# A Review of the Existing Literature on Al's Energy and Water Implications

Sarah An, Taylor Backer, Lucie Batista, Emilynne Newsom, Trevor Shepherd\* ucs24@hmc.edu Harvey Mudd College Claremont, California, USA

> Lucie Batista lbatista@hmc.edu Harvey Mudd College Claremont, California, USA

Sarah An saan@hmc.edu Harvey Mudd College Claremont, California, USA

Emilynne Newsom enewsom@hmc.edu Harvey Mudd College Claremont, California, USA

Lynn Kirabo lkirabo@hmc.edu Harvey Mudd College Claremont, California, USA Trevor Shepherd

Taylor Backer tbacker@hmc.edu

Scripps College

Claremont, California, USA

tshepherd@hmc.edu Harvey Mudd College Claremont, California, USA

# Abstract

As artificial intelligence (AI) continues to expand in both its reach and capability, its environmental impact, particularly energy and water usage, is becoming an urgent concern. This literature review synthesizes research from 31 written sources published between 2006 and 2025 to examine AI's environmental impact. It focuses on energy consumption, water usage, and the role of data centers. We use qualitative coding to categorize our findings into four main areas: energy, water, data centers, and solutions. Our results show a large increase in publications after 2020, as well as the largest category being energy. The literature highlights AI's significant and growing environmental impact, while identifying strategies to mitigate it. Some of these strategies include algorithmic improvements, alternative data center cooling systems, and strategic data center placement. Additionally, we discuss the limitations of current research and suggest areas of further research. We conclude by highlighting the need for both the development of stronger ethical frameworks and policy intervention, which can guide the development of sustainable AI.

# **CCS** Concepts

• Social and professional topics  $\rightarrow$  Sustainability.

#### Keywords

Artificial Intelligence, Data Centers, Energy, Literature Review, Sustainability, Water

# 1 Introduction

Artificial intelligence (AI) poses very exciting possibilities, but not without significant environmental costs. Behind every use of AI lies an energy and water intensive infrastructure that is rapidly expanding. As AI models become increasingly complex and widespread, J. Pablo Ortiz-Partida jportiz@ucs.org Union of Concerned Scientists USA

the computational resources required to develop and operate them grow exponentially alongside their environmental impact. These environmental costs, however, are often obscured from public view. From the extraction of raw materials needed to manufacture AI hardware, to the electricity and water used to train and deploy large models, the hidden infrastructure of AI is resource-intensive and unsustainable in its current form. For example, according to the OECD AI Policy Observatory, by 2027 global AI operations may require 4.2-6.6 billion cubic meters of water withdrawal, which is freshwater taken from ground or surface water sources and conveyed to a place of use [31]. This would equate to more than half of the United Kingdom's total annual water withdrawal [31]. This daunting demand in just a few years from now highlights the urgency of needing more effective and transparent development of AI systems as more and more users will be using generative AI in the near future.

As concerns over sustainability rise, it is critical to understand both the scale and nature of AI's environmental impact. This literature review analyzes 31 scholarly articles published between 2006 and 2025 to uncover how the academic community is responding to these emerging challenges. We examine trends in publication over time, disciplinary focuses, and thematic overlap between topics such as energy, water use, data centers, and proposed solutions. Our goal is to analyze the following objectives to provide valuable insights:

- To investigate the current state of AI research focusing on AI's unsustainable development, directly implemented solutions, algorithmic and training modifications, locations of data centers, and cooling of data centers.
- (2) To examine areas of further discussion centering on new ethical frameworks and the necessity of further reporting and transparency from technology corporations.

<sup>\*</sup>All authors contributed equally to this research.

The paper's organization will be as follows: Section 2, Background and Related Work, outlines the Union of Concerned Scientists' mission and relevant historical context of the rapid expansion of artificial intelligence, giving a more in depth look at the problem beyond the introduction. Section 3, Methods, describes how we found our sources to review and how we determined a reasonable time frame of sources to review. Section 4, Analysis, presents how we analyzed papers from our research database by describing our use of qualitative coding to create categories. Section 5, results, focuses on trends in the papers we analyzed by quantifying which themes appeared throughout and focuses on discoveries we found important. Section 6, discussion, examines the implications of our literature review results, present limitations of our literature review, and discuss future pathways of research in this field.

