

White Paper on:

AI and Environment



White Paper on:

AI and Environment

Nature and People Foundation

China Institute, Fudan University

Pahle India Foundation

in partnership with

Dynamic Coalition on Environment of the Internet Governance Forum

July 2026

List of experts (alphabetic order):

- **Polina Burkova** – Moscow State Institute of International Relations;
- **Sofia Denisova** – AI&ESG Expert;
- **Liu Dian** – Associate Research Fellow, China Institute, Fudan University;
- **Nitish Dogra** – MD, MPH, Fulbright Fellow, Scientific & Technology Lead, I-LEAP (India-Lead Elimination Action Partnership), Pahle India Foundation;
- **Victoria Elias** – Director of Conservation, Nature and People Foundation;
- **Alexey Kokorin** – Climate Expert, Nature and People Foundation;
- **Sergey Rybakov** – General Director, Nature and People Foundation, and Co-Chair, Dynamic Coalition on Environment, Internet Governance Forum;
- **Jeffrey D. Sachs** – University Professor and Director, Center for Sustainable Development, Columbia University, and President, UN Sustainable Development Solutions Network;
- **Zhang Weiwei** – Director of the China Institute, Fudan University, and Board Member, the China National High-end Think Tank Council.

Table of Contents

Foreword (Jeffrey D. Sachs)	7
Foreword	9
Executive Summary	11
Section 1. Aim and Objectives	14
The Dual Challenge	15
Why This Paper — and What Makes It Different	16
Scope, Objectives, and Structure	19
Timeliness	19
A Note on Approach	20
Section 2. AI Models	21
2.1 Life Cycle Assessment Overview for AI Models	22
Why Life Cycle Assessment Matters for AI and Who Should Conduct It	22
Three Phases of AI’s Environmental Lifecycle	23
The State of LCA Practice for AI	24
Methodological Challenges, Disagreements, and Data Deficits	25
Towards Standardization	26
2.2 Carbon Footprint of AI Models and Data Centers	27
The Growth Trajectory	27
Carbon Footprint Is a Design Variable	27
Inference: The Larger and Overlooked Share	28
The Transparency Deficit	29
2.3 Water Footprint of AI Models and Data Centers	30
The Scale of the Problem	30
Two Types of Water Consumption	30
Geographic Concentration and Water Stress	30
Innovation and Alternatives	31
Measurement Gaps	31
2.4 Metrics and Tools for Sustainable Assessment	32
The Conceptual Foundation: Green AI	32
The Measurement Tool Ecosystem	32
Standardized Metrics and Frameworks	33
The Adoption Gap	33
Institutional Momentum	34
Section 3. Climate Change	34
3.1 Why Do We Care About Climate Change	36
The International Framework	36
The AI–Climate Nexus	37

3.2 How AI Is Contributing and Could Contribute to Address Climate Change	38
Climate Modelling and Weather Prediction	38
Emissions Monitoring and Tracking	39
Energy System Optimisation	39
Wildfire Prediction and Disaster Response	40
Carbon Capture and Materials Discovery	40
3.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments	41
Corporate Climate Targets in the AI Sector	41
Frameworks and Accountability Mechanisms	41
Instruments from BRICS and the Global South	42
Section 4. Biological Diversity	43
4.1 Why Do We Care About Biological Diversity	44
The International Framework	44
The Scale of the Crisis	45
4.2 How AI Is Contributing and Could Contribute to Address Conservation and Sustainable Use of Biological Diversity	46
Monitoring and Species Identification	46
The Biodiversity Policy Analyzer	47
Scientific Assessment and Attribution	47
AI for Biodiversity Finance	47
GPAI and UNDP Frameworks	47
4.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments	48
Ethical Principles for AI in Biodiversity	48
Data Sovereignty and Indigenous Rights	48
Capacity Building and the 2030 Timeline	49
Section 5. Desertification and Land Degradation: Challenges and Mitigation	50
5.1 Why Do We Care About Desertification and Land Degradation	51
The International Framework	51
The Scale of the Challenge	51
5.2 How AI Is Contributing and Could Contribute to Address Desertification and Land Degradation: Challenges and Mitigation Issues	52
The International Drought Resilience Observatory (IDRO)	52
Land Degradation Monitoring	52
Sand and Dust Storm Monitoring	53
5.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments — Challenges and Mitigation Issues	54
Persistent Challenges	54
Pathways Forward	54
Section 6. Availability and Sustainable Management of Water Resources	55
6.1 Why Do We Care About Water Resources	56
The International Framework	56
The Water–AI Nexus: A Dual Relationship	56

