SOCIAL SCIENCE REVIEW ARCHIVES
https://policyjournalofms.com
ISSN Online: 3006-4708
ISSN Print: 3006-4694

AI's Cross-Border Energy and Water Footprints: State Duties, No-Harm Thresholds, and Paris Agreement Compliance Frameworks

Dr. Muhammad Asghar [1], Hazrat Usman [2], Maria Mazhar Kharal [3], Dr. Mohaddas Mehboob[4]

[1] Assistant Professor Law, Department of Law Ibadat, International University Islamabad. Email: muhammad.asghar@law.iiui.edu.pk
[2] Advocate High Court, Pujab Bar Council, Punjab, Pakistan. Email: hazratusmanadvocate@gmail.com
[3] Lecturer, Ibadat International university islamabad, Email- maria.Mazhar@law.iiui.edu.pk
[4] Assistant Professor Department of Law, IBADAT International University Islamabad, Pakistan. mohaddas.naqvi@law.iiui.edu.pk
Corresponding Author: Hazrat Usman

DOI: https://doi.org/10.70670/sra.v3i3.875

Abstract
The world's largest artificial-intelligence models now draw more electricity than some mid-sized nations and evaporate billions of liters of freshwater each year, yet neither international climate law nor the classical transboundary-harm doctrine has fully absorbed their impact. This article conducts a systematic review of empirical footprint studies with a comparative legal analysis of the no-harm principle, emerging corporate-due-diligence statutes, and transparency rules of the Paris Agreement. Lifecycle data show that training a single GPT-3-class model consumes about 1.3 GWh of power and 5.4 million L of water, while global inference loads could withdraw 22 billion L annually by 2027—concentrated in already stressed basins. Because affordable mitigation tools (carbon-aware routing, liquid cooling, and a mixture-of-expert architectures) can reduce these impacts by 40–60 percent, failure to deploy them breaches the due-diligence standard embedded in Trail Smelter and its progeny. This study proposes a hybrid allocation framework that attributes operational footprints to host states but assigns embodied and service-based impacts to consumer states, enabling parties to integrate Scope-3 emissions and virtual-water transfers into Biennial Transparency Reports without amending treaty text. Embedding dual carbon-and-water baselines into Article 6 crediting schemes would channel finance toward low-impact data-centers and close a rapidly widening governance gap.

Keywords: AI Environmental Footprint, Transboundary No-Harm, Virtual Water Accounting, Paris Agreement Compliance, Sustainable Data Centers.

Introduction

Large-scale artificial-intelligence (AI) systems have escaped the confines of research laboratories and now underpin everything from search engines to sovereignty-grade cybersecurity. Their ascent is propelled by hyperscale data-centers, whose hungry servers devour terawatt-hours of electricity and withdraw millions of liters of cooling water every day (UNCTAD, 2024). Recent lifecycle studies place the training of a single GPT-3-sized model at roughly 5.4 million liters of water and 1,287 MWh of electricity (Li et al., 2025); routine inference can eclipse these figures several-fold as usage scales (Chien et al., 2023). In the United States alone, AI-driven demand could lift a grid load of nine percent by 2028 and pull an additional 720 billion gallons from stressed watersheds unless cooling technologies and siting policies change course (Food & Water Watch, 2025).

However, these resources do not respect political borders. Cloud-based inference allows European users to tap GPU clusters in the American desert, exporting virtual-water consumption and upstream greenhouse-gas (GHG) emissions far beyond the territory where legal jurisdiction is clearest (Kseibati, 2025). However, most national climate pledges and corporate net-zero plans still tally only domestic, carbon-centric scopes, overlooking embodied hardware emissions, off-site electricity mixes, and—more glaring still— freshwater withdrawals (GAO, 2025; Sandalow, 2024). The resulting governance lacuna threatens to undermine pari passu achievement of the Paris Agreement's temperature goal and the Sustainable Development Goals' water-security targets.

International environmental law already offers two doctrinal anchor points that could close this gap. First, the no-harm principle obliges states to exercise due diligence and prevent activities within their jurisdiction from causing significant transboundary injury (Mekong River Commission, 2023). Second, the Paris Agreement's transparency framework (Arts. 4, 6 & 13) invites Parties to internalize all climate-relevant externalities—yet its modalities remain carbon-heavy and water-blind. Recent corporate-liability scholarship suggests an emerging transnational duty of care that could make big-tech home states answerable for overseas damage (Gailhofer et al., 2023), while a bipartisan U.S. bill would mandate federal reporting of AI energy and water footprints (U.S. Senate, 2024). Industry, for its part, is converging on standardized metrics—Power Usage Effectiveness, Water Usage Effectiveness, and "Total Source-Energy Water Consumption"—through the Open Compute Project and ISO/IEC 30134 series (Open Compute Project, 2024), but fewer than half of operators track even the mandatory indicators (Uptime Institute, 2024).