#### 2 Background and Related Work

2.0.1 Background on Clinic Sponsor. The Union of Concerned Scientists (UCS) is a national nonprofit organization founded in 1969 by scientists and students at the Massachusetts Institute of Technology, using science to achieve a better world. Their mission consists of putting rigorous, independent science into action, developing solutions and advocating for a healthy, safe, and just future. UCS focuses on a variety of issues, including the following:

- Combating climate change and seek to alleviate harm caused by the heat, sea level rise, and other consequences of runaway emissions
- They strive to develop sustainable ways to feed, power, and transport ourselves
- They work to reduce the existential threat of nuclear war
- They fight back when powerful corporations or special interests mislead the public on science
- They ensure their solutions advance racial and economic equity

While UCS tackles a wide range of pressing global challenges, the rapid expansion of artificial intelligence (AI) presents new environmental concerns regarding energy and water usage that are relevant to UCS's goals of advancing sustainability and combating climate change. As AI technologies scale, so do their demands on energy and water resources, highlighting the need to assess the sustainability of these models in terms of their water and energy consumption. Addressing these concerns aligns with UCS's commitment to promoting a healthy, safe, and just future.

2.0.2 Historical Account of Artificial Intelligence. To fully encompass how artificial intelligence currently impacts the world, we must understand how artificial intelligence came to be. In the 1950s, Alan Turing released a landmark research paper, "Computer Machinery and Intelligence", that laid the foundation for discussion about computer intelligence and its measurability, now known as "The Turing Test" [36]. In the 1960s, researchers and professors worked together to discuss the ethical and technical expectations for computer intelligence. By the 1980s, artificial intelligence prompted breakthroughs in computer vision systems and robotics. By the 2000s, artificial intelligence became commonplace in homes globally and internationally through the release of Amazon Alexa, and NASA's Rover exploration that used artificial intelligence to navigate the rocky terrain of Mars [26]. In recent years, OpenAI's initial release of ChatGPT 3.0 sparked excitement within the artificial intelligence space, which later boomed into a global sensation in 2022 with their public ChatGPT 3.5 model. The initial model of ChatGPT 3.5 gathered over one million subscriptions in the first five days of its release, showing public interest in engaging with artificial intelligence [28]. ChatGPT focuses on generative artificial intelligence (Gen AI), which can create new combinations and transformations of text, images, video, and other content based on patterns learned from existing human-created works. Generative artificial intelligence is a subsection within the study of artificial intelligence, and many forms beyond Gen AI are actively in research.

2.0.3 Current State of Artificial Intelligence's Sustainability Impact. The growth of AI offers vast potential for innovation, yet reveals critical environmental vulnerabilities. The substantial energy and water consumption required for AI infrastructure, specifically data centers, raises significant concerns about AI's sustainability, such as increased greenhouse gas emissions, water scarcity, and strain on local energy grids. Data centers are integral to the functioning of our digital infrastructure, yet they demand substantial quantities of both water and energy for their operation. These facilities rely heavily on electricity to power their servers and technology (and in doing so generate significant amounts of heat), while sophisticated cooling systems, which are essential for maintaining operational temperatures, frequently consume large volumes of water. Inadequate attention to these impacts can lead to unsustainable practices, where the environmental costs such as energy and water usage outweigh the benefits of AI advancements. When society lacks comprehensive policies and regulations that address the environmental implications of AI, the insufficiently regulated growth of data centers and resource-intensive processes can exacerbate global warming and ecological degradation. Our literature review focuses on exploring the following two threats that are of immediate concern: The increasing energy demand associated with AI data centers, which contributes to growing carbon emissions. The strain on water resources necessary for cooling the vast infrastructure that powers AI systems.