6.2 How AI Is Contributing and Could Contribute to Address Water Resources Issues	58
Flood Prediction and Early Warning	58
Urban Water Infrastructure and Leak Detection	59
Precision Agriculture and Irrigation Optimisation	59
Water Quality Monitoring and Prediction	59
Groundwater and Watershed Management	60
The IDRO Connection	60
6.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments	61
Corporate Water Stewardship in the AI Sector	61
International Accountability Frameworks	62
The Regulatory Landscape	62
Section 7. Policy and Regulatory Considerations	63
7.1 Comparative Overview	64
Cross-Cutting Comparison	64
7.2 Jurisdiction Profiles	67
BRICS	67
China	67
European Union	69
India	69
Russia	70
United States	71
7.3 Key Regulatory Gaps	73
Disclosure Gap	73
Inference Gap	73
Water Gap	73
Geographic Equity Gap	73
Application Accounting Gap	74
Public Accountability Gap	74
Section 8. Role of Public Sector, Private Sector, Civil Society, and International Organizations	75
8.1 Role of Public Sector	76
Regulatory Frameworks	76
Public Procurement	76
Public Research Investment	77
Direct Environmental Constraints	77
8.2 Role of International Organizations	78
ITU (International Telecommunication Union)	78
UNEP (United Nations Environment Programme)	78
UNFCCC (United Nations Framework Convention on Climate Change)	78
CBD (Convention on Biological Diversity)	78
UNCCD (United Nations Convention to Combat Desertification)	79
UNESCO (United Nations Educational, Scientific and Cultural Organization)	79
OECD (Organisation for Economic Co-operation and Development) and GPAI (Global Partnership on Artificial Intelligence)	79
IGF Dynamic Coalition on Environment	79

8.3 Role of Private Sector	80
AI Developers	80
Hardware Manufacturers	80
Cloud Infrastructure Providers	81
Sustainability Innovators	81
8.4 Role of Civil Society	82
Academic and Research Institutions	82
Indigenous Peoples and Local Communities	82
Communities Affected by AI Infrastructure	83
Section 9. Next Steps	84
9.1 Challenges	85
Challenge 1. Data on AI's energy, carbon and water footprint remain incomplete and often non-comparable	85
Challenge 2. Environmental impacts are geographically uneven and can create equity problems	85
Challenge 3. AI infrastructure is expanding faster than regulatory and measurement systems	85
Challenge 4. AI-for-environment applications may generate rebound effects if efficiency gains lower costs and increase consumption	86
Challenge 5. Environmental AI systems require high-quality data and long-term institutional maintenance	86
Challenge 6. International cooperation is complicated by digital sovereignty, export controls and uneven access to computing power	86
9.2 Emerging Trends and Future Directions	87
Standardized AI Environmental Reporting and Accountability	87
Life-Cycle Assessment for Foundation Models	87
Carbon- and Water-Aware Workload Scheduling	87
Model Compression and Task-Specific Small Models	87
Clean-Power and Computing-Power Coordination	88
AI and Digital Twins for Environmental Monitoring and Governance	88
Open Environmental Data and Capacity Building for the Global South	88
9.3 Recommendations	89
For Governments	89
For AI Developers and Technology Companies	89
For Hardware Manufacturers	90
For International Organizations	90
For Civil Society and Research Institutions	90

Foreword

Jeffrey D. Sachs

Every so often a general purpose technology (GPT) arrives that not merely adds to the economy but fundamentally reorganizes it. The steam engine, electricity, and the computer each did so, and artificial intelligence now belongs unmistakably in their company. What sets AI apart is the sheer pace of its arrival and impacts. As a GPT, it is setting records on every axis at once: its rate of improvement, the speed of its adoption across the world, the breadth of its applications, and the scale of its likely impacts on economies, societies, and the natural environment. A technology moving this rapidly leaves our institutions of understanding and governance struggling to keep pace. This White Paper on AI and the Environment is a landmark effort to help close that gap, and it succeeds admirably.

Its first great virtue is timeliness. The infrastructure decisions that will determine AI's environmental footprint for decades — where data centers are built, how they are powered, how their thirst for water and materials is managed — are being made right now. The paper's key message is that AI's environmental impact is not a fixed cost but a governance choice. Whether AI ends up, on balance, a friend or a foe of the environment will be decided not mainly by the technology itself but by the design decisions, disclosure rules, and policy frameworks that the world adopts in the next few years. To have a state-of-the-art reference on these questions at this moment is a service of enormous practical value.

Its second virtue is comprehensiveness. The paper holds together, in a single frame, several topics that are mostly discussed separately. AI is both a growing source of environmental pressures and simultaneously among the most powerful tools we possess for providing sustainable solutions to environmental harms. The white paper follows AI across its full life cycle — energy, carbon, water, critical materials, and electronic waste — rather than focusing on carbon alone. It systematically examines, domain by domain, on AI's double-edged role in climate change, biological diversity, land degradation, and water resources, both as a potential threat and as a potential source of new solutions. In each domain, the paper links to the relevant

multilateral frameworks. And it carries the analysis all the way through to recommended solutions, the governance of those potential solutions, and the policy implications at every level of government from local to global, including the regulation of individual firms to national strategies to global coordination under UN treaties and bodies. It accomplishes all of this without sacrificing rigor or comprehensiveness.