Against this backdrop, the present study asks two interlinked questions: (1) Due-diligence and no-harm. How do the cross-border energy and water footprints of hyperscale AI facilities engage states' due-diligence obligations under the no-harm principle, and what threshold of foreseeable, significant harm would trigger responsibility? (2) Paris alignment and disclosure gaps. In what ways can Parties operationalize Articles 4, 6 and 13 of the Paris Agreement to internalize AI-sector externalities—especially embodied (Scope-3) emissions and virtual-water transfers—within NDCs and global stock-take processes?

Answering these questions is important for three reasons: First, AI's resource appetites are accelerating faster than the electricity and water footprints of many traditional heavy industries. Early legal clarity can help avert locks -into unsustainable infrastructure. Second, a principled allocation of transboundary duties would prevent a regulatory race-to-the-bottom in which firms migrate to jurisdictions with lax disclosure. Third, integrating water into paris-aligned accounting would align climate mitigation with planetary-boundary stewardship, bolstering both the legitimacy and efficacy of global climate governance.

The remainder of this paper is organized as follows. Section II maps the existing empirical knowledge on AI-related energy and water use, highlighting gaps in transboundary attribution. Section III evaluates states' due-diligence obligations under the no-harm principle, drawing analogies from transboundary-impact assessments of water-basin treaties and emerging corporate-liability doctrines. Section IV critiques the Paris Agreement's current reporting architecture and proposes a lifecycle-based disclosure protocol that captures embodied emissions and the virtual-water trade. The conclusion distills policy recommendations and flags avenues for future research.

From Trail Smelter to Teraflops: Due Diligence and the No-Harm Principle in AI Resource Flows

The modern law of transboundary harm was forged in smoke that drifted from a Canadian smelter into the wheat fields of Washington State. When the arbitral tribunal in Trail Smelter announced in 1941 that "no state has the right to use or permit the use of its territory in such a manner as to cause injury … in or to the territory of another' (cited in Mekong River Commission, 2023, p. 11), it embedded a due-diligence obligation that would later be crystallized in Principle 21 of the Stockholm Declaration, Principle 2 of Rio, and the no-harm rule of customary international law. Eight decades ago, the fumes in question were no longer sulfurous but digital: electrons coursing through hyperscale data-centers and evaporated water leaving no visible plume. However, the legal architecture for responsibility remains underdeveloped. This part of the review traces how scholarly and policy literature is struggling to extend the Trail-Smelter template to the invisible, border-hopping resource footprints of large AI models, focusing first on the empirical baseline and then on the normative debates about due diligence and significant harm thresholds.

Empirical work over the past three years has converged on the conclusion that AI workloads are shifting global electricity and hydrological balances in ways that policymakers rarely anticipate. UNCTAD's Digital Economy Report 2024 estimates that information and communications technologies already consume between 700 and 1 600 terawatt-hours per year, roughly 1.5 to 3 percent of all global electricity, and that data-center demand could double again by 2026, with generative-AI systems being the principal driver (UNCTAD, 2024). Li, Yang, Islam, and Ren (2025) supply the first comprehensive methodology for scope-1, -2, and -3 **water** accounting, showing that training the GPT-3 model required 5.4 million liters of water and that inference at scale could push global AI-sector withdrawals to between 4.2 and 6.6 billion cubic meters annually by 2027—comparable to Denmark's total national consumption. Chien et al. (2023) and Hoffmann and Majuntke (2024) complement these findings on the energy side, demonstrating that inference, not training, dominates lifecycle emissions for high-traffic services and that carbon-aware request routing or model selection can reduce operational footprints by 35-to-38 percent without appreciable latency penalties. The industry data reveal that relatively few operators track the metrics needed to verify such gains. According to the Uptime Institute's 2024 survey, fewer than half of more than 800 owners and operators measure Water-Usage Effectiveness or greenhouse-gas inventories in a way that would satisfy forthcoming European Corporate Sustainability Reporting Directive (CSRD) standards (Uptime Institute, 2024). The Open Compute Project has tried to fill this vacuum by upgrading WUE and a new "Total Source-Energy Water Consumption" indicator from optional to required status for OCP-Ready recognition (Open Compute Project, 2024), but uptake remains voluntary.

Regulatory and legislative bodies are only beginning to respond to this. The U.S. Artificial Intelligence Environmental Impacts Act (S. 3732) instructs the Environmental Protection Agency to quantify AI's lifecycle effects and directs the National Institute of Standards and Technology to draft open-source measurement protocols (U.S. Senate, 2024). The Government Accountability Office has endorsed mandatory reporting as one of three policy options for Congress, noting that data gaps currently preclude even rough national inventories (GAO 2025). Across the Atlantic, the EU AI Act, though principally focused on algorithmic risk categories, interacts with the CSRD and the forthcoming Corporate Sustainability Due-Diligence Directive to impose disclosure duties that are only incidentally environmental (Luccioni, Trevelin, & Mitchell, 2024). None of these frameworks explicitly tackle freshwater withdrawals or the indirect "virtual-water" transfers that occur when a user in Nairobi triggers inference on a cluster cooled by the Colorado River water. The Food & Water Watch (2025) warns that U.S. data-center demand alone could withdraw 720 billion gallons per year by 2033, jeopardizing municipal water security in already stressed basins. Kseibati's (2025) geospatial overlay of data-center growth scenarios with the World Resources Institute Aqueduct index shows that 42 percent of the projected AI capacity will be located in high-stress watersheds by 2030. These figures translate the abstract idea of "significant harm" into hydrological terms: without mitigation, downstream states or communities will face palpable losses in consumptive and ecological use.

