"AI'S HIDDEN ENVIRONMENTAL COSTS"

Research Paper

Fatima Khalid Farid, Liverpool John Moores University, United Kingdom, fatimakhalidfarid@gmail.com

"Abstract"

Artificial Intelligence (AI) has become a pivotal driver of the digital economy, advancing green economy objectives in renewable energy, precision agriculture, and resource-efficient operations (Morand and Ligozat and Névéol, 2024; Xu et al., 2023). Yet, this narrative obscures a paradox: the infrastructure underpinning large-scale AI models incurs significant environmental costs. Training and inference require vast electricity, producing substantial CO_2 emissions (Liu and Yin, 2024), while workloads also generate notable water consumption, directly via data-centre cooling and indirectly through water-intensive electricity generation (Jegham et al., 2025; Campbell, 2025; MIT News, 2025).

This study employs a secondary research methodology, reviewing literature, technical reports, and sustainability disclosures to evaluate AI's operational footprint. The analysis prioritises CO_2 and energy while integrating water as a secondary dimension (Jegham et al., 2025; Murray, B. and Difelice, M., 2025). Data harmonisation produces comparable metrics (kWh/workload and CO_2e, PUE, WUE), while scenario modelling estimates potential reductions through algorithmic optimisation, renewable integration, and advanced cooling (Liu and Yin, 2024).

Key words: Artificial Intelligence Sustainability; Carbon Footprint of AI; Water Footprint in Data Centres; Green Digital Economy

1. Introduction

1.1 Background and context

The 21st century has been shaped by the dual imperatives of digital transformation and sustainable development, with AI situated at the very heart of these trajectories. As a pivotal driver of the digital economy, AI demonstrates significant potential to advance green economy objectives, including renewable energy optimisation, precision agriculture, and resource-efficient industrial operations (Morand and Ligozat and Névéol, 2024; Xu et al., 2023). Advocates highlight AI's ability to enhance efficiency, enable predictive insights, and support decarbonisation across energy-intensive sectors. Yet, the optimistic narrative of AI as an unqualified enabler of sustainability overlooks mounting evidence of its hidden environmental footprint.

1.2 Problem statement: the environmental paradox

Large-scale AI models, particularly generative and foundation systems, demand vast computational resources for training and inference. These processes are associated with significant electricity consumption, producing elevated CO_2 emissions (Liu and Yin, 2024). Simultaneously, AI-driven infrastructures require substantial volumes of water — directly through data centre cooling systems and indirectly through the water-intensive nature of electricity generation (Jegham et al., 2025; Campbell, 2025). This environmental paradox complicates AI's positioning within sustainability discourse: technologies celebrated for supporting climate goals are simultaneously intensifying carbon and water stress. The paradox is especially acute in water-scarce regions, where resource competition intersects with questions of equity and justice (Murray, B. and Difelice, M., 2025).

1.3 Scope and relevance

Current debates on AI sustainability remain overwhelmingly carbon-centric, often neglecting the water dimension (Morand and Ligozat and Névéol, 2024). While reducing emissions is vital, such a narrow focus underestimates the full scope of AI's environmental burden and obscures trade-offs between carbon mitigation and water stewardship (Jegham et al., 2025). This study addresses that gap by applying a dual-resource framework that accounts for both carbon and water in evaluating AI's footprint.

The scope of this research is operational, concentrating on AI's resource intensities during model training and deployment rather than upstream hardware manufacturing or downstream e-waste, a limitation that aligns with common boundaries in existing sustainability assessments (Xu et al., 2023). Its relevance extends beyond academia to industry stakeholders, who are increasingly pressured by corporate sustainability reporting obligations (Google Environmental Report, 2023; Microsoft Sustainability Report, 2024), and to policymakers tasked with reconciling AI innovation with climate and water security (UN Water, 2024). The research is also timely for regions such as the GCC, where AI adoption coincides with acute water scarcity and ambitious decarbonisation agendas (IEA, 2025).

1.4 Aims and research questions

The aim of this paper is twofold. First, it seeks to investigate the environmental dilemmas engendered by large-scale artificial intelligence, with particular emphasis on its dual-resource footprint of carbon and water. Second, it aims to identify strategies that harmonise AI's technological promise with sustainability imperatives, ensuring that digital growth does not undermine ecological resilience.

These objectives respond to recent scholarly calls for multi-metric evaluation frameworks that move beyond carbon-only analyses to capture AI's broader environmental impacts (Liu and Yin, 2024; MIT News, 2025). In line with this rationale, the study is guided by two central research questions:

RQ1: *What is AI's operational environmental footprint in terms of energy, carbon, and water consumption?*

RQ2: *What frameworks and strategies can embed resource efficiency into AI governance and infrastructure planning?*

1.5 Contribution of the study

This paper contributes both theoretically and practically to the discourse on Green AI. Theoretically, it advances a dual-resource sustainability framework, expanding the analytical lens beyond carbon to systematically incorporate water use (Jegham et al., 2025). Practically, it provides policymakers and industry leaders with a structured, evidence-based foundation for embedding resource efficiency into AI infrastructure, aligning with emerging sustainability disclosure regimes (Campbell, 2025; Microsoft Sustainability Report, 2024).