The White Paper's third great virtue is the distinctive partnership of authorship. The White Paper is produced jointly by the Dynamic Coalition on Environment of the Internet Governance Forum, the Pahlé India Foundation, Fudan University, and the Nature and People Foundation, thereby integrating the perspectives and policy experience of China, India, and Russia, alongside the better-known cases of the United States and the European Union. This matters greatly. AI is governed very differently by its major developers and host nations, and it has become a central issue in geopolitics — partly because of its monumental economic weight, and partly because of its rapidly expanding and potentially very dangerous role in security and military affairs. An account of AI and the environment written from a single vantage point would be not only incomplete but misleading. By treating AI governance as the genuinely multipolar endeavor it is, this paper offers analysis that simply does not exist elsewhere.

The global task ahead in AI governance is formidable. We need serious scenario-building and long-range planning; we need regional integration of infrastructure and standards; and we need global governance that reaches across hardware, software, and applications alike. This will be difficult, not only because of the pace and scope of change, but also because the stakes — economic, political, and military — are so high, and because the incentives to compete threaten to overwhelm the urgent needs to cooperate. Yet such difficulties are not arguments for delay. The alternative to shared governance is an AI arms race whose environmental and geopolitical costs we cannot afford and should not risk. We must find ways to steer this extraordinary technology toward the common good, and toward the Sustainable Development Goals and Multilateral Environmental

Agreements that express our shared aspirations for a prosperous, peaceful, and sustainable world.

The White Paper is a hugely valuable reference work at this critical juncture. It deserves the close attention by many communities of practice, including the AI developers and hyper-scalers whose engineering choices carry such weight; national policy authorities and regulators; scholars, researchers, and think tanks; civil society organizations; and the international and United Nations bodies charged with forging common rules. To all of them the White Paper offers a powerful combination of technical depth, breadth of coverage, and international perspective. I congratulate the authors and partner institutions on this important achievement, and I heartily commend

this work to readers everywhere who share the conviction that the future of AI, and the future of our planet, can and must be built together for the common good.

Jeffrey D. Sachs

University Professor and Director, Center for Sustainable Development, Columbia University

President, UN Sustainable Development Solutions Network

Foreword

Artificial intelligence is reshaping the world, and the effort to govern it is reshaping the global landscape in turn. In just two or three years, a technical issue once dominated by a handful of countries has become a genuinely multipolar endeavor — one that bears on the placement of energy and computing power, on the making of standards and rules, and on how countries at different stages of development decide, at one and the same table, the relationship between this technology and humanity's future. It is at this moment of profound reorganization that research institutions from China and Russia have come together under the framework of the Dynamic Coalition on Environment of the United Nations Internet Governance Forum, and, together with partners from India, have completed this white paper.

This reorganization is seen with particular clarity in China. As an important pole of global AI development, China has risen rapidly not only in technology and industry but also in advancing ideas of governance. The World Artificial Intelligence Conference (WAIC) in Shanghai is a microcosm of this: from an industry gathering it has grown into a global platform convening governments, international organizations, industry, and academia, its influence expanding year on year. Its 2024 Shanghai Declaration on Global AI Governance called for AI to be applied to environmental protection, resource utilization, energy management, and biodiversity; its 2025 Global AI Governance Action Plan advocated joint AI energy- and water-efficiency standards and green computing technologies. Nor does China's exploration stop at advocacy: it has embedded environmental requirements into the very expansion of computing power — setting green targets for data centers, guiding eastern computing demand toward the renewable-rich west through the «Eastern Data, Western Computing» program, integrating AI with the energy system, and deploying AI for biodiversity monitoring, forest patrol, and environmental enforcement. Sustainable AI, this experience shows, cannot rest on model efficiency alone; it requires spatial planning, power-system coordination, infrastructure standards, and institutional capacity.

Other emerging forces are entering the stage rapidly. The India AI Impact Summit in New Delhi in February 2026 signaled a marked rise in the Global South's voice in AI governance, with investment pledges approaching 250 billion US dollars; the BRICS+ AI Alliance, launched in Moscow in December 2024, and the AI cooperation framework adopted at the Kazan BRICS Summit represent a new axis of cooperation. From Shanghai to New Delhi to Moscow, a line of cooperation unlike those of the past is taking shape — one that demands analysis capable of genuinely reflecting the priorities and realities of China, India, Russia, and the broader BRICS+ grouping. At the institutional level, too, the architecture of governance is forming quickly, from a now-permanent Internet Governance Forum to the inaugural Global Dialogue on AI Governance convening in Geneva in July 2026. An international order for governing AI's environmental impact is being built from the ground up — yet its environmental dimension remains underdefined.