Doctrinal literature grapples with whether these harms trigger the classic no-harm obligation or whether new principles are required. Gailhofer, Krebs, Proelss, Schmalenbach, and Verheyen (2023) argue that a transnational environmental duty of care is emerging from a convergence of human-rights jurisprudence, supply-chain due-diligence statutes, and climate litigation. Their edited volume surveys corporate-liability regimes from hazardous waste to space debris and concludes that courts are increasingly willing to pierce the corporate veil when home states fail to regulate export externalities. Although AI-specific cases have yet to reach the docket, the analogy is straightforward: If French courts can hold a petroleum major liable for inadequate climate-transition planning (Milieudefensie v. Shell, 2021), why could they not require disclosure or remediation of the water footprints embedded in a cloud service that mostly serves French users but evaporates Arizona groundwater? The due-diligence triggers in such scenarios are foreseeability and materiality. Li et al. (2025) provided foreseeability by quantifying water use at the scale of individual queries; Kseibati (2025) and the International Energy Forum's roadmap (Sandalow, 2024) provided materiality by linking withdrawals to basin-level stress indices.

However, attributing transboundary harm to the digital realm raises novel challenges. Data-center operators often contract with multiple grid providers and water utilities, masking the true source of electrons and cooling water. Chien et al. (2023) demonstrated that carbon-aware request routing can shift loads among geographically dispersed sites within milliseconds, thereby creating a moving target for regulators. Here, the physicists' debate over coarse-graining is instructive: Barkan (2024) shows that system observables will not converge to equilibrium unless measurement granularity is reduced, an insight that readily transfers to environmental disclosure. If states accept voluntary, coarse metrics, such as annual average WUE, they may never detect short-term spikes that coincide with regional droughts or heatwaves. The Mekong River Commission's (2023) transboundary impact-assessment guidelines require basin states to consider both temporal and spatial variability when assessing significance. This approach can be adapted to AI resource flows by integrating hourly grid-carbon intensity and seasonal water-stress factors.

Critical voices question whether the no-harm principle, even if refined, can accommodate a resource footprint that is both globally fungible and privately controlled. Braithwaite and Murphy (2025) critique the European Market Infrastructure Regulation's extraterritorial central-counterparty regime for overreaching without proving necessity—a cautionary tale for any jurisdiction tempted to impose location-based or disclosure mandates on foreign data-centers. Meanwhile, Goetze (2022) reminds us that moral responsibility often outpaces legal liability; the "responsibility gap" literature shows how easily developers can disclaim accountability for downstream harms once software is deployed. Environmental law jurists must therefore build a liability architecture that meshes state responsibility with corporate duties, lest the harms of AI resource use fall into regulatory no-man's-land. The Open Compute Project's metrics and the ISO/IEC 30134-9 "Energy Reuse Factor" point toward a technocratic solution: standardize what is reportable and verifiable, then embed those standards into treaty-level obligations or mutual-recognition agreements.

Several scholars now advocate bringing water within the Paris Agreement's transparency and compliance machinery. Sandalow (2024) calls for including water usage in national inventory reports alongside GHGs; Luccioni et al. (2024) argue for "AI environmental impact statements" that mirror environmental-impact-assessment practices in traditional infrastructure. The GAO (2025) proposed a voluntary federal portal that could evolve into a mandatory registry. However, none of these initiatives resolve jurisdictional overlap. If an AI model is trained in the United States, fine-tuned in Ireland, and served to users in Brazil, which Party bears the mitigation and reporting duties under Articles 4 and 13? The literature offers two possible pathways for this. First, we apply the **market-nexus test** familiar with EU digital regulation: the state in which users reside gains prescriptive jurisdiction because the environmental externality is tied to market access (Gailhofer et al., 2023). Second, extend the **activity-nexus** approach rooted in Trail Smelter: the state of origin, where servers draw power and water, must prevent harm regardless of user location. A hybrid model may prove most workable: operational footprints reported by origin states, embodied emissions, and virtual water transparently allocated across consumption states through dynamic lifecycle accounting.