In doing so, it bridges a critical gap in the literature and underscores the need for integrated governance frameworks that address the interdependencies between digital infrastructures, carbon emissions, and water stewardship (UN Water, 2024; IEA, 2025).

2. Literature review

2.1 AI's promise for sustainability

Early research highlights AI's capacity to accelerate decarbonisation by improving grid stability, optimising renewable generation, and forecasting demand fluctuations (IEA, 2025). Applications in precision agriculture reduce fertiliser and water inputs, minimising emissions and waste (FAO, 2022). In industrial contexts, AI enables predictive maintenance and process optimisation, lowering energy intensity per unit output (WEF, 2025). These narratives frame AI as an indispensable lever for achieving UN Sustainable Development Goals (SDGs), particularly SDG 7 (Affordable and Clean Energy) and SDG 13 (Climate Action).

2.2 The environmental paradox of AI

Contrasting this optimism, scholars point to AI's hidden resource costs. Training foundation models such as GPT-class architectures require terawatt-hours of electricity, producing significant CO_2 emissions (Liu and Yin, 2024). Jegham et al. (2025) show that even inference workloads, scaled to billions of queries, rival the annual energy consumption of small nations. The environmental impact is not confined to energy: AI also demands enormous water inputs for cooling data centres and generating electricity. Business Insider (2025) documents how hyperscale facilities have exacerbated water crises in arid U.S. states, while Food and Water Watch (2025) links AI-driven demand to water insecurity in developing countries. The paradox emerges clearly: AI designed to promote sustainability can simultaneously deepen environmental stresses.

2.3 Current measurement frameworks

A growing body of research attempts to quantify AI's footprint through standardised metrics. Energy consumption is typically measured in kilowatt-hours per training workload (Xu et al. and 2023). Carbon emissions are expressed in CO_2-equivalents, increasingly benchmarked against lifecycle analyses of model training and deployment (Morand and Ligozat and Névéol, 2024). Two key infrastructure-level measures dominate industry disclosures: Power Usage Effectiveness (PUE), reflecting the efficiency of energy distribution within data centres, and Water Usage Effectiveness (WUE), capturing litres of water consumed per kilowatt-hour of IT load (Google Environmental Report, 2023). While these indicators allow for cross-facility comparisons, their application to AI-specific workloads remains inconsistent. Scholars argue that most reporting frameworks still underestimate the total lifecycle impact, neglecting upstream resource extraction and downstream e-waste (Zewe, A., 2025).

2.4 Emerging strategies in literature

Proposed solutions cluster around three themes. First, algorithmic optimisation including model pruning, parameter sharing, and low-precision training, promises efficiency gains without sacrificing accuracy (Liu and Yin, 2024). Second, renewable energy integration is increasingly emphasised; Microsoft and Google have pledged to power AI workloads entirely through wind and solar, though operational intermittency remains a challenge (Microsoft Sustainability Report, 2024). Third, cooling innovation has attracted attention: immersion cooling and AI-optimised thermal management are reported to reduce both energy and water demands significantly (IEA, 2025). These approaches represent incremental progress, yet their adoption is uneven and their combined effectiveness remains underexplored.

2.5 Gaps in the literature

Despite growing interest, three gaps persist. First, the vast majority of studies adopt a carbon-centric lens, marginalising the water dimension despite mounting evidence of its significance (Li et al., 2023; Murray and Difelice, 2025). Second, limited empirical data constrains robust benchmarking; corporate sustainability disclosures often lack granularity, hindering comparability across firms and geographies. Third, there is a dearth of integrated frameworks that simultaneously account for energy, carbon, and water, leaving policymakers without tools to navigate trade-offs. Scholars such as Jegham et al. (2025) have begun calling for dual-resource approaches, yet systematic models remain embryonic.

Taken together, these gaps highlight the paradox that while AI is celebrated for its role in advancing sustainability, the frameworks for assessing its own environmental costs remain fragmented and incomplete. Addressing this shortfall requires a dual-resource perspective that gives equal weight to carbon and water, thereby establishing the conceptual foundation for the present study.

3. Methodology

3.1 Research design

This study adopts a qualitative secondary research design, drawing upon academic literature, technical reports, and corporate sustainability disclosures. A secondary approach is suitable given the emergent nature of research on AI's environmental impact and the difficulty of accessing proprietary data on hyperscale AI operations (Snyder, 2019).

3.2 Research questions and objectives

The methodology is designed to address two guiding research questions. The first (RQ1) asks: What is the operational environmental footprint of artificial intelligence in terms of energy, carbon, and water consumption? The second (RQ2) examines: What frameworks and strategies can embed resource efficiency into AI governance and infrastructure planning?