This is the gap this white paper seeks to address. To the best of our knowledge, it is the first report to approach the relationship between AI and international sustainable development from the perspective of the BRICS countries and the Global South, bringing the environmental footprint of AI and its environmental benefits within a single analytical framework. Earlier discussions tended to examine only the cost of AI, or only its promise as an environmental solution, seldom weighing the two together; earlier analyses were shaped mainly by US, European, and to some extent Chinese perspectives, with few voices genuinely arising from a multipolar world. It is these gaps the paper hopes to fill — a task undertaken jointly because each partner brings an irreplaceable perspective: the China Institute of Fudan University, that of one of the largest contributors to global economic growth and a leading AI power, with China's experience of embedding environmental requirements into computing expansion; the Nature and People Foundation, that of a country with one of the world's largest digital ecosystems and long engagement across the three Rio conventions, contributing tools and cold-region monitoring applications underrepresented in global

datasets; and the partners from India, that of the world's most populous country and one of its fastest-growing digital economies. The shared conviction is simple: in a multipolar world, effective AI governance cannot be designed by one group of countries and exported to the rest.

On this basis the paper develops its analysis. It insists on examining the environmental costs and benefits of AI within a single framework, assessing impact across the full lifecycle — from critical-mineral extraction and chip manufacture, through training, inference, cooling, and water use, to hardware retirement and electronic waste. Its judgment is that AI's footprint is determined not by model size but by choices within human control — model architecture, hardware efficiency, grid carbon intensity, and data center siting — and that the burden is shifting from one-time training toward inference performed billions of times a day. It extends the discussion from climate to biological diversity, land degradation and desertification, and water resources, each aligned with the relevant multilateral framework. In every domain AI proves double-edged: it strengthens climate modeling, emissions monitoring, biodiversity observation, drought and flood early warning, and basin-scale water management as no other technology can, even as the data centers behind these applications draw electricity and water from the very systems under stress. The recurring lesson is that AI's environmental benefits are real but conditional — credible only when it genuinely reduces environmental pressure rather than shifting cost elsewhere, and only when the compute it requires is weighed against the savings it delivers.

On policy, the paper offers a systematic comparison across six jurisdictions — the BRICS grouping, China, the European Union, India, Russia, and the United States — and finds three contrasting governance models emerging: an infrastructure-planning-led path, a law-centered path, and a path led by voluntary corporate commitments. Each addresses part of the challenge; none integrates the two domains fully. Comparing them is more instructive than ranking them. It also names the gaps common to every jurisdiction: no AI-specific environmental disclosure; AI and environmental governance still siloed; the growing impact of inference unmeasured; water governance lagging energy governance; supply-chain impacts opaque; and the uneven geographic distribution of costs and benefits scarcely addressed. From these it distills six cross-cutting challenges — incomplete and non-

comparable data, geographically uneven impacts and equity, infrastructure outpacing governance, rebound effects, the need for high-quality data and sustained institutions, and cooperation under geopolitical strain — that together define the agenda for action.

We offer this white paper to two occasions in 2026: the inaugural Global Dialogue on AI Governance and the AI for Good Global Summit in Geneva, and the World Artificial Intelligence Conference in Shanghai. We chose this moment because 2026 is a pivotal year: the infrastructure decisions whose consequences will last for decades are being made now, while the frameworks capable of guiding them are not yet in place. The window is real, and it will close. We offer no single right answer and advocate no one policy position; we have gathered the best available evidence, marked where consensus and uncertainty lie, and set out options adaptable to different contexts. The value of this paper lies not in how many conclusions it reaches, but in bringing, for the first time, judgments drawn from different stages of development, energy endowments, and governance traditions to one and the same table. How the relationship between humanity's most powerful technology and its most essential foundation unfolds will be decided not by technology alone, but by the choices we make together in the years ahead. We offer this shared document as a contribution to that process, and invite our readers to take part in a critical and constructive spirit.

Signed jointly,

Zhang Weiwei

Director, China Institute, Fudan University, and Board Member, the China National High-end Think Tank Council

Sergey Rybakov

Director General, Nature and People Foundation, and Co-Chair, Dynamic Coalition on Environment, Internet Governance Forum

On behalf of the Dynamic Coalition on Environment of the Internet Governance Forum and all contributing experts

July 2026

Executive Summary

Artificial intelligence is now widely recognized as both a contributor to environmental harm and a powerful tool for addressing it. What has been missing is an analysis that considers these two dimensions together, extends the discussion beyond climate to the full range of environmental domains, and reflects the priorities of the countries where AI infrastructure is expanding most rapidly. This white paper, authored by the Dynamic Coalition on Environment of the Internet Governance Forum in partnership with the Pahlé India Foundation, Fudan University, and the Nature and People Foundation, sets out to fill that gap.

Its central argument is that the environmental impact of AI is **not a fixed cost but a governance choice**. Whether AI's contributions to solving environmental problems ultimately outweigh its contributions to creating them will be determined not by technology alone, but by the design decisions, disclosure rules, and policy frameworks adopted in the next few years — while infrastructure decisions with decades-long consequences are still being made.

At the same time, improvements in AI efficiency do not automatically translate into lower environmental impacts. As AI becomes cheaper and more accessible, growing demand for computation may offset efficiency gains — a phenomenon commonly referred to as the rebound effect (or Jevons paradox).