The energy and water demand are categorized as scope 1, 2, and 3 emissions. The Greenhouse Gas protocol refers to the categorization of companies emissions [4]. Scope 1 refers to direct emissions from company-controlled sources, such as on-site electricity generators and heating/cooling systems; scope 2 describes the indirect emissions from purchased energy, including energy used from the power plants that data centers rely on; scope 3 includes other indirect emissions in the AI lifecycle. For example: electricity used to manufacture semiconductors or chips, as well as energy used to extract raw materials for manufacturing. These scopes help categorize the emissions footprint of AI systems, but they only begin to capture the scale of their environmental burden.

#### 3 Methods

To gather research for this literature review, we used a plethora of resources to help us understand the water and energy usage associated with artificial intelligence. Our primary method involved searching peer-reviewed academic databases, including but not limited to Google Scholar, JSTOR, and the Claremont Colleges Library A Review of the Existing Literature on AI's Energy and Water Implications

for relevant articles. Additionally, we were recommended other academic articles from our advisor, Dr. Lynn Kirabo, and our liaison, Dr. Pablo Ortiz-Partida, who are both climate science experts. They recommended research that they believed would be most informative in guiding us towards our goal of understanding water and energy usage of AI. Their insights as experts in the field helped us identify other key studies and emerging research. Our search strategy was built around keyword combinations including: AI, Artificial Intelligence, Water Usage, Energy Consumption, Environmental Impact, Data Centers, and Sustainability. These keywords were changed based on the specific focus of each database or platform. To ensure the validity and relevance of our sources, we prioritized documents and data published within the last ten years. Given the fast-paced advancements in AI models and the associated increase in computational demand, it was critical to source the most up-to-date research. Older sources were only included if they provided context necessary for understanding the progression of AI's environmental footprint. This approach ensured that our literature review drew from a credible and up-to-date evidence base.

## 4 Analysis

In order to analyze our collected resources, we had at least two people read each article, to get multiple perspectives on one article. This gave us a solid base of notes to work with, and a working knowledge of the subject. After our initial research, we also began summarizing each article we read in 1-2 main points every week, to get an overview of them. Using these notes, we created four categories based on the reoccurring subjects: Energy, Water, Data Centers, and Solutions. "Energy" covered any notes related specifically to the effect AI development would have on energy usage, or any specific numbers on AI's energy usage. Similarly, "Water" covered anything related to the water usage of AI, and what effect AI development might have on AI's water usage. "Data Centers" covered anything related to the functioning or growth of data centers not already covered under the first two categories, and "Solutions" covered any proposed changes or suggestions that might reduce AI's energy or water usage.

Using these categories, we labeled each paper: if a paper discussed one of the categories, it received that label. Two people looked at every article, in order to account for differences in perspective when labeling. We then compared the results from each person, discussed the differences, and came to a conclusion on the final labeling. Using this data, we counted the number of papers discussing each subject, and where the most and least overlap between categories was. This gave us a numerical way of comparing categories, to see which areas the current research focused on, and which areas got less attention.

Lastly, we developed insights from the articles. To do this, we summarized each article in one concise main point. We then used affinity diagramming to organize the main points of each article in order to understand general themes within the research. This organization allowed us to cluster the articles into groups centered around specific ideas, and highlighted patterns within the current discussion on AI. The themes that emerged from this analysis form our results.

Journal Subject	Papers
Computer Science	8
STS	7
Energy	4
Environmental Science	3
General Science	2
Water Industry	2
Other	4

Table 1: Most papers came from journals focused on some element of technology, with environmental science and energy concerns coming second.

#### 5 Results

## 5.1 Overall Statistics

In our review, we examined 31 articles.

*Publication Date.* We analyzed the number of publications surrounding the environmental impact of AI over the period of 2006-2025. We observed that there was a significant increase in research on AI's environmental impact starting 2020. We attribute this rise in the number of articles to the recent boom of generative AI and potentially our recency bias in the process of gathering relevant publications.

*Journal Subject Matter.* Among the journals publishing this research, the most common focus was computer science, followed by science, technology, and society (STS). The next most common disciplines were energy and environmental science. Two journals fell into the category of general science, and two focused on the water industry.