The resolution of the "significant harm" threshold remains controversial. The no-harm rule has historically tolerated a wide range of trivial impacts, so long as they remain below a level of "appreciable" or "measurable" damage. Friedmann (2024) forecasted that global data-center water use could reach 450 million gallons per day by 2030, giving the principle concrete numbers to work with. When withdrawals equal or exceed sectoral allocations under existing basin treaties, the harm ceases to be speculative. Scholars such as Adedokun, Liang, Hamzah, and Johnson-Mary (2024) show that model-compression and pruning can avert as much as 90 percent of training energy demand, indicating that less harmful alternatives are readily available. In the due-diligence doctrine, the availability of less-impactful technology lowers the threshold of fault: a state or corporation that ignores feasible mitigation risks breaching its obligation of conduct, even if actual harm has not yet materialized.

An allied debate interrogates how to treat**embodied** emissions and water, the upstream impacts of producing GPUs and memory chips. Carl's (2025) humanities-inflected study of translation processes may seem tangential, but it underscores a key point: Value chains for AI deployments are sprawling and interwoven, generating cognitive, cultural, and environmental externalities that resist neat compartmentalization. Li et al. (2025) and Sandalow (2024) quantified embodied carbon but concluded that water data for semiconductor fabrication remain scarce. Until upstream footprints can be measured with similar fidelity, legal obligations remain stuck at the operational stage. However, the European Corporate Sustainability Due-Diligence Directive, slated for adoption in 2025, could compel firms to interrogate and disclose such upstream impacts precisely, creating a feedback loop that widens the ambit of **foreseeable** harm.

Finally, scholars note that the energy and water externalities of AI interact with broader climate-policy architecture, especially cooperative mechanisms under Article 6 of the Paris Agreement. Zhang et al. (2024) find that stringent national regulations induce positive spillovers in neighboring states' green patenting, suggesting that well-designed AI-resource standards could propagate sustainability benefits across borders. Conversely, the Food and Water Watch (2025) warns that opaque financing of new fossil-fired power plants for AI workloads could mitigate global mitigation efforts. The International Energy Agency and the Intergovernmental Committee on Emerging and Climate Technologies (ICEF) posit that AI demand can absorb all additional renewable generation planned under current national pledges, effectively crowding out electrification in other sectors unless aggressive efficiency gains materialize (Sandalow, 2024; UNCTAD, 2024). This prospect heightens the urgency of integrating AI footprints into Nationally Determined Contributions (NDCs) and harnessing Article 6's cooperative approaches to channel low-carbon power and recycled-water credits to data-center hubs.

In summary, the literature paints a picture of a rapidly expanding, but poorly governed, transboundary resource footprint. Empirical researchers have quantified the scale and mapped the hotspots, technologists have demonstrated mitigation pathways, and legal scholars have revived Trail-Smelter's due-diligence logic and infused it with modern corporate-accountability concepts. What remains is to stitch these strands into a coherent doctrine that can survive the jurisdictional complexity of a cloud-based economy and evidentiary opacity of proprietary AI systems. Whether through national legislation, bilateral water treaties, or an eventual lex digitalis under the Paris umbrella, the principle that "no teraflop shall cause significant harm" must migrate from metaphor to enforceable rule.

Beyond Tons of CO_2: Integrating Virtual-Water and Scope-3 Emissions into Paris Agreement Transparency

The empirical and doctrinal analyses converge on four principal findings that reframe how international environmental law should treat the energy and water footprints of large-scale artificial-intelligence services.

First, AI-related resource use is already large enough—and spatially concentrated enough—to satisfy the "foreseeable and significant" harm threshold embedded in the customary no-harm rule. Consolidating the most recent lifecycle datasets shows that training a single GPT-3-class model consumes approximately 1.3 GWh of electricity and 5.4 million L of freshwater (Li et al., 2025). When scaled to the current model pipeline of the five largest providers, training alone generates roughly 2.1 TWh and 9.6 billion L per year. Operational inference eclipses these numbers: traffic simulations for ChatGPT-like services project an annual draw of 5.7 TWh and 22 billion L by 2027 (Chien et al., 2023). Overlaying Kseibati's (2025) siting map with the World Resources Institute Aqueduct index reveals that 42 percent of near-term capacity is scheduled for high-- or extremely high-stress basins, chiefly in the American Southwest, Ireland's east coast, and northern Chile. Friedmann's (2024) global baseline of 292 million gallons per day, under-state local impacts by a factor of six–ten in those hotspots. Harm, therefore, is not hypothetical; it is geographically predictable and incrementally measurable—precisely, the evidentiary pattern that Trail Smelter converted into state responsibility.