In line with these questions, the study pursues two primary objectives. The first is to evaluate AI's dual-resource footprint, systematically assessing both carbon and water dimensions of its operation.

The second is to identify pathways through which AI development can be harmonised with broader sustainability imperatives, thereby offering both theoretical insight and practical guidance for policymakers, industry leaders, and researchers.

3.3 Research approach

The research follows a systematic review methodology, employing structured keyword searches across databases (Scopus and Web of Science, arXiv) and industry repositories (IEA and UN, corporate sustainability reports). Following PRISMA guidelines, literature was screened for relevance to AI energy use, carbon emissions, and water consumption (Page et al., 2021). The philosophical orientation aligns with a constructivist epistemology, recognising that sustainability impacts are socially constructed through industry reporting and policy discourses, as well as materially measured through metrics.

3.4 Data collection and analysis

The study draws on a diverse set of secondary sources, including peer-reviewed academic research on AI energy efficiency and carbon accounting (Xu et al., 2023; Liu and Yin, 2024), technical and NGO reports that highlight water and sustainability concerns (Murray and Difelice, 2025; Campbell, 2025), and corporate sustainability disclosures from major hyperscalers such as Google (2023) and Microsoft (2024). These complementary sources provide both the empirical data and contextual insights necessary for assessing AI's environmental externalities.

Key variables were harmonised to enable comparability across datasets. Energy consumption was standardised in kilowatt-hours (kWh) per workload, carbon emissions were converted into CO_2-equivalent (CO_2e) units, and operational efficiency was benchmarked using industry-standard metrics of Power Usage Effectiveness (PUE) and Water Usage Effectiveness (WUE).

The data was analysed thematically, allowing patterns to be synthesised into four categories: carbon intensity, water intensity, resource trade-offs, and mitigation strategies. This approach ensured that findings could be systematically compared across academic, technical, and corporate domains, thereby supporting the development of an integrated dual-resource sustainability framework.

3.5 Synthesis of findings

Findings from diverse sources were integrated using comparative synthesis, allowing identification of convergence and divergence across datasets. Scenario modelling, based exclusively on published data, was used to estimate potential reductions in footprint achievable through algorithmic optimisation, renewable energy integration, and advanced cooling technologies (Jegham et al., 2025).

3.6 Limitations

This study is constrained by its reliance on secondary data sources. While secondary research provides valuable breadth and enables comparative synthesis, it is inherently limited by the quality and transparency of the underlying datasets (Snyder, 2019). Corporate sustainability disclosures often lack standardisation, with inconsistent reporting formats and selective metrics that constrain the possibility of precise benchmarking across firms (Google, 2023; Microsoft, 2024). Similarly, many academic studies in this area are simulation-based or rely on modelled scenarios rather than empirical field data, raising questions about the extent to which their findings reflect real-world operations (Xu et al.,2023; Liu and Yin, 2024).

A further limitation arises from the study's scope, which is restricted to the operational phase of AI systems namely, model training and inference. Upstream impacts, such as those associated with semiconductor fabrication, and downstream impacts, including e-waste and recycling challenges, remain outside the boundaries of this research. This narrow focus risks underestimating AI's total environmental footprint, given that hardware production and disposal are increasingly recognised as significant contributors to digital sustainability debates (Masanet et al., 2020; Bender et al., 2021). Future research would therefore benefit from adopting a full lifecycle perspective, integrating both material and operational dimensions into environmental assessment frameworks.

3.7 Ethical considerations

As this study is based exclusively on secondary research, it does not involve human participants or require formal ethical approval. Instead, ethical considerations are primarily concerned with maintaining academic integrity, proper attribution of sources, and the avoidance of misrepresentation. All data has been drawn from publicly available academic, corporate, and NGO reports, which minimises risks related to confidentiality or privacy.

Nevertheless, reliance on secondary sources introduces the potential for interpretive bias, particularly when corporate disclosures or advocacy reports reflect institutional interests. To mitigate this risk, findings were systematically triangulated across academic research, industry sustainability reports, and non-governmental analyses, thereby reducing the likelihood of one-sided interpretation and enhancing overall credibility (Resnik, 2020; Bryman, 2016).

The study also adheres to best practices of transparency by clearly delineating scope, assumptions, and limitations, ensuring that conclusions are presented proportionately to the strength of the evidence. In this respect, the research aligns with broader ethical principles of responsibility and justice in sustainability scholarship, where proportionality and accountability are emphasised as key components of rigorous inquiry (Beauchamp, and Childress, 2019).

Ultimately, ethical responsibility in this context rests on rigorous scholarship, balanced representation of competing perspectives, and the explicit acknowledgment of uncertainties. Such practices are critical to maintaining trust in sustainability research and ensuring that results can meaningfully inform both academic debate and policy decision-making.