*Labels and Overlap.* In our four categories, we found 12 pertaining to water, 19 pertaining to energy, 13 pertaining to data centers, 15 pertaining to solutions, and 4 that pertained to something outside of our four labels. We found that the most common topic brought up in our articles was energy, followed by solutions. Similarly, energy had the most overlap, with 10 articles that discussed both data centers and energy, and 10 articles that discussed both energy and solutions. Of the 27 articles not classified as "other", all but one discussed one of either energy or water, and no article discussed only data centers alone.

#### 5.2 Qualitative

#### 5.2.1 Current State of AI Research.

*AI's Unsustainable Development.* AI's rapid growth may present further challenges to sustainable development. In its current state, AI development consumes too many resources at too high a rate to be sustainable [11]. Therefore, in any discussion of AI's potential to improve sustainability, we must also discuss the sustainability of AI itself [37]. In order to build the devices and chips that run AI, critical and limited minerals must be extracted from the earth in large quantities, straining existing resources [8]. Furthermore, as AI develops, it requires additional compute, which in turn requires more energy. Areas where data centers and other computing facilities were increasing also saw the fastest increase in energy demand [2].



Figure 1: Most papers discussed energy, while about half discussed solutions.



Figure 2: Energy had the most overlap with other topics.

Using renewable energy sources could mitigate the environmental impact of that demand, but this comes with its own challenges: the rising demand for energy, combined with the inherent traits of renewable energy (for example, the intermittent nature of wind and solar energy generation), make it harder to meet sustainability goals [14]. AI also consumes significant amounts of water: for example, ChatGPT 3.0 requires about 500mL of water per medium length response [12, 23]. While most of the research acknowledges that there exist some potential solutions to this consumption, implementing any of these may be challenging because of the additional investment and time required [12]. How much AI increases energy consumption will depend on what choices consumers and AI developers make [9], but environmentally friendly options may not be feasible in some areas or for some companies, or may face resistance due to the money and effort required [12]. Furthermore, because development is a global process, the consequences of AI development must be addressed globally [11].

*AI's Potential Environmental Good.* However, AI does not need to be environmentally costly, and some argue for the environmental good that AI could do. Due to improvements in machine learning and energy efficiency, the carbon footprint of machine learning may decrease over time, and current predictions of machine learning's

potential carbon footprint may be overestimations because they do not account for these improvements [29]. Other areas where AI could help improve sustainability include water allocation and efficiency [10] and energy efficiency in buildings [40]. AI can also be used to enable new research in energy [18]. This potential to improve sustainability exists alongside AI's current unsustainability. AI can be used in sustainable development, but this requires further resources and better AI models [20]. For example, implementing certain best practices in AI model development makes it more likely that AI's carbon footprint will go down over time, instead of being outpaced by AI's growth [29].

5.2.2 Directly Implemented Solutions. AI's energy and water consumption is directly tied to the energy and water consumption of the data centers that run their computations. Therefore, by reducing the energy and water consumption of data centers, we can reduce the environmental impact of AI. Suggestions for how to achieve this fell into three main categories: changing the algorithms so the calls require fewer resources, changing the data center locations and operating times so they consume fewer resources, and changing the cooling systems of data centers to consume fewer resources.

Algorithmic and Training Modifications. By altering when and how we train AI models, we can reduce their energy consumption. Training models during times when the demands of data centers generate the least amount of carbon can reduce the carbon emissions of AI [16]. We can also adjust how much traffic goes to specific data centers and the order in which bandwidth is allocated to data center jobs to save power, reducing the environmental impact [38]. If researchers take into account the energy consumption of their models while training and building them, this will also allow for better, more energy-efficient models [34].