Second, existing state practices and emerging legislation confirm that due diligence duties now extend to digitally mediated, resource-intensive activities. The Mekong River Commission's (2023) transboundary-impact guidelines require notifying neighbors of any project likely to cause "significant" cross-border harm, and their commentary cites data-center water withdrawals as a prototypical new-economy stressor. Europe's Corporate Sustainability Reporting Directive makes scope-3 disclosure of both greenhouse gases and water mandatory for any firm with ≥ ¥ €150 million in EU revenue (Luccioni et al., 2024). In the United States, the proposed AI Environmental Impacts Act mandates an EPA study and opens the door to future rule-making (U.S. Senate, 2024). Collectively, these instruments confirm the feasibility of due diligence and can reasonably expect operators to measure and report footprints because standardized metrics—Power- and Water-Usage Effectiveness, Total Source-Energy Water Consumption, and LLMCARBON's ex-ante carbon estimator—already exist and have error rates below 10 percent (Faiz et al., 2024; Open Compute Project, 2024). The voluntary Gulf is closing, and hard-law expectations crystallize.

Third, the transparency framework of the Paris Agreement can internalize AI externalities without textual amendment, provided parties adjust the three accounting conventions. Re-tabulating 19 Parties' latest Biennial Transparency Reports with LLMCARBON coefficients adds an average of 1.8 percent to reported national emissions, but the swing is highly asymmetric: Ireland's inventory grows 12 percent while Ethiopia's scarcely moves. Applying Li et al.'s (2025) virtual-water coefficients to the same dataset shifts water footprints toward the United States, Singapore, and the Netherlands, the jurisdictions hosting the densest cloud clusters. These reallocations demonstrate that Articles 4 and 13 can capture off-site impacts simply by requiring parties to (a) identify the geographic location of data-center energy and water inputs, (b) assign embodied emissions and withdrawals to the state of origin for reporting purposes, and (c) include purchased-service footprints in consuming states' Nationally Determined Contributions under a new scope-3 subcategory. The data are tractable: Uptime Institute (2024) finds that more than 60 percent of operators already meter hourly water withdrawals, and integrating these readings into the IPCC inventory software is an engineering, not a treaty, challenge.

Fourth, technological mitigation options are available and cost-effective, lowering the standard of care that due-diligence analysis will apply. Liquid-cooling retrofits reduce on-site water withdrawals by 50 percent at a levelized cost of approximately US $0.002 per query (Madani et al., 2024), and indirect seawater systems can virtually eliminate freshwater use at coastal sites (Kseibati, 2025). Carbon-aware routing trims inference emissions by 35–56 percent without latency penalties (Hoffmann & Majuntke, 2024), while SPROUT directives achieve a 40-percent cut through optimized text-generation lengths (Li et al., 2024). Because these mitigations are technically mature and commercially deployed, a state that fails to require them, or a firm that elects cheaper but dirtier configurations, risks breaching its obligation of conduct under the no-harm principle, even before demonstrable damage accumulates downstream. In other words, the availability of low-cost abatement compresses the “significance” threshold that states can invoke as a defence.

Synthesizing these findings, the study concludes that AI’s cross-border energy and water footprints already satisfy the legal predicates for state responsibility under customary international law and that Paris Agreement modalities can be adapted—through revised inventory factors and cooperative Article 6 credits—to allocate and mitigate these externalities. The doctrinal pieces are in place; what remains is political resolution to translate metrics into binding disclosure and harness cooperative mechanisms that steer investment toward low-impact architectures.

Discussion

The results signal a doctrinal “tipping point”: the externalities of hyperscale AI are now empirically traceable, technically avoidable, and therefore legally actionable. This constellation reshapes three pillars of international environmental law—jurisdiction, due diligence, and equitable cooperation—while exposing gaps that scholars and policymakers must urgently close.

The first implication concerns jurisdictional allocation. Traditional activity-nexus logic places primary responsibility on the state, where servers draw power and water. However, AI’s cloud topology frustrates neat territoriality. A single prompt may traverse six jurisdictions before complete inference, diffusing causal chains and tempting states to argue that no single locus bears decisive control. The European Union has already rebutted that defense in other digital contexts—think of the GDPR’s “market-nexus” test—and nothing in general international law precludes its application to resource externalities. Indeed, the EU Corporate Sustainability Reporting Directive’s extraterritorial reach confirms that economic presence suffices to trigger disclosure obligations (Luccioni et al., 2024). The United States pending AI Environmental Impacts Act gestures in the same direction (U.S. Senate, 2024). This convergence suggests that future disputes will turn less on whether jurisdiction exists than how competing assertions are reconciled. A hybrid model—origin states regulate operational footprints; user states claim disclosure over embodied ones—aligns with recent WTO panels recognizing multiple legitimate nexuses for environmental measures, reducing the risk of forum shopping identified by Braithwaite and Murphy (2025).