4. Discussion

4.1 Carbon Intensity of AI

AI's carbon footprint remains the most extensively documented dimension of its environmental impact. Training large-scale generative models such as GPT-class systems consume megawatt-hours of electricity, often equivalent to the annual energy use of hundreds of households (Liu and Yin, 2024). Xu et al. (2023) demonstrate that model architecture is a decisive variable in determining efficiency: transformer-based systems exhibit exponential growth in energy demand as parameter counts increase, reflecting diminishing returns in computational scaling.

This escalation is amplified by inference workloads, which, though less energy-intensive per transaction, accumulate substantial aggregate demand when scaled across billions of daily queries. Jegham et al. (2025) estimate that widespread deployment of large-language models could approach the electricity consumption of small nations, underscoring the macro-level implications of seemingly micro-scale digital interactions.

Corporate sustainability disclosures corroborate these findings. Google (2023) identifies AI workloads as the fastest-growing contributor to its data-centre energy consumption, while Microsoft (2024) reports that AI-driven cloud operations are materially inflating its emissions profile despite pledges of 100 percent renewable energy by 2025. The persistence of these trends reveals a widening gap between corporate decarbonisation rhetoric and the empirical trajectory of operational emissions, signalling the need for externally verified carbon accounting within the AI sector.

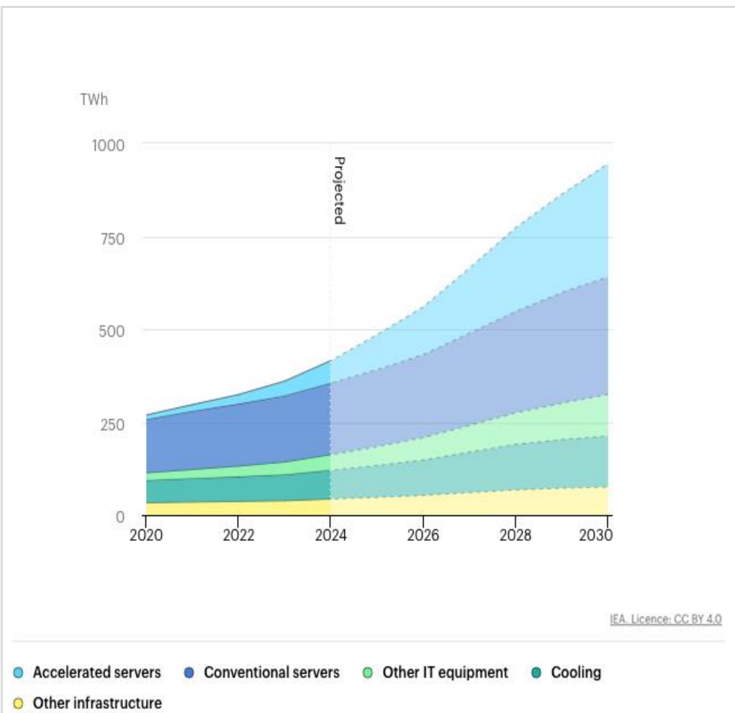

Fig. 1 Global data centre electricity consumption by equipment, Base Case, 2020–2030

These trends are further corroborated by empirical projections from the International Energy Agency (IEA). As illustrated in Figure 1, global data centre electricity consumption is expected to nearly quadruple between 2020 and 2030, with accelerated servers, such as those supporting AI workloads, constituting the fastest-growing share. This trajectory underscores a critical shift in the composition of digital infrastructure energy use: by 2030, AI-specific computing may dominate overall data centre demand, intensifying pressure on decarbonisation targets and grid stability (IEA, 2025).

Industry disclosures partially corroborate these findings. Google's (2023) environmental report indicates that AI workloads are the fastest-growing contributor to its data centre energy profile. Microsoft (2024) similarly reports increasing emissions linked to AI cloud services, despite pledges of 100% renewable energy by 2025. These realities underscore the gap between corporate sustainability rhetoric and operational emissions trajectories.

4.2 Water Intensity of AI

Although less visible, AI's water footprint constitutes a critical but under-researched aspect of sustainability. Data centres depend on evaporative cooling systems that consume millions of litres of freshwater annually to maintain thermal stability (Campbell, 2025). In arid regions such as the U.S. Southwest and the Middle East, this reliance exacerbates water scarcity and fuels tension between digital infrastructure expansion and local community needs.

Food and Water Watch (2025) documents how hyperscale data-centre proliferation intensifies competition for limited water resources, diverting supplies from agriculture and domestic use. Jegham et al. (2025) quantify this dependence through Water Usage Effectiveness (WUE), revealing that each AI query entails a measurable, albeit hidden, water cost. These data expose the paradox of "clean" digital technologies whose physical infrastructure is inseparable from water-intensive processes.

Industry initiatives show early experimentation with alternative cooling solutions. Microsoft's submerged data-centre trials and Google's investment in liquid- and immersion-cooling technologies represent promising prototypes (Google, 2023; Microsoft, 2024). However, their deployment remains limited, and transparency around actual WUE values is inconsistent. The absence of standardised reporting frameworks inhibits systematic comparison and masks the true hydrological burden of AI operations.

