discount rates. (b) Adjusted water impact (AWI) of Google's datacenters under different discount rates, with shaded areas representing variability across different future scenarios.

future generations. In contrast, U.S. federal budgeting often uses a 7% rate to favor near-term returns [5].

Figure 8(b) shows how adjusted water impact (AWI) for Google datacenters changes under different discount rates. The solid lines represent the median AWI across all future scenarios; the shaded areas reflect uncertainty among the three future pathways [23]. As expected, lower discount rates place more weight on long-term water stress, resulting in higher AWI and greater uncertainty with larger shaded areas. Notably, the relative comparison between datacenters can shift depending on the discount rate. For example, as the discount rate decreases, TX's AWI increases faster than NV2's, indicating that TX has higher long-term water stress. When long-term impacts are prioritized (i.e., at low discount rates), NV2 becomes the more sustainable datacenter location.

A more striking shift appears between NV and VA3. At low discount rates, NV has lower AWI, suggesting that NV is more sustainable in long term. However, as the discount rate increases, NV's AWI surpasses VA3's. This implies that when short-term water stress is more heavily weighted, VA3 is more sustainable.

Takeaway 4: The choice of discount rate significantly alters datacenter sustainability rankings. A site that is sustainable in the long term (e.g., NV over VA3) may appear less favorable when short-term impacts are prioritized.

5 Case Study III: Semiconductor Fab Plants

In this section, we use SCARF to evaluate the long-term water impact of semiconductor fabrication plants (fabs), which often operate over multiple decades. For instance, several Intel fabs were built in the late 20th century and are still in active use today [3].

5.1 Evaluation Methodology

We analyze the annual water consumption of Intel's U.S. semiconductor fabs. We focus on Intel because it provides detailed public reporting, including site-specific data on annual water consumption [16] and energy consumption [15].

Adjusted water impact. Given the decades-long lifespan of fab plants, SCARF evaluates the annual adjusted water impact using




Fig. 9. (a) Annual water consumption of Intel's fab plants, (b) AWI of Intel's fab plants using the 3% discount rate.

long-term water stress aggregation, based on the Business-as-Usual scenario and a 3% discount rate, as in §4.

5.2 Evaluation Results

Figure 9 shows Intel's U.S. fabs in three main regions: Chandler in Arizona (AZ), Hillsboro in Oregon (OR), and Ronler Acres in New Mexico (NM). Both Arizona and Oregon host two campuses. Figure 9(a) illustrates the annual water consumption at each fab location. The OR2 site reports the highest water consumption with 1.82×10^{10} liters per year, followed by AZ2 in Arizona.

When SCARF accounts for water stress, the adjusted water impact (AWI) in Figure 9(b) shows significant differences. Despite consuming less water, the Arizona fabs have much higher AWI due to the region's severe water stress. In contrast, Oregon and New Mexico fabs show far lower AWI. This highlights how water stress, not just consumption volume, drives environmental impact.

Takeaway 5: Semiconductor fabs consume a large amount of water. Intel's Arizona fabs have higher environmental burdens than those in Oregon or New Mexico due to higher local water stress, despite lower water consumption. It emphasizes the need to factor in regional water availability in siting and operations.

6 Conclusions

We introduce SCARF, the first general water impact evaluation framework for computing by considering spatial and temporal variation in water stress. Using SCARF, we assess the water impact across three layers of the computing stack: LLM serving (software), datacenters (computing infrastructure), and semiconductor fabs (hardware manufacturing). Our findings highlight the opportunities for deploying more water-sustainable computing systems.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback. This work is supported by NSF CCF-2413870.

References

- [1] 1974. Water Resources Development Act of 1974. <https://www.congress.gov/bill/93rd-congress/house-bill/10203/text>. Public Law 93-251.
- [2] 2021. *Taiwan is Facing a Drought, and It Has Prioritized Its Computer Chip Business Over Farmers*. <https://www.nytimes.com/2021/04/08/business/taiwan-is-facing-a-drought-and-it-has-prioritized-its-computer-chip-business-over-farmers.html> The New York Times, Accessed: 2025-05-09.
- [3] 2024. Intel Plans Massive Fab Expansion in Oregon. <https://www.tomshardware.com/news/intel-plans-massive-fab-expansion-in-oregon>. Accessed: 2024-06-25.