The dual nexus: footprint and contribution

The paper's analytical foundation is an integrated assessment of AI's environmental life cycle, examined across energy, carbon, water, critical materials, and electronic waste rather than carbon alone. Its key conclusion is that AI's footprint is not determined by model size but by choices within human control — model architecture and optimisation, hardware efficiency, the carbon intensity of the electricity grid, and the geographic siting of data centers — each capable of changing environmental impact by large factors. The same analysis shows that the environmental burden is shifting from model training towards inference as AI is deployed at scale, and

that water and embodied-material impacts remain far less visible than energy. This establishes the basis for treating measurement, disclosure, and design as the primary levers of environmental governance.

The four thematic domains examined in the paper — climate change, biological diversity, land degradation and desertification, and water resources, each aligned with the relevant multilateral framework — confirm that AI is genuinely double-edged in every one of them. AI is strengthening climate and weather modelling, independent emissions monitoring, biodiversity observation, drought and flood early-warning, and basin-scale water management in ways no other technology matches; at the same time, the data centers that make these applications possible draw electricity and water from the very systems under stress, and concentrate that demand in regions already facing scarcity. The recurring lesson across all four domains is that AI's environmental benefits are real but conditional: they are credible only when the additional computing demand required to produce them is measured against the savings they deliver, and only when applications that genuinely reduce environmental pressure are distinguished from those that merely shift costs elsewhere.

A multipolar perspective

This multipolar perspective also reflects the paper's broader objective of comparing governance approaches rather than ranking them according to a single model. The distinctive contribution of this paper lies in bringing together evidence and policy experience from India, China, and Russia alongside the more familiar United States and European cases, treating AI governance as a genuinely multipolar endeavor. This produces analysis that does not exist elsewhere in the literature: how a rapidly developing economy such as India confronts the tension between digital-infrastructure growth, water stress, and an environmental-clearance system never designed for cloud-based AI; how China has begun attaching environmental requirements directly to the expansion of computing power through infrastructure planning, ecological-monitoring deployment, and explicit links between AI policy and energy-system

transformation; and how Russia frames low-carbon energy as the enabling foundation for AI compute while contributing tools and cold-region environmental applications underrepresented in global datasets. These nationally grounded cases, contributed directly by the partner institutions, are the empirical heart of the paper.

A comparative reading of AI-environment governance

The paper's comparative analysis across six jurisdictions — the BRICS grouping, China, the European Union, India, Russia, and the United States — yields one of its most useful findings: three contrasting governance models are emerging, each addressing part of the challenge and none integrating the two domains fully. China's approach is **infrastructure-planning-led**, embedding environmental conditions into the approval and siting of computing capacity. The European Union's is **law-centered**, distributing obligations across AI, energy, and sustainability-reporting legislation. The United States' is **corporate-led**, relying on the voluntary commitments of a small number of dominant firms. Comparing these models is more instructive than ranking them: an effective international framework will likely need elements of all three.

Against this backdrop, the paper identifies the regulatory gaps common to every jurisdiction examined. No country currently requires AI-specific environmental disclosure; AI governance and environmental governance remain institutionally siloed; the rapidly growing impact of inference is largely unmeasured and unregulated; water governance lags far behind energy governance; supply-chain and embodied impacts remain opaque; and the uneven geographic distribution of AI's costs and benefits is rarely addressed by any framework. These gaps, not any single statistic, define the agenda for action.

Six cross-cutting challenges

Synthesizing across all domains, the paper identifies six challenges that any credible response must confront:

1. **Incomplete and non-comparable data.** The environmental costs of AI are becoming measurable, but its claimed benefits remain difficult to verify — there is still no recognized framework for validating AI-enabled environmental outcomes against the additional computing demand required to produce them.
2. **Geographically uneven impacts and equity.** The burdens of mineral extraction, water-intensive computing, and electronic waste fall disproportionately on regions and communities that share least in AI's benefits — including through the underrepresentation of non-English and indigenous knowledge.
3. **Infrastructure outpacing governance.** AI infrastructure is being built faster than the measurement and regulatory systems needed to manage it, while responsibility remains fragmented across institutions — narrowing the window for proactive governance.
4. **Rebound effects.** Efficiency gains that lower the cost of AI can increase total consumption, partially or wholly offsetting environmental benefits, so governance must address absolute resource use, not only relative efficiency.
5. **Data quality and institutional maintenance.** Environmental AI depends on high-quality data and sustained institutional capacity that are weakest precisely where environmental risks are greatest; it must be treated as continuous public infrastructure, not a one-time deployment.
6. **Cooperation under geopolitical strain.** Export controls, digital-sovereignty concerns, and uneven access to computing power complicate cooperation exactly when shared environmental challenges most require it — risking a new digital-environmental divide.