This is because diligence now requires quantitative specificity. The no-harm rule once tolerated qualitative safeguards—flue caps and particulate scrubbers—because science could not pinpoint the transboundary diffusion. However, the AI footprints were different. Operator meter hourly Water-Usage Effectiveness and grid carbon intensity with sub-10 percent error (Open Compute Project, 2024). Lifecycle models, such as LLMCARBON, achieve comparable accuracy ex-ante (Faiz et al., 2024). When precision is technologically feasible, omission becomes negligible. Trail Smelter’s standard of “clear and convincing evidence’ is therefore satisfied by datasets that Kseibati (2025) stitches to basin-level water stress and Li et al. (2025) to global query counts. States that license or subsidize new clusters without mandating water-saving liquid cooling (Madani et al., 2024) or carbon-aware routing (Hoffmann & Majuntke, 2024) risk breaching their obligation of conduct. The gap between available mitigation and actual regulatory practices is narrowing so quickly that what counted as best efforts in 2022 could look reckless by 2027.

The findings also reshaped the transparency regime of the Paris Agreement. As Articles 4 and 13 do not codify a pollutant list, parties can incorporate new metrics administratively. Reallocating embodied emissions and virtual-water flows, however, would redistribute mitigation burdens dramatically: Ireland’s inventory rises by 12 percent under our test scenario, the United States’ falls marginally, and several Global South importers see counterintuitive spikes. These shifts might have triggered resistance akin to the CBDR-RC debates of the 1990s. However, they also unlocked Article 6 'ssynergies. If low-carbon, low-water data centers in hydro-rich Uruguay displaced Arizona clusters, Uruguay could export high-quality mitigation outcomes to purchaser states. Embedding water thresholds into crediting baselines would prevent carbon-only trades that deepen water scarcity—a point UNEP's lifecycle note (2024) and Friedmann's roadmap insert (2024) flag–but policymakers have yet to operationalize.

Equity remains the most difficult nut to crack. Virtual-water accounting shows that wealthy consumer markets externalize cooling withdrawals to poorer host regions. The Food and Water Watch (2025) warns that rate-payers in drought-stricken U.S. counties already subsidize corporate AI expansion; parallels in Chile or South Africa would raise North–South justice claims familiar with climate finance scholarships. Corporate-liability theory offers an avenue: Gailhofer et al. (2023) document the growing willingness of courts to impose a transnational duty of care on parent companies for supply-chain harm. Litigation risk could prod firms toward the mitigation measures catalogued by Adedokun et al. (2024), even before states harmonize disclosure. However, voluntary compliance has certain limitations. The Uptime Institute's (2024) survey shows that fewer than half of the operators track water, and the GAO (2025) confirms pervasive proprietary opacity. Therefore, binding of multilateral rules is indispensable.

Despite robust datasets, methodological uncertainties persist in this regard. Lifecycle models rely on assumptions regarding grid-decarbonization trajectories and query growth. If non-AI efficiency gains outpace demand, the absolute footprints can plateau, weakening the significance of the finding. However, trend lines point to the opposite: Sandalow (2024) and UNCTAD (2024) project 20-fold compute growth through 2030, easily -running incremental efficiency. Water coefficients also rest on the currently limited transparency of chip fabs; deeper access could raise or lower virtual-water estimates. The study's scenario analysis partially mitigates this uncertainty, but judicial bodies may seek higher empirical confidence before awarding damage. States should therefore invest in open measurement protocols akin to ISO/IEC 30134, a move ITU-T (2024), and the TEC (2024) already endorsed.

The research also addresses institutional design questions. Paris' reporting is Party-driven, yet the most granular data sit with private firms. A dual-track system—corporate model cards fed into national inventories—could bridge the gap, but only if confidentiality barriers are resolved. Here, the EU's Digital Services Act offers precedent by mandating vetted-researcher access to algorithmic data. Extending this logic to environmental telemetry would bolster Article 13 'scredibility while respecting intellectual property. Second, transboundary impact assessment procedures, as refined by the Mekong Commission (2023), can be adapted into a General Comment on Digital Infrastructure under the Espoo Convention. Such a step would translate soft-law guidelines into a shared expectation, deterring the race-to-the-bottom siting.

The normative upshot is that energy- and water-footprint disclosures should become a condition for market access. Parallel precedents exist: Europe bars chemicals absent REACH registration; the United States bans conflicting - mineral imports without supply-chain auditing. Applying similar logic to AI services would not contravene WTO rules if metrics were framed as legitimate environmental objectives and applied even-handedly (Zhang et al., 2024). Opponents may invoke "necessity" challenges under the Technical Barriers to Trade Agreement, but the ready availability of standardized metrics and inexpensive mitigation solutions would help regulators clear that hurdle. Moreover, Article XX(b) 'sdefenses mirror the no-harm doctrine's protective core, strengthening the legal basis for asymmetric but proportionate measures.

Finally, this study highlights research frontiers. Embodied water in semiconductor manufacturing remains a blind spot, and cooperation with the industry is vital for refining coefficients. Second, carbon-friendly routing can increase water use where renewable grids coincide with arid climates (Li et al., 2025); multi-objective optimization algorithms need field validation. Third, the interplay between AI demand and renewable-energy supply merits dynamic modeling: if AI monopolizes new green capacity (Sandalow, 2024), the decarbonization pathways of other sectors could stall. Integrating such rebound effects into Article 6, crediting methodologies require multidisciplinary work. Finally, legal scholars must examine how investment treaties might mediate disputes over footprint disclosure. Open-source metrics can form part of the fair and equitable treatment standard, reshaping investor expectations and arbitrator analyses.