Location of Data Centers. The location of data centers also plays a heavy part in their environmental impact, so their environmental cost can be mitigated by choosing data center locations with this in mind. Data centers' water consumption can be reduced by accounting for where and when water is used: for example, a data center running in a cooler location or at night will need less water for cooling [30]. Moving data centers to areas that are less water-stressed can also reduce the impact of those center on their surrounding environment [32]. Data centers' carbon emissions are also influenced by their location. Moving data centers to electricity grids that draw from renewable power means that the data center is more likely to be environmentally friendly. Still, simply putting data centers in grids that use more renewable energy does not guarantee lower carbon emissions, and data centers must account for the operational limits of their power system in order to reduce emissions [1]. Furthermore, location and time do not necessarily guarantee better outcomes on all fronts: for example, a data center running at night may be able to use less water, but will not be able to make use of solar power as effectively [23]. Therefore, changing data centers' location in order to reduce their emission requires taking into account a number of different factors.

*Cooling*. Another way of reducing data centers' resource consumption is to improve their cooling systems, as cooling data centers also requires a significant amount of water and energy. Improving a data center's cooling system reduces its energy consumption [19], and using different cooling methods can save both energy and water [33]. There are a few ways to improve on a data center's cooling system: for example, by using liquid cooling, which improves on air cooling due to liquid's larger heat capacity [25]. Another method of improving involves air-side free cooling, and using the outside environmental conditions to reduce the energy and water needed [33]. These solutions also become more effective given the right weather and temperature conditions [33], which ties back into the importance of a data center's location.

#### 5.2.3 Areas of Further Discussion.

New Ethical Frameworks. Many of the ethical frameworks emerging in response to AI's environmental impact focus on the impact AI has on communities, and highlight underrepresented ways of discussing the topic. Currently, discussion of AI ethics centers around the inevitability of AI, and what AI will become: for example, people focus on the possibility of developing sentient AI, and the ethical concerns involved in that. However, focusing so much of the energy around AI ethics on this kind of discussion takes attention away from other forms of climate discussion, and the ethical consequences of AI's current environmental harm [41]. One proposed method is to look at the "Good Way" of building things, a Lakota concept that lays out all the areas affected by the building of something new [22]. This concept could serve as a framework for building ethical AI, including taking into account environmental concerns. Indigenous understandings of consciousness are often undervalued [24], so these understandings could also provide another way of looking at AI. Furthermore, because the impact of AI are so widespread, AI ethics requires a focus on bigger-picture systemic analysis [6]. For similar reasons, AI's environmental harms also require a new ethical framework that focus on communities, and the impact AI's growth has on them [21].

Necessity of Further Reporting. Lack of transparency is a crucial concern in this field. In our research, we found multiple papers calling attention to technology corporations not publishing full information on their AI models. Neither scope-1 nor scope-2 water usage is included in AI's model cards, which hinders the efforts of creating accurate optimization strategies for water efficiency. [23] Further research is needed to quantify the impacts of AI over its entire lifecycle with greater accuracy, and critically, the lack of available and detailed information for that research prevents the creation of effective policy to address the environmental cost of AI [11]. Similarly, researchers can build more energy-efficient models, but in order to do so they need more information about the current energy consumption of AI [34, 35]. Implementing any of the more hardware-focused solutions, like implementing better cooling, would also require better metrics and analysis [17]. In order for data centers to improve statistics, like water efficiency, they must report it [27]. Furthermore, additional information would also allow consumers to better assess and compare models for computational efficiency, providing an economic avenue with which to push AI towards sustainability [34]. Currently, what we can observe of AI's environmental impact is incomplete, and requires us to look at data outside of what companies report, such as the satellite data around data center locations [7].

#### 6 Discussion

Research on the water and energy cost of AI has grown alongside its rapid development over the past few years. In this section, we discuss the implications of our literature review, present limitations of our results, and discuss future pathways of research in this field.

Balancing Technological Innovation with Environmental Sustainability. With the unsustainable growth of AI, it is questionable whether the supposed benefits of AI outweigh its damage to the environment. AI is claimed to have many benefits, including being used as an efficient mechanism to improve sustainability in various industries. However, while some sources are optimistic that carbon emissions from model training will eventually decrease [29], other sources argue that, given the significant environmental cost of AI, societies should choose carefully whether certain AI models need to be developed or used at all [37]. Despite the potential good that AI can do to the environment through optimizing various sustainability measures [20], such benefits are justified only when the AI itself is sustainable.