From analysis to action – initiatives for Geneva and WAIC

The paper translates this analysis into recommendations addressed to each category of actor — governments, AI developers, hardware manufacturers, international organizations, and civil society — built around mandatory and AI-specific disclosure, the integration of environmental criteria into AI governance, and equitable capacity-building for the Global South. Building on these, the Coalition proposes seven concrete initiatives for international follow-up in the discussions at the Global Dialogue on AI Governance (Geneva) and the World Artificial Intelligence Conference (Shanghai):

1. Develop a **common reporting template** for AI-related data-center energy, carbon, and water indicators, including PUE, WUE, CUE, renewable-energy share, and server utilisation.

2. Create a **comparative policy map** covering China, India, Russia, the EU, and the United States, with attention to how AI regulation interacts with climate and environmental policy.
3. Launch a **BRICS+ and Global South case repository** on AI for climate, biodiversity, land, and water governance.
4. Encourage **dual disclosure** — public agencies and companies reporting both the environmental footprint of AI systems and the environmental benefits claimed by AI applications.
5. Promote **pilot projects on carbon- and water-aware AI workload scheduling** across regions with different grid mixes and water-stress levels.
6. Use **international forums** to test the measurement framework and expand cooperation on implementation cases.
7. **Develop a common glossary** of AI-environment terminology and reporting metrics (including PUE, WUE, CUE, LCA, DSI, NBSAP and related concepts) to improve international comparability).

Section 1.

Aim and Objectives



The Dual Challenge

Artificial intelligence is transforming the global economy at a pace that few technologies have matched. From drug discovery and climate modelling to precision agriculture and disaster response, AI systems are demonstrating an extraordinary capacity to address problems that have long defied conventional approaches. Moreover, AI holds genuine promise as a tool for helping humanity overcome the obstacles that stand in the way of a sustainable model of existence — accelerating the clean-energy transition, monitoring fragile ecosystems, optimizing the use of scarce natural resources, and strengthening resilience to environmental shocks. Yet there is a paradox at the heart of this promise: the very infrastructure that powers these capabilities — vast networks of data centers, accelerated computing hardware, and energy-intensive training pipelines — is itself becoming a significant source of environmental pressure. A technology capable of helping to safeguard environmental stability is, at the same time, placing new strain upon it.

The most studied problem stemming from massive AI growth is its greenhouse gas emissions and impact on climate change. The statistic most commonly used to illustrate this connection is the growth in global electricity consumption. According to the International Energy Agency's Energy and AI report (April 2025), global data center electricity consumption reached approximately 415 TWh in 2024, accounting for around 1.5% of global electricity demand. This figure has been growing at 12% per year since 2017 — more than four times faster than total global electricity consumption. Under the IEA's base case scenario, data center electricity consumption is projected to more than double, reaching 945 TWh by 2030 — roughly equivalent to

Japan's entire current electricity consumption.¹ AI is identified as the most important driver of this growth.

The carbon emissions from training frontier models have risen steeply: from 0.01 tons for AlexNet (2012) to 588 tons for GPT-3 (2020), 5,184 tons for GPT-4 (2023), and 8,930 tons for Llama 3.1 405B (2024).²

Beyond energy and carbon, the broader ecological impacts of AI — including effects on biodiversity and ecosystems — remain insufficiently studied, with most research still narrowly focused on compute-related metrics. Emerging literature highlights that AI's environmental footprint is systemic and embedded in global socio-economic infrastructures, producing crosssectoral and interdependent risks. These impacts are also spatially uneven: while benefits are concentrated in advanced economies, environmental and resource-related costs are disproportionately externalized to resource-intensive regions and peripheral actors, particularly in the Global South.³ In this context it is more important than ever to promote a truly international collaboration.

At the same time, the environmental crises that AI could help address are intensifying. The Intergovernmental Panel on Climate Change has emphasized the narrowing window for limiting global warming to 1.5°C⁴. The Kunming-Montreal Global Biodiversity Framework has set ambitious targets for halting biodiversity loss by 2030.⁵ The United Nations Convention to Combat Desertification reports that up to 40% of the world's land is already degraded.⁶ And the UN World Water Development Report warns that around 4 billion people live in waterstressed countries,⁷ with the estimated AI-related water consumption that could reach 6.6 billion cubic meters by 2027.⁸

¹ International Energy Agency. "Energy and AI", 10 April 2025.

² Stanford HAI, AI Index Report 2025, Chapter 1: Research and Development. <https://hai.stanford.edu/aiindex/2025-ai-index-report>

³ Obasesam Okoi. Artificial Intelligence, the Environment and Resource Conflict: Emerging Challenges in Global Governance. June 27, 2025. <https://balsilliepapers.ca/wp-content/uploads/2025/06/Balsillie-Paper-Okoi.pdf>

⁴ IPCC. "Global Warming of 1.5 °C" <https://www.ipcc.ch/sr15/>

⁵ Decision Adopted by The Conference of the Parties to the Convention on Biological Diversity.15/4. Kunming-Montreal Global Biodiversity Framework. <https://www.cbd.int/doc/decisions/cop-15/cop-15-dec-04-en.pdf>

⁶ UNCCD. Global Land Outlook. Second Edition. Summary for Decision Makers. https://www.unccd.int/sites/default/files/2022-04/GLO2_SDM_low-res_0.pdf