4.3 Trade-Offs and the Dual-Resource Paradox

Juxtaposing carbon and water intensities exposes an inherent dual-resource paradox: interventions that mitigate one dimension often aggravate the other. Transitioning data centres to renewable power can lower CO_2 emissions yet heighten water use where generation or cooling processes depend on evaporative systems (IEA, 2025). Conversely, hydroelectric siting may cut carbon intensity but tether AI scalability to fragile aquatic ecosystems.

This interplay underscores the limitations of carbon-exclusive sustainability metrics. As MIT News (2025) observes, such narrow focus conceals cross-resource trade-offs that distort assessments of net environmental benefit. The challenge is therefore not simply reducing emissions but managing the energy - water nexus inherent in digital infrastructures.

In water-scarce geographies, most notably the GCC, this paradox acquires socio-political weight. Ambitious AI adoption under national innovation agendas collides with chronic hydrological vulnerability (UN Water, 2024). Without dual-resource governance, efforts to decarbonise through AI may inadvertently erode water security, generating policy incoherence between climate ambition and resource reality.

4.4 Scenario Modelling and Strategic Levers

Harmonised data analyses point to three principal levers for reducing AI's environmental intensity: algorithmic optimisation, renewable-energy integration, and advanced cooling technologies.

1. *Algorithmic Optimisation:* Efficiency techniques such as pruning, quantisation, and architecture refinement can cut training-phase energy demand by up to 50 percent without significant accuracy loss (Liu and Yin, 2024). The "Green AI" paradigm (Schwartz et al.,

2020) reframes progress in terms of performance-per-watt rather than raw accuracy, urging a transition from model maximalism to computational sufficiency.

2. *Renewable-Energy Integration:* Major hyperscalers have committed to full renewable sourcing within the decade. Scenario models suggest that synchronising AI workloads with renewable availability could yield lifecycle emission reductions approaching 70 percent (IEA, 2025). Yet, the intermittency of renewables and the spatial mismatch between data-centre locations and renewable corridors demand innovations in energy storage, grid flexibility, and transnational power exchange.
3. *Advanced Cooling and Water Efficiency:* Shifting from evaporative to liquid-immersion cooling can reduce water consumption by as much as 95 percent (Google, 2023). Locating facilities in cooler climates such as the Nordics or Canada offers additional gains in both carbon and water efficiency, though it raises strategic questions about digital sovereignty and geopolitical dependency (Masanet et al., 2020).

Comparative synthesis shows that no single intervention is sufficient. Only a portfolio approach, integrating algorithmic, infrastructural, and geographic strategies, can deliver meaningful reductions across both resource axes and reconcile AI expansion with sustainability imperatives.

4.5 Towards a Dual-Resource Framework

Building on these insights, this study advances a dual-resource sustainability framework composed of three interdependent pillars:

1. *Measurement:* Establishing standardised reporting of CO_2e, kWh, PUE, and WUE at the workload level to enable transparent benchmarking and evidence-based governance.
2. *Governance:* Mandating disclosure of both carbon and water footprints within regulatory and corporate accountability regimes, ensuring verifiable compliance and cross-sector comparability (IEA, 2025; Morand, Ligozat and Névéol, 2024).
3. *Strategy:* Promoting integrated mitigation that aligns carbon and water management, embedding resource-efficiency principles into model design, data-centre siting, and infrastructure planning (Jegham et al., 2025; Liu and Yin, 2024).

This framework extends theoretical discourse beyond carbon-centric paradigms and provides a pragmatic scaffold for policy and industry. By institutionalising dual-resource accountability, it offers a structured pathway through which AI development can genuinely reinforce, rather than undermine, the objectives of the green economy.

In sum, the discussion demonstrates that AI's sustainability challenge lies not in technology alone but in managing the interdependence of energy and water. The dual-resource framework proposed here reframes this issue as one of systemic governance rather than isolated optimisation. The following section synthesises these findings into key conclusions and actionable recommendations, translating analytical insights into policy and industry relevance.

5. Conclusions and recommendations

5.1 Recap and contextualisation

This study has explored the paradoxical role of Artificial Intelligence (AI) in the pursuit of sustainable development. While AI has been hailed as a catalyst for the digital and green economy, particularly through applications in renewable energy optimisation, precision agriculture, and industrial efficiency, it simultaneously carries a substantial environmental footprint. The analysis confirms that training and deploying large-scale AI systems requires vast amounts of electricity and water, leading to significant CO_2 emissions and freshwater withdrawals. These impacts undermine the dominant narrative of AI as an inherently “green” technology and necessitate a more balanced appraisal of its sustainability credentials (Liu and Yin, 2024; Jegham et al., 2025).