More Focus Needed on Sustainability in AI Research and Development. As we conducted our literature review, we realized that research on the sustainability of AI has yet to catch up with the growth of AI itself. Most AI researchers and companies do not report their specific energy and water consumption, forcing research on the energy and water implications of AI to draw from indirect data and resort to careful estimations [23, 32]. Even then, not all stages of AI's energy and water consumption can be calculated. Most of the research that attempts to quantify the energy and water consumption of AI focus on such consumption within data centers. The energy and water used in other steps of AI development, e.g. resource extraction, manufacturing, and disposal of AI hardware, is largely unaccounted for when calculating the energy and water cost of AI models [11, 23]. To further this area of research, it is important for AI researchers and companies to share more information on their AI models pertaining to their environmental cost, including the specific energy and water usage at each step of AI development. A few sources have pointed out that AI research and development tends to prioritize accuracy of AI models over their environmental cost [7, 35]. In the AI industry, the drive for profit largely outweighs the considerations for the environmental impact. While attempts to improve efficiency of AI and data centers exist, they are not enough to solve the problem [5]. With the growing demand of AI, companies might use the savings to further increase production of AI according to the Jevons Paradox, and drive up the total energy and water use of AI [13, 39].

*Environmental Regulations Needed to Hold AI Companies Accountable.* It is imperative to hold AI companies accountable for the environmental consequences of their development of AI. To achieve this goal, new regulatory frameworks need to be established. Currently, there are limited regulations around the environmental impact of AI, and those that exist are not strictly enforced [11]. To protect the environment from harm caused by AI's extractive use of water and energy, it is critical to understand the experiences and needs of communities directly impacted by the growth of AI [21]. Therefore, we first suggest that new policies should evaluate AI's energy and water impact on local communities, particularly the electricity grid

and local water supply. Second, policies should enforce reporting of the energy and water consumption of AI models made by AI companies, including specific energy and water cost in each step of AI development. Third, we suggest placing limits on the energy and water usage of data centers to a level that is evaluated to be sustainable.

#### 6.1 Limitations and Future Work

In our literature review, we focused on the energy and water consumption of AI. Nevertheless, the development of AI affects the environment through other pathways unaccounted for in this report, including but not limited to the ecological damage caused by mining materials for AI hardware [40] and degradation of air quality associated with data center operations [15]. Due to the focus on energy and water, there are also publications about AI's environmental impact in general that we did not include in our literature analysis but still helped form our opinions about sustainable AI [3, 5, 13, 39]. The majority of the resources we analyzed focused on energy and water used during the training and inference stages of AI development. Research on the energy and water cost of AI in other stages remains lacking. Another pathway for future research is to analyze how the different types of environmental impact are distributed across different countries and regions. Finally, while broad overviews of the environmental damage across the globe exist, more detailed cost-benefit analyses would be greatly helpful [11].

#### Acknowledgments

We thank Sheridan Dorsey and the Harvey Mudd Writing Center for their assistance in the writing process. We would also like to thank Morgan McArdle and the rest of Harvey Mudd's Computer Science Clinic team for their help facilitating this program.