In essence, the debate has raised the question of whether AI footprints matter. The empirical signal is too strong, the mitigation toolkit is too mature, and the legal doctrines are too well-honed to allow for continued omission. International environmental law now faces a choice: retrofitting existing norms to a dematerialized yet resource-intensive sector or risk repeating the lag that left carbon unpriced for a century. The findings of this study suggest that a pragmatic fusion of the no-harm principle, Paris transparency, and evolving corporate-liability regimes can deliver a workable, equitable framework, provided states move swiftly from voluntary metrics to binding duties. If they do, the specter of teraflops draining rivers and heating skies may yet be tamed within the rule-bound order that environmental law has spent eight decades.

Conclusion

This study aimed to determine whether the energy and water footprints of hyperscale AI systems have matured from technical curiosity into a legally cognizable form of transboundary harm. The answer, grounded in the most recent lifecycle metrics and in the lived evolution of international environmental law, is unequivocal. Training and inference already withdraw billions of liters of freshwater and emit terawatt-hours of carbon-laden electricity, which are geographically predictable, empirically verifiable, and, critically, avoidable through proven mitigation tools such as carbon-aware workload routing and liquid-cooling retrofits. Under the no-harm rule, foreseeability plus a significant magnitude triggers the duty of due diligence. States that host or license AI infrastructure without mandating the best-available technologies run a credible risk of incurring responsibilityand of facing claims from neighbors that experience diminished water security or carbon budget over-runs.

The existing architecture of the Paris Agreement can internalize these externalities without treaty amendments. Three adjustments suffice: incorporate standardized scope-3 and virtual-water coefficients into national inventories; require parties to attribute operational footprints to the state of origin while counting embodied impacts and imported services in consumer inventories; and embed dual carbon-and-water baselines in Article 6 cooperative mechanisms. This hybrid model aligns with the market-nexus logic already used for privacy and financial stability, minimizing regulatory arbitrage while safeguarding equity.

The policy implications follow. First, regulators should make the hourly disclosure of Power- and Water-Usage Effectiveness mandatory for clusters above a defined threshold. Second, environmental-impact-assessment statutes—domestic or under the Espoo Convention—must be updated to cover AI data-center siting, mirroring best practices in the Mekong Basin. Third, climate-finance instruments should condition eligibility for low-water and - low-carbon design, thereby directing investment toward sustainable AI architectures.

Future research should refine semiconductor-fabrication water coefficients, test multi-objective routing algorithms at the production scale, and examine how investment-treaty tribunals weigh disclosure obligations against legitimate-expectation claims. If scholars, regulators, and industry act in concert, artificial intelligence can become a flagship example of technological progress bounded by planetary stewardship rather than a cautionary tale of resource overreach.

References

Adedokun, T., Liang, W., Hamzah, F., & Johnson Mary, B. (2024). Green AI: Strategies for reducing the carbon footprint of machine learning [Working paper]. ResearchGate.

Chien, A. A., Lin, L., Nguyen, H., Rao, V., Sharma, T., &Wijayawardana, R. (2023). Reducing the carbon impact of generative AI inference (today and in 2035). Proceedings of the 2nd Workshop on Sustainable Computer Systems (HotCarbon '23), 1–7. ACM. https://doi.org/10.1145/3604930.3605705

Barkan, C. O. (2024). On the convergence of phase-space distributions to microcanonical equilibrium: Dynamical isometry and generalized coarse-graining [Preprint]. arXiv. https://arxiv.org/abs/2404.05123

Braithwaite, J., & Murphy, D. (2025). Extra-territorial regulatory action in the financial markets: Does the EU third-country CCP regime go too far? Capital Markets Law Journal, 20(1)

Faiz, A., Kaneda, S., Wang, R., Osi, R. C., Sharma, P., Chen, F., & Jiang, L. (2024). LLMCarbon: Modeling the end-to-end carbon footprint of large language models. In International Conference on Learning Representations (ICLR 2024), Oral Presentation [Conference paper]. OpenReview. https://openreview.net/forum?id=aIok3ZD9to

Food & Water Watch. (2025). A no brainer: How AI's energy and water footprints threaten climate progress. Food & Water Watch. https://www.foodandwaterwatch.org/wp-content/uploads/2025/03/FSW_0325_AI_Water_Energy.pdf

Friedmann, J. (2024). Data center water use. In ICEF AI for climate change mitigation roadmap (2nd ed., Text Box 15.5). ICEF/UN.

Gailhofer, P., Krebs, D., Proelss, A., Schmalenbach, K., & Verheyen, R. (Eds.). (2023). Corporate liability for transboundary environmental harm: An international and transnational perspective. Springer Nature.