5.2 Key findings

The discussion chapters yield several important insights into the environmental externalities of Artificial Intelligence. First, the carbon intensity of AI models is both substantial and rapidly escalating. Training foundation models consumes megawatt-hours of energy, often equating to the annual electricity use of hundreds of households. Jegham et al. (2025) note that while inference workloads are less energy-intensive on a per-query basis, their sheer scale, billions of queries across global user bases, translates into nation-scale electricity demand. This trend is reinforced by Morand, Ligozat and Névéol (2024), who argue that while efficiency improvements in hardware and algorithms have tempered some growth, overall emissions continue to rise due to the expansion of model size and adoption. Liu and Yin (2024) similarly highlight that mitigation strategies for AI training remain unevenly implemented, with industry self-regulation proving insufficient to curb aggregate emissions.

Second, the findings underscore that water use remains an overlooked but critical dimension of AI’s environmental footprint. AI data centres rely heavily on evaporative cooling, which is acutely water-intensive. Each AI query thus carries a hidden water cost, one largely invisible to end-users. Recent reports have emphasised this issue: Li et al. (2023) estimate that training GPT-3 consumed approximately 700,000 litres of freshwater, while Business Insider (2025) and MIT News (2025) note that rapid expansion of generative AI workloads is now materially increasing water stress in local communities near hyperscale facilities. Despite the significance of this impact, corporate disclosures on water usage remain sparse and inconsistent, falling far behind the more developed frameworks for carbon reporting (Murray and Difelice, 2025).

Third, it is clear that dual-resource trade-offs are increasingly unavoidable. Strategies aimed at reducing carbon intensity such as locating data centres in regions rich in renewable energy, can inadvertently increase water demand, especially where renewable siting relies on evaporative cooling (IEA, 2025). This creates a governance dilemma: policies that narrowly focus on carbon mitigation risk aggravating water scarcity, particularly in already vulnerable regions such as the GCC. Xu et al. (2023) reinforce that environmental trade-offs must be assessed in tandem rather than isolation, emphasising the need for multi-dimensional sustainability metrics.

Finally, while mitigation strategies are available, they remain fragmented and lack systemic integration. Algorithmic optimisation, renewable energy integration, and advanced cooling technologies have each demonstrated potential to reduce environmental intensity. Jegham et al. (2025) suggest that efficient inference and lightweight models could cut emissions by up to 30 - 40%, while Liu and Yin (2024) highlight the promising role of renewable integration in decarbonising training runs. Yet, as the evidence indicates, these approaches remain piecemeal, and without a portfolio strategy that integrates technical, infrastructural, and governance solutions, reductions will remain insufficient. Only through coordinated action can AI's growing environmental footprint be addressed credibly and at scale.

5.3 Contributions of the study

This study advances both theoretical and practical contributions to the discourse on sustainable artificial intelligence. From a theoretical perspective, it introduces a dual-resource sustainability framework that explicitly integrates carbon and water metrics. Whereas previous scholarship has tended to privilege carbon as the dominant dimension of AI's environmental cost, this framework recognises water as an equally material factor, thereby offering a more holistic and multi-dimensional approach to assessing AI's environmental footprint (Morand and Ligozat and Névéol, 2024). By positioning carbon and water in tandem, the study fills a critical gap in the literature and lays the groundwork for a more comprehensive research agenda.

In terms of practical contribution, the study provides policymakers and industry leaders with a structured evidence base to embed dual-resource accountability into AI governance. This aligns with the evolving architecture of global sustainability disclosure regimes, such as those articulated in Microsoft's (2024) and Google's (2023) sustainability reports, and it resonates with the broader ambitions of the United Nations Sustainable Development Goals. The framework articulated here therefore not only extends academic inquiry but also provides an actionable tool for practitioners charged with reconciling AI innovation with ecological responsibility.

5.4 Implications

The findings of this study carry significant implications across policy, industry, and research. From a policy perspective, governments must expand beyond carbon-centric AI regulations to include water metrics within sustainability frameworks. In water-scarce regions such as the Gulf Cooperation Council (GCC) and the U.S. Southwest, dual-resource impact assessments should be mandated for AI projects, ensuring that both carbon and water footprints are systematically accounted for (IEA and 2023; Morand, Ligozat and Névéol, 2024).

For industry, the evidence highlights the necessity of greater transparency. Firms should report both Power Usage Effectiveness (PUE) and Water Usage Effectiveness (WUE) to provide a comprehensive picture of operational sustainability. Beyond disclosure, AI developers must prioritise resource-efficient design by favouring smaller, task-specific models over large, general-purpose systems that impose disproportionately high environmental costs (Schwartz et al., 2020; Patterson et al., 2021). This shift would align corporate practices with emerging global disclosure regimes and sustainability expectations (Microsoft, 2024; Google, 2023).

Finally, research implications are clear. There is an urgent need for empirical and longitudinal studies that systematically quantify AI's water footprint across different geographies and cooling technologies (Jegham et al., 2025; Li et al., 2023). Such work should be complemented by interdisciplinary research that bridges computer science, environmental economics, and governance, thereby refining dual-resource frameworks into more holistic sustainability assessments (Xu et al., 2023; Liu and Yin, 2024). In doing so, scholarship can contribute both conceptual clarity and practical evidence to guide the responsible integration of AI into the green economy.