#### References

- Yasmine Abdennadher, Julia Lindberg, Bernard C Lesieutre, and Line Roald. 2022. Carbon Efficient Placement of Data Center Locations. In 2022 North American Power Symposium (NAPS). 1–6. doi:10.1109/NAPS56150.2022.10012198
- U.S. Energy Information Administration. 2024. Commercial electricity demand grew fastest in states with rapid computing facility growth. https://www.eia. gov/todayinenergy/detail.php?id=62409
- [3] Peter Dauvergne and. 2022. Is artificial intelligence greening global supply chains? Exposing the political economy of environmental costs. *Review of International Political Economy* 29, 3 (2022), 696–718. doi:10.1080/09692290.2020. 1814381 arXiv:https://doi.org/10.1080/09692290.2020.1814381
- Kara Anderson. 2025. https://greenly.earth/en-us/blog/company-guide/whatis-the-greenhouse-gas-protocol
- [5] Noman Bashir, Priya Donti, James Cuff, Sydney Sroka, Marija Ilic, Vivienne Sze, Christina Delimitrou, and Elsa Olivetti. 2024. The Climate and Sustainability Implications of Generative AI. An MIT Exploration of Generative AI (mar 27 2024). https://mit-genai.pubpub.org/pub/8ulgrckc.
- [6] Larissa Bolte and Aimee van Wynsberghe. 2024. Sustainable AI and the third wave of AI ethics: a structural turn. Partners Universal International Innovation Journal (PUIIJ) (July 2024). doi:10.1007/s43681-024-00522-6
- [7] Srija Chakraborty. 2024. Towards A Comprehensive Assessment of Al's Environmental Impact. arXiv preprint arXiv:2405.14004 (2024).
- [8] Kate Crawford. 2021. Earth. Yale University Press, 23–51. http://www.jstor.org/ stable/j.ctv1ghv45t.4
- [9] Alex de Vries. 2023. The growing energy footprint of artificial intelligence. Joule 7, 10 (2023), 2191–2194. doi:10.1016/j.joule.2023.09.004
- [10] Michael T. Eaton and Stephanie McLain Barker. 2006. current issues: Artificial Intelligence Helps JEA Optimize Water Resources. *Journal (American Water Works Association)* 98, 6 (2006), 40–42. http://www.jstor.org/stable/41312620
- [11] Sophia Falk, Aimee van Wynsberghe, and Lisa Biber-Freudenberger. 2024. The attribution problem of a seemingly intangible industry. Environmental Challenges