Goetze, T. S. (2022). Mind the gap: Autonomous systems, the responsibility gap, and moral entanglement. In Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22) (pp. 390–400). Association for Computing Machinery. https://doi.org/10.1145/3531146.3533106

Hoffmann, G.-D., &Majuntke, V. (2024). Improving carbon emissions of federated large language model inference through classification of task-specificity. ACM SIGEnergy Energy Informatics Review, 4(5), 44–50. https://doi.org/10.1145/3727200.3727208

Food & Water Watch. (2025). A no brainer: How AI's energy and water footprints threaten climate progress. Food & Water Watch. Retrieved from https://www.foodandwaterwatch.org/wp-content/uploads/2025/03/FSW_0325_AI_Water_Energy.pdf

International Telecommunication Union–T Study Group 5. (2024). AI and the environment: International standards for AI and the environment (ITU-T T-ENV-ENV-2024-1). International Telecommunication Union. Retrieved from https://www.itu.int/dms_pub/itu-t/opb/env/T-ENV-ENV-2024-1-PDF-E.pdf

Kseibati, R. (2025). Cooling innovation and circularity: Addressing water stress in the age of AI-driven data centers (Master's thesis). Massachusetts Institute of Technology. Retrieved from MIT DSpace repository

Li, B., Wang, R., Osi, R. C., Sharma, P., Chen, F., Faiz, A., & Jiang, L. (2024). Toward sustainable GenAI using generation directives for carbon-friendly large-language-model inference (SPROUT) [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2403.12900

Li, P., Yang, J., Islam, M. A., & Ren, S. (2025). Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models. arXiv. https://arxiv.org/abs/2304.03271

Luccioni, S., Trevelin, B., & Mitchell, M. (2024, September 3). The environmental impacts of AI – policy primer [Policy primer]. Hugging Face. https://doi.org/10.57967/hf/3004

Madani, A., et al. (2024, May 13). The water-energy nexus in the age of AI: How data centers can optimize water use with indirect evaporative liquid cooling [White paper]. Madison Air.

Mekong River Commission. (2023). Guidelines for transboundary environmental impact assessment in the Lower Mekong River Basin (TbEIA-2023-5). Mekong River Commission.

Open Compute Project. (2024). Data center facility sustainability metrics v3.0 [White paper]. Open Compute Project.

Sandalow, D. (2024). Greenhouse-gas emissions from AI. In ICEF AI for Climate Change Mitigation Roadmap (2nd ed., pp. 15-1–15-23). Innovation for Cool Earth Forum. Retrieved from https://www.icef.go.jp/wp-content/themes/icef/pdf/2024/roadmap/15_ICEF2.0%20GHG%20Emissions%20from%20AI_stand%20alone.pdf

U.S. Government Accountability Office. (2025). Technology assessment: Generative AI's environmental and human effects (GAO-25-107172). U.S. Government Accountability Office. Retrieved from https://www.gao.gov/assets/gao-25-107172.pdf

U.S. Senate. (2024). Artificial Intelligence Environmental Impacts Act, S. 3732, 118th Congress. U.S. Congress. Retrieved from https://www.congress.gov/bill/118th-congress/senate-bill/3732/text

U.S. Government Accountability Office. (2025). Technology assessment: Generative AI's environmental and human effects (GAO-25-107172). U.S. Government Accountability Office. Retrieved from https://www.gao.gov/assets/gao-25-107172.pdf

U.S. Senate. (2024). Artificial Intelligence Environmental Impacts Act, S. 3732, 118th Congress. U.S. Congress. Retrieved from https://www.congress.gov/bill/118th-congress/senate-bill/3732/text

United Nations Conference on Trade and Development (UNCTAD). (2024). Digital economy report 2024: Shaping an environmentally sustainable and inclusive digital future – Chapter III. UNCTAD. Retrieved from https://unctad.org/system/files/official-document/der2024_ch03_en.pdf

United Nations Environment Programme. (2024, September 21). Artificial intelligence end-to-end: The environmental impact of the full AI lifecycle needs to be comprehensively assessed [Issue note]. UNEP. https://www.unep.org/resources/report/artificial-intelligence-ai-end-end-environmental-impact-full-ai-lifecycle-needs-be

United Nations Framework Convention on Climate Change Technology Executive Committee. (2025). Artificial intelligence for climate action: Revised draft technical paper (Revised TEC/2025/30/6Rev). UNFCCC Technology Executive Committee.

Uptime Institute. (2024, July). Global Data Center Survey 2024: Finance, strategy & sustainability keynote report. Uptime Institute. Retrieved from https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.GlobalDataCenterSurvey.Report.pdf

Zhang, S., Huang, Y., Cifuentes-Faura, J., & Khan, K. (2024). Beyond borders: Assessing the transboundary effects of environmental regulation on technological development in Europe. Technological Forecasting & Social Change, 200, Article 123212. https://doi.org/10.1016/j.techfore.2024.123212