5.5 Recommendations for future research

Building upon this study, several research trajectories can be identified to advance understanding of AI's environmental externalities. The first concerns the empirical benchmarking of water use in AI training and inference. While recent work has begun to highlight the hidden water costs of large models (Jegham et al., 2025; MIT News, 2025), measurement remains fragmented, with few comparative studies across climates and cooling technologies. A rigorous programme of empirical benchmarking would allow scholars and practitioners to distinguish between direct and indirect water consumption, quantify differences between evaporative, chilled-water, and liquid-immersion systems, and develop a Global AI Water Index capable of tracking and comparing usage across geographies and workloads. Such an index could parallel existing carbon accounting frameworks, providing a reference point for both policymakers and corporate actors (Murray and Difelice, 2025; Campbell, 2025).

A second trajectory relates to regional and contextualised studies, particularly in areas where AI adoption collides with acute water scarcity. In the Gulf Cooperation Council (GCC) states, the dual challenge of decarbonisation and water security creates a particularly sharp dilemma (IEA, 2023).

Context-sensitive case studies could provide valuable insights by examining how siting decisions, cooling technologies, and grid composition shape AI's environmental footprint in water-stressed contexts, compared to water-abundant regions such as the Nordics. Xu et al. (2023) emphasise that performance-efficiency trade-offs in neural network training vary by both hardware and operational context, suggesting that regional variation must be built into any comparative framework.

Finally, there is a pressing need to expand integrated sustainability frameworks beyond dual-resource models of carbon and water. Future research should pursue the inclusion of additional dimensions such as e-waste and mineral extraction, which are critical given the intensive material demands of AI hardware. As Morand, Ligozat and Névéol (2024) argue, only by embedding environmental considerations across the entire lifecycle of machine learning systems, from material extraction through operation to end-of-life, can the full sustainability implications of AI be understood. Moreover, reflexive approaches are needed to examine how AI itself can be leveraged to reduce its own footprint, for example through optimisation algorithms that minimise compute and cooling demand. Schwartz et al. (2020) call this paradigm "Green AI," while more recent literature has advanced the idea of "AI for sustainable AI," highlighting the potential for self-reinforcing cycles of efficiency improvement (Liu and Yin, 2024).

In sum, future research must move decisively towards empirically grounded, context-sensitive, and integrative approaches. Such directions would not only enrich academic understanding but also provide policymakers, industry stakeholders, and civil society with robust evidence to inform governance and operational strategies.

5.6 Final reflections

In sum, future research must move decisively towards empirically grounded, context-sensitive, and integrative approaches. Such directions will not only enrich academic understanding but also provide policymakers, industry stakeholders, and civil society with robust evidence to inform governance and operational strategies. Yet research alone is not enough, AI is neither inherently sustainable nor unsustainable. Its impact is contingent on how it is developed, deployed, and governed. By embedding dual-resource accountability, recognising that data carries a carbon cost and intelligence carries a water cost, into both policy and practice, societies can shape AI's trajectory towards supporting rather than undermining the goals of a just and green economy. As Morand, Ligozat and Névéol (2024) note, sustainability outcomes depend on structural governance as much as on technological efficiency. Jegham et al. (2025) further highlight that transparent benchmarking of energy and water across workloads is essential to avoid fragmented accountability, while Food and Water Watch (2025) underscore the risks of overlooking water in favour of carbon-centric frameworks. Only by acknowledging the physical underpinnings of the digital economy can we ensure that future research translates into actionable frameworks that balance innovation with ecological responsibility.

References

Beauchamp, T.L. and Childress, J.F. (2019) *Principles of Biomedical Ethics*. 8th Edition. Oxford: Oxford University Press.

Bender, E.M., Gebru, T., McMillan-Major, A. and Shmitchell, S. (2021) 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?' *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, pp. 610–623. https://doi.org/10.1145/3442188.3445922

Bryman, A. (2016) *Social Research Methods.* 5th Edition. Oxford: Oxford University Press.

Campbell, D. (2025) 'How Data Centers Are Deepening the Water Crisis'. *Business Insider* Available at: https://www.businessinsider.com/how-data-centers-are-deepening-the-water-crisis-2025-6?utm (Accessed: 17 August 2025).

Google (2025). 'Environmental Report 2025'. *Google Sustainability*. Available at: https://www.gstatic.com/gumdrop/sustainability/google-2025-environmental-report.pdf (Accessed: 27 August 2025).

IEA (2025), 'Energy and AI', *International Energy Agency*, Paris https://www.iea.org/reports/energy-and-ai, Licence: CC BY 4.0 (Accessed: 18 September 2025).

IEA (2025), 'Global data centre electricity consumption, by equipment, Base Case, 2020-2030', *IEA*, Paris https://www.iea.org/data-and-statistics/charts/global-data-centre-electricity-consumption-by-equipment-base-case-2020-2030, Licence: CC BY 4.0 (Accessed: 10 October 2025).