- [4] Amazon Web Services. 2023. 2023 Amazon Sustainability Report – AWS Summary. <https://sustainability.aboutamazon.com/2023-amazon-sustainability-report-aws-summary.pdf>.
- [5] Nicole T Carter, Adam C Nesbitt, et al. 2016. Discount rates in the economic evaluation of US Army Corps of Engineers projects. *Congressional Research Service* 44594 (2016).
- [6] Li Chen and Aaron P Wemhoff. 2022. Characterizing Data Center Cooling System Water Stress in the United States. In *ASHRAE Annual Conference papers*.
- [7] Datacenters.com. 2025. Datacenters.com. <https://www.datacenters.com/>.
- [8] Clément Desroches, Martin Chauvin, Louis Ladan, Caroline Vateau, Simon Gosset, and Philippe Cordier. 2025. Exploring the sustainable scaling of AI dilemma: A projective study of corporations' AI environmental impacts. *arXiv preprint arXiv:2501.14334* (2025).
- [9] Patrick Donovan. 2013. Fundamentals of Managing the Data Center Life Cycle for Owners. *APC, Schneider Electric White Papers* (2013).
- [10] Equinix. 2013. Everything You Need to Know About DC11: Q&A with Data Center Guru Jim Farmer. <https://blog.equinix.com/blog/2013/05/10/everything-you-need-to-know-about-dc11-qa-with-data-center-guru-jim-farmer/>
- [11] Wedan Emmanuel Gnibga, Andrew A Chien, Anne Blavette, and Anne Cécile Orgerie. 2024. FlexCoolDC: datacenter cooling flexibility for harmonizing water, energy, carbon, and cost trade-offs. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 108–122.
- [12] Google LLC. 2024. Google Environmental Report 2024. <https://www.gstatic.com/gumdrop/sustainability/google-2024-environmental-report.pdf>.
- [13] Pranjoy Sen Gupta, Md Rajib Hossen, Pengfei Li, Shaolei Ren, and Mohammad A Islam. 2024. A dataset for research on water sustainability. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 442–446.
- [14] World Resources Institute. 2024. Aqueduct Water Risk Atlas. <https://www.wri.org/applications/aqueduct/water-risk-atlas>. Accessed: 2025-05-05.
- [15] Intel Corporation. 2024. *Explore Intel's U.S. Semiconductor Manufacturing Expansion*. <https://www.exploreintel.com/> Accessed: 2025-05-04.
- [16] Intel Corporation. 2024. *Intel 2023–24 Corporate Responsibility Report*. <https://csrreportbuilder.intel.com/pdfbuilder/pdfs/CSR-2023-24-Full-Report.pdf> Accessed: 2025-05-04.
- [17] Mohammad A Islam, Kishwar Ahmed, Hong Xu, Nguyen H Tran, Gang Quan, and Shaolei Ren. 2016. Exploiting spatio-temporal diversity for water saving in geo-distributed data centers. *IEEE Transactions on Cloud Computing* 6, 3 (2016), 734–746.
- [18] Mohammad A Islam, Shaolei Ren, Gang Quan, Muhammad Z Shakir, and Athanasios V Vasilakos. 2015. Water-constrained geographic load balancing in data centers. *IEEE Transactions on Cloud Computing* 5, 2 (2015), 208–220.
- [19] Ian James. 2024. *Groundwater depletion is worsening worldwide. Parts of California are suffering rapid declines*. <https://www.latimes.com/environment/story/2024-01-25/groundwater-depletion-is-worsening-worldwide> Los Angeles Times, Accessed: 2025-05-04.
- [20] Scott Jasechko, Hansjörg Seybold, Debra Perrone, Ying Fan, Mohammad Shamsudhuha, Richard G Taylor, Othman Fallatah, and James W Kirchner. 2024. Rapid groundwater decline and some cases of recovery in aquifers globally. *Nature* 625, 7996 (2024), 715–721.
- [21] Yankai Jiang, Rohan Basu Roy, Raghavendra Kanakagiri, and Devesh Tiwari. 2025. WaterWise: Co-optimizing Carbon- and Water-Footprint Toward Environmentally Sustainable Cloud Computing. In *Proceedings of the 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming* (Las Vegas, NV, USA) (PPoPP '25). Association for Computing Machinery, New York, NY, USA, 297–311. doi:10.1145/3710848.3710891
- [22] Matti Kummu, Joseph HA Guillaume, Hans De Moel, Stephanie Eisner, Martina Flörke, Miina Porkka, Stefan Siebert, Ted IE Veldkamp, and Philip J Ward. 2016. The world's road to water scarcity: shortage and stress in the 20th century and pathways towards sustainability. *Scientific reports* 6, 1 (2016), 38495.
- [23] Samantha Kuzma, Marc FP Bierkens, Shivani Lakshman, Tianyi Luo, Liz Saccoccia, Edwin H Sutanudjaja, and Rens Van Beek. 2023. *Aqueduct 4.0: Updated decision-relevant global water risk indicators*. World Resources Institute Washington, DC, USA.
- [24] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *SOSP*.
- [25] Bernhard Lehner and Günther Grill. 2013. Global river hydrography and network routing: baseline data and new approaches to study the world's large river systems. *Hydrological Processes* 27, 15 (2013), 2171–2186.
- [26] Nuoa Lei, Jun Lu, Arman Shehavi, and Eric Masanet. 2025. The water use of data center workloads: A review and assessment of key determinants. *Resources, Conservation and Recycling* 219 (2025), 108310.
- [27] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. 2024. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI

- Models. *Commun. ACM* (2024).
- [28] Pengfei Li, Jianyi Yang, Adam Wierman, and Shaolei Ren. 2024. Towards environmentally equitable AI via geographical load balancing. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*. 291–307.
- [29] Jordan Macknick, Robin Newmark, Garvin Heath, and KC Hallett. 2011. Review of operational water consumption and withdrawal factors for electricity generating technologies. (2011).
- [30] Iman Mallakpour, Mojtaba Sadegh, and Amir AghaKouchak. 2018. A new normal for streamflow in California in a warming climate: Wetter wet seasons and drier dry seasons. *Journal of hydrology* 567 (2018), 203–211.
- [31] Microsoft Corporation. 2024. Efficiency and Sustainability at Microsoft Datacenters. <https://datacenters.microsoft.com/sustainability/efficiency/>
- [32] Microsoft Corporation. 2024. Microsoft 2024 Environmental Sustainability Report. <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Microsoft-2024-Environmental-Sustainability-Report.pdf>.
- [33] Jacob Morrison, Clara Na, Jared Fernandez, Tim Dettmers, Emma Strubell, and Jesse Dodge. 2025. Holistically Evaluating the Environmental Impact of Creating Language Models. In *The Twelfth International Conference on Learning Representations (ICLR)*.
- [34] Leonardo Nicoletti, Michelle Ma, and Dina Bass. 2025. *AI Is Draining Water From Areas That Need It Most*. <https://www.bloomberg.com/graphics/2025-ai-impacts-data-centers-water-data> Bloomberg, Accessed: 2025-05-09.
- [35] Greenhouse Gas Protocol. 2011. Greenhouse gas protocol. *Sector Toolsets for Iron and Steel-Guidance Document* (2011).
- [36] Paul Reig. 2013. What's the difference between water use and water consumption? <https://www.wri.org/insights/whats-difference-between-water-use-and-water-consumption>.
- [37] Varun Sakalkar, Vasileios Kontorinis, David Landhuis, Shaohong Li, Darren De Ronde, Thomas Blooming, Anand Ramesh, James Kennedy, Christopher Malone, Jimmy Clidaras, et al. 2020. Data center power oversubscription with a medium voltage power plane and priority-aware capping. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. 497–511.
- [38] ShareGPT. [n.d.]. ShareGPT - Share and Save Your Conversations with AI. <https://sharegpt.com/>.
- [39] Arman Shehabi, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu Bakkar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, Dale Sartor, et al. 2024. 2024 United States Data Center Energy Usage Report. (2024).
- [40] Md Abu Bakar Siddik, Arman Shehabi, and Landon Marston. 2021. The environmental footprint of data centers in the United States. *Environmental Research Letters* 16, 6 (2021), 064017.
- [41] Md Abu Bakar Siddik, Arman Shehabi, Prakash Rao, and Landon T Marston. 2024. Spatially and temporally detailed water and carbon footprints of US Electricity generation and use. *Water Resources Research* 60, 12 (2024), e2024WR038350.
- [42] Nicholas Stern. 2007. *The economics of climate change: the Stern review*. cambridge University press.
- [43] Richard SJ Tol. 2011. The social cost of carbon. *Annu. Rev. Resour. Econ.* 3, 1 (2011), 419–443.
- [44] Martin L Weitzman. 1994. On the "environmental" discount rate. *Journal of Environmental Economics and Management* 26, 2 (1994), 200–209.
- [45] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruizhe Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuxiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. 2024. Qwen2 Technical Report. *arXiv preprint arXiv:2407.10671* (2024).
- [46] An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuxiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2024. Qwen2.5 Technical Report. *arXiv preprint arXiv:2412.15115* (2024).
- [47] Lindsey MW Yasarer, Jason M Taylor, James R Rigby, and Martin A Locke. 2020. Trends in land use, irrigation, and streamflow alteration in the Mississippi River Alluvial Plain. *Frontiers in Environmental Science* 8 (2020), 66.
- [48] Adam Zewe. 2025. *Explained: Generative AI's Environmental Impact*. <https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117> MIT News, Accessed: 2025-05-04.
- [49] Raymond Zhong and Amy Chang Chien. 2021. *Drought in Taiwan Pits Chip Makers Against Farmers*. <https://www.nytimes.com/2021/04/08/technology/taiwan-drought-tsmc-semiconductors.html> The New York Times, Accessed: 2025-05-09.
- [50] Guido Zuccon, Harrison Scells, and Shengyao Zhuang. 2023. Beyond CO2 Emissions: The Overlooked Impact of Water Consumption of Information Retrieval Models. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval (Taipei, Taiwan) (ICTIR '23)*. Association for Computing Machinery, New York, NY, USA, 283–289. doi:10.1145/3578337.3605121