16 (2024), 101003. doi:10.1016/j.envc.2024.101003

- [12] A.Shaji George, A.S.Hovan George, and A.S.Gabrio Martin. 2023. The Environmental Impact of AI: A Case Study of Water Consumption by Chat GPT. *Partners Universal International Innovation Journal (PUIIJ)* 01, 02 (April 2023). doi:10.5281/zenodo.7855594
- [13] Mario Giampietro and Kozo Mayumi. 2018. Unraveling the complexity of the Jevons Paradox: The link between innovation, efficiency, and sustainability. *Frontiers in Energy Research* 6 (2018), 349753.
- [14] S&P Global. 2024. Datacenters and energy 101: Powering through renewable intermittency.
- [15] Yuelin Han, Zhifeng Wu, Pengfei Li, Adam Wierman, and Shaolei Ren. 2024. The Unpaid Toll: Quantifying the Public Health Impact of AI. arXiv:2412.06288 [cs.CY] https://arxiv.org/abs/2412.06288
- [16] Akshaya Jagannadharao, Nicole Beckage, Dawn Nafus, and Scott Chamberlin. 2023. Timeshifting strategies for carbon-efficient long-running large language model training. *Innovations in Systems and Software Engineering* (2023), 1–15.
- [17] Coles Jennings. 2017. As Data Center Growth Surges, So Does the Energy Requirement. *The Military Engineer* 109, 710 (2017), 51–53. https://www.jstor. org/stable/26464625
- [18] Donghan Jin, Raffaella Ocone, Kui Jiao, and Jin Xuan. 2020. Energy and AI. Energy and AI 1 (2020), 100002. doi:10.1016/j.egyai.2020.100002
- [19] John Judge, Jack Pouchet, Anand Ekbote, and Sachin Dixit. 2008. Reducing data center energy consumption. Ashrae Journal 50, 11 (2008), 14.
- [20] Arpan Kumar Kar, Shweta Kumari Choudhary, and Vinay Kumar Singh. 2022. How can artificial intelligence impact sustainability: A systematic literature review. *Journal of Cleaner Production* 376 (2022), 134120. doi:10.1016/j.jclepro. 2022.134120
- [21] Sebastián Lehuedé. 2024. An elemental ethics for artificial intelligence: water as resistance within AI's value chain. AI & SOCIETY (2024), 1–14.
- [22] Jason Edward Lewis, Angie Abdilla, Noelani Arista, Kaipulaumakaniolono Baker, Scott Benesiinaabandan, Michelle Brown, Melanie Cheung, Meredith Coleman, Ashley Cordes, Joel Davison, et al. 2020. Indigenous protocol and artificial intelligence position paper. (2020).
- [23] Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. 2025. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. arXiv:2304.03271 [cs.LG] https://arxiv.org/abs/2304.03271
- [24] Radmila Lorencova and Radek Trnka. 2023. Variability in cultural understandings of consciousness: a call for dialogue with native psychologies. *Journal of Consciousness Studies* 30, 5-6 (2023), 232–254.
- [25] G. I. Meijer. 2010. Cooling Energy-Hungry Data Centers. Science 328, 5976 (2010), 318–319. http://www.jstor.org/stable/40544557
- [26] Tim Mucci. 2024. https://www.ibm.com/think/topics/history-of-artificialintelligence
- [27] David Mytton. 2021. Data centre water consumption. npj Clean Water 4, 1 (2021), 11.
- [28] Sarah O'Neill. 2023. https://www.lxahub.com/stories/the-history-of-openai
- [29] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. 2022. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Computer* 55, 7 (2022), 18–28. doi:10.1109/MC.2022.3148714
- [30] Shaolei Ren. 2013. Optimizing Water Efficiency in Distributed Data Centers. In 2013 International Conference on Cloud and Green Computing. 68–75. doi:10.1109/ CGC.2013.19
- [31] Shaolei Ren. 2023. How much water does AI consume? The public deserves to know. https://oecd.ai/en/wonk/how-much-water-does-ai-consume
- [32] 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.
- [33] Luis Silva-Llanca, Carolina Ponce, Elizabeth Bermúdez, Diego Martínez, Andrés J. Díaz, and Fabián Aguirre. 2023. Improving energy and water consumption of a data center via air free-cooling economization: The effect weather on its performance. *Energy Conversion and Management* 292 (2023), 117344. doi:10. 1016/j.enconman.2023.117344
- [34] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and Policy Considerations for Deep Learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Anna Korhonen, David Traum, and Lluís Marquez (Eds.). Association for Computational Linguistics, Florence, Italy, 3645–3650. doi:10.18653/v1/P19-1355
- [35] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2020. Energy and policy considerations for modern deep learning research. In *Proceedings of the* AAAI conference on artificial intelligence, Vol. 34. 13693–13696.
- [36] Alan M. Turing. 1950. COMPUTING MACHINERY AND INTELLIGENCE. Mind LIX, 236 (October 1950), 433–460. https://doi.org/10.1093/mind/LIX.236.433
- [37] Aimee Van Wynsberghe. 2021. Sustainable AI: AI for sustainability and the sustainability of AI. AI and Ethics 1, 3 (2021), 213–218.
- [38] Ting Wang, Yu Xia, Jogesh Muppala, and Mounir Hamdi. 2018. Achieving Energy Efficiency in Data Centers Using an Artificial Intelligence Abstraction Model. IEEE Transactions on Cloud Computing 6, 3 (2018), 612–624. doi:10.1109/TCC.

A Review of the Existing Literature on AI's Energy and Water Implications

2015.2511720

- [39] Dustin Wright, Christian Igel, Gabrielle Samuel, and Raghavendra Selvan. 2023. Efficiency is not enough: A critical perspective of environmentally sustainable AI. arXiv preprint arXiv:2309.02065 (2023).
  [40] Biao Yan, Fei Hao, and Xi Meng. 2021. When artificial intelligence meets building energy efficiency, a review focusing on zero energy building. Artificial Intelligence

*Review* 54, 3 (2021), 2193–2220.
[41] Nicolas Zehner and André Ullrich. 2024. Dreaming of AI: environmental sustainability and the promise of participation. *AI & SOCIETY* (2024), 1–13.

Received 25 April 2025