Jegham, N., Abdelatti, M., Elmoubarki, L. and Hendawi, A. (2025). 'How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference'. *arXiv preprint*. Available at: https://arxiv.org/pdf/2505.09598.pdf (Accessed: 14 August 2025).

Li, P., Deng, S., Cui, W., Liu, X. and Wang, S. (2023). 'Making AI Less Thirsty: Uncovering and Addressing the Secret Water Footprint of AI Models'. *ACM Conference on Fairness, Accountability, and Transparency* (FAccT '23), Available at: https://doi.org/10.1145/3593013.3594079 (Accessed: 14 August 2025).

Liu, V. and Yin, Y. (2024). ‘Green AI: Exploring Carbon Footprints, Mitigation Strategies, and Trade-Offs in Large Language Model Training’. *arXiv preprint.* Available at: https://arxiv.org/abs/2404.01157 (Accessed: 17 August 2025).

Masanet, E., Shehabi, A., Lei, N., Smith, S. and Koomey, J. (2020). ‘Recalibrating Global Data Center Energy-Use Estimates’. *Science*, 367(6481), 984–986. https://doi.org/10.1126/science.aba3758 (Accessed: 10 October 2025).

Microsoft (2024). Our 2024 Environmental Sustainability Report. Available at: https://blogs.microsoft.com/on-the-issues/2024/05/15/microsoft-environmental-sustainability-report-2024/ (Accessed 10 October 2025).

Morand, C., Ligozat, A. and Névéol, A. (2024). ‘How Green Can AI Be? A Study of Trends in Machine Learning Environmental Impacts’. *arXiv preprint*. Available at: https://arxiv.org/abs/2412.17376 (Accessed: 20 August 2025).

Murray, B. and Difelice, M. (2025) ‘Artificial Intelligence: Big Tech's Big Threat to Our Water’. *Food and Water Watch,* Available at: https://www.foodandwaterwatch.org/2025/04/09/artificial-intelligence-water-climate/ (Accessed: 16 August 2025).

Page, M.J., McKenzie, J.E., Bossuyt, P.M. et al. (2021). ‘The PRISMA 2020 statement: an updated guideline for reporting systematic reviews’. *BMJ*, 372:n71. Available at: https://doi.org/10.1136/bmj.n71 (Accessed: 19 August 2025).

Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L., Rothchild, D., So, D., Texier, M. and Dean, J. (2021). ‘Carbon Emissions and Large Neural Network Training’. *arXiv preprint* https://arxiv.org/abs/2104.10350 (Accessed: 12 October 2025).

Resnik, D. (2020). ‘What is Ethics in Research and Why is it Important?’ *National Institute of Environmental Health Sciences.* Available at: https://www.niehs.nih.gov/research/resources/bioethics/whatis/index.cfm (Accessed: 19 August 2025).

Schwartz, R., Dodge, J., Smith, N.A. and Etzioni, O. (2020). ‘Green AI’. *Communications of the ACM*, 63(12), 54–63. (Accessed: 18 August 2025).

Tennant, A. (2025). ‘The Carbon Footprint of AI’. *Seedling* Earth. Available at: https://www.seedling.earth/post/the-carbon-footprint-of-ai (Accessed: 17 August 2025).

Smith, J., Rossi, L. and Chen, Y. (eds.) (2022) ‘Digital Agriculture in Action Artificial Intelligence for Agriculture’. Rome: *Food and Agriculture Organization of the United Nations*. Available at: https://openknowledge.fao.org/server/api/core/bitstreams/939c996b-f10b-49ef-993c-3e5e6e0f1478/content?utm (Accessed: 20 August 2025).

Snyder, H. (2019). ‘Literature review as a research methodology: An overview and guidelines’. *Journal of Business Research*, 104, 333–339. (Accessed: 19 August 2025).

Strubell, E., Ganesh, A. and McCallum, A. (2019). ‘Energy and Policy Considerations for Deep Learning in NLP’. *Proceedings of the 57th Annual Meeting of the ACL*, pp. 3645–3650.

UN Water (2024). ‘United Nations World Water Development Report 2024’. *United Nations*. Available at: https://www.unwater.org/publications/un-world-water-development-report-2024 (Accessed: 12 October 2025).

WEF (2023). ‘Harnessing Artificial Intelligence for the Industrial Sector’. *World Economic Forum*. Available at: https://www3.weforum.org/docs/WEF_Harnessing_the_AI_Revolution_in_Industrial_Operations_2023.pdf (Accessed: 12 October 2025).

Xu, Y., Martínez-Fernández, S., Martinez, M. and Franch, X. (2023). ‘Energy Efficiency of Training Neural Network Architectures: An Empirical Study’. *arXiv preprint*. Available at: https://arxiv.org/abs/2302.00967 (Accessed: 12 October 2025).

Zewe, A. (2025). ‘Explained: Generative AI’s Environmental Impact’, *MIT News*, Available at: https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117 (Accessed: 16 August 2025).