⁷ The United Nations World Water Development Report 2026, Water for All People: Equal Rights and Opportunities. <https://unesdoc.unesco.org/ark:/48223/pf0000397159>

⁸ Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. 26 March 2025. <https://arxiv.org/pdf/2304.03271>

The goals identified in the 2030 Agenda for Sustainable Development have dual relationship with AI, which is simultaneously a tool and a risk factor: SDG 7 (affordable and clean energy) is challenged by data centers' surging power demand; SDG 9 (industry, innovation, and infrastructure) underpins the case for responsible AI development; SDG 13 (climate action) frames both AI's emissions contribution and its mitigation potential; SDG 14 (life below water) and SDG 15 (life on land) are affected by the mining, water consumption, and habitat disruption associated with AI infrastructure, while also standing to benefit from AI-powered monitoring and conservation.⁹

This creates a fundamental tension: the technology best positioned to help solve environmental crises is simultaneously contributing to them. Navigating this tension — maximising AI's positive environmental potential while minimising its ecological footprint — is one of the defining governance challenges of the current decade.

Why This Paper – and What Makes It Different

Several important contributions have advanced our understanding of the AI–environment nexus. The OECD's *Measuring the Environmental Impacts of Artificial Intelligence Compute and Applications* (2022) established an early methodological framework.¹⁰ The UASDF *White Paper on Global AI Environmental Impact* (2024) introduced the AI Green Index for quantifying carbon and water footprints.¹¹ The UNFCCC Technology Executive Committee's technical paper on *AI for Climate Action* (2024–2025) provided a comprehensive review of AI applications for climate mitigation and adaptation, with particular attention to least developed countries and small island developing states.¹² The ITU's *AI for Good* platform has catalyzed practical partnerships between AI innovators and development practitioners.¹³

This paper builds on these foundations but addresses five significant gaps that, taken together, define its distinctive contribution.

The cost–benefit divide. Existing documents tend to examine only one side of the equation: either AI's environmental footprint, or AI's potential as a tool for environmental solutions. No comprehensive document has yet brought both dimensions together within a single analytical framework. This paper treats the environmental cost of AI and the environmental benefit of AI as inseparable aspects of a single governance challenge. Sections 2 through 6 systematically examine both sides for each environmental domain.

The thematic gap. Climate change has dominated the discussion, understandably given its urgency. But AI's environmental implications extend well beyond climate to biodiversity conservation, land degradation, desertification, and freshwater management — areas covered by distinct international conventions (CBD, UNCCD) and requiring distinct policy responses.

These dimensions remain largely absent from existing white papers. Sections 3 through 6, each addressing one of these four domains, are designed to close this gap, with a consistent threepart structure: (a) framing the environmental challenge; (b) analyzing AI's current and potential contributions; (c) mapping sustainable solutions and corporate pathways.

The geographic gap. Existing analyses have been shaped primarily by US, European, and to some extent Chinese perspectives. Yet AI governance is increasingly a multipolar endeavor. The India AI Impact Summit 2026, held in New Delhi in February, signaled a decisive shift towards Global South leadership, with investment pledges approaching \$250 billion.¹⁴

The BRICS+ AI Alliance, launched in December 2024 at the AI Journey conference in Moscow with participation of more than 20 technology companies

⁹ UN. Transforming our world: the 2030 Agenda for Sustainable Development <https://docs.un.org/en/A/RES/70/1>

¹⁰ OECD. *Measuring the Environmental Impacts of Artificial Intelligence Compute and Applications*. November 2022. https://www.oecd.org/en/publications/measuring-the-environmental-impacts-of-artificialintelligence-compute-and-applications_7babf571-en.html

¹¹ UASDF. *White Paper on Global Artificial Intelligence Environmental Impact*. September 2024. <https://www.researchgate.net/publication/384364115>

¹² UNFCCC Technology Executive Committee. *Technical Paper on AI for Climate Action*. 2024–2025. <https://unfccc.int/tclear/tec/AI4climate.html>

¹³ ITU. *AI for Good*. <https://aiforgood.itu.int/>

¹⁴ DFRLab. "From Delhi to Geneva: What the AI Impact Summit reveals." March 2026. <https://dfrlab.org/2026/03/04/what-the-ai-impact-summit-reveals/>

from six countries, represents a new axis of cooperation in AI development.¹⁵

The 2024 BRICS Summit in Kazan adopted a framework program for AI cooperation, including the BRICS AI Research and Innovation Center and a BRICS Digital Cloud Corridor.¹⁶

These developments demand analysis that reflects the priorities and realities of India, China, Russia, and the broader BRICS+ grouping. This paper provides such analysis through its partnership structure:

- **Pahlé India Foundation** brings the perspective of the world's most populous country of 1.4 billion people and 886 million internet users,¹⁷ one of the fastest-growing digital economies¹⁸ and a rising AI power with IndiaAI Mission. India's rapid AI expansion has implications not only for economic development but also for energy demand, carbon emissions, water consumption, and long-term digital infrastructure sustainability. The Foundation is an independent “think-and-action tank” producing evidence-based research on economic reform, energy, and sustainable development — expertise that informs the paper's analysis of how rapidly developing economies can reconcile AI-driven growth with environmental resilience.¹⁹
- **Fudan University** contributes the expertise of one of China's leading research institutions, based in Shanghai — the host city of the World Artificial Intelligence Conference (WAIC), which in 2024 issued the Shanghai Declaration on Global AI Governance calling for AI to be applied to “environmental protection, resource

utilisation, energy management, and biodiversity promotion”.²⁰ A member of China's elite C9 League (often described as China's Ivy League) and of the BRICS Universities League, with recognized research strength in environmental science, public policy, and artificial intelligence, Fudan anchors the paper in the perspective of the world's secondlargest economy and a leading AI power.²¹

- **The Nature and People Foundation**²² brings the Russian perspective — that of a country with one of the world's largest digital ecosystems and where Russian remains among the most widely used content languages on the web.²³ The Foundation is the first and, to date, only Russian non-governmental organization to obtain observer status at a Conference of the Parties to the UN Convention to Combat Desertification (UNCCD), and is actively engaged across all three Rio conventions — on climate change (UNFCCC), biological diversity (CBD), and desertification (UNCCD).²⁴ It contributes both this international environmental-governance experience and the perspective of Russian institutions — including Sber, Yandex, and the AI Research Institute (AIRI).²⁵

The inclusion of these perspectives is not merely additive. It reflects a fundamental conviction: in a multipolar world, effective AI governance cannot be designed by one group of countries and exported to the rest. It must be co-created through inclusive dialogue that respects the sovereignty, priorities, and contributions of all stakeholders. The BRICS Leaders' Statement on the Global Governance of Artificial Intelligence, adopted following the 2025 Rio de Janeiro summit, places special emphasis on equitable geographic representation in global AI

¹⁵ iGlobenews. “The BRICS+ AI Alliance: Towards A Multipolar Focal Point in Artificial Intelligence.” April 2025. <https://www.iglobenews.org/brics-ai-alliance/>

¹⁶ BRICS Council. “Joint BRICS Projects in Artificial Intelligence.” 2025. <https://bricscouncil.ru/en/analytics/sovместnye-proekty-stran-briks-v-oblasti-iskusstvennogo-intellekta>

¹⁷ Government of India. The Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023). New Delhi: Ministry of Electronics and Information Technology; 2023.

¹⁸ Indian Council for Research on International Economic Relations. The Digital Economy in India: Growth, Transformation, and Future Prospects. New Delhi: Indian Council for Research on International Economic Relations; 2019.

¹⁹ Pahlé India Foundation. <https://pahleindia.org/>

²⁰ People's Republic of China. Ministry of Foreign Affairs. Full text: Shanghai Declaration on Global AI Governance. https://www.mfa.gov.cn/eng/xw/zyxw/202407/t20240704_11448351.html

²¹ Fudan University. <https://www.fudan.edu.cn/en/>

²² Nature and People Foundation. <https://naturepeople.ru/>

²³ W3Techs. Internet Society Foundation. Russian is approximately the second–third most used content language on websites. <https://www.isocfoundation.org/2023/05/what-are-the-most-used-languages-on-the-internet/>

²⁴ UN News. Nature and People Foundation: the climate agenda unites all countries of the world, interview, December 2024, in Russian. <https://news.un.org/ru/story/2024/12/1459236>

²⁵ Sber / AIRI. Eco2AI: <https://github.com/sb-ai-lab/Eco2AI>; Eco4cast: <https://github.com/AIRIInstitute/eco4cast>

governance — a principle this paper embodies in its authorship and analysis.²⁶

The regulatory gap. The policy landscape is evolving faster than the literature can track.

The EU AI Act entered application in 2025, but dropped proposed requirements for CO₂ emissions transparency in its final text.²⁷ China has issued multiple AI governance regulations since 2023. India's DPDP Act and forthcoming AI governance framework are shaping a distinctive approach. Russia updated its National Strategy for AI Development through 2030 in 2024.²⁸

The interplay between AI regulation, climate policy, and environmental law — across these and other jurisdictions — has not yet been systematically mapped. Section 7 provides this comparative analysis.

The institutional gap. The architecture for AI governance at the international level is itself in flux. The WSIS+20 review, concluded in December 2025, made the Internet Governance Forum a permanent forum of the United Nations.²⁹

The UN General Assembly adopted Resolution A/RES/79/325 in August 2025, establishing the Independent International Scientific Panel on AI and the Global Dialogue on AI Governance.³⁰

The inaugural Global Dialogue, co-facilitated by Estonia and El Salvador, will convene in Geneva in July 2026 alongside the AI for Good Global Summit.³¹

Separately, the World Artificial Intelligence Conference (WAIC) in Shanghai continues to serve as a key platform for international AI cooperation; at WAIC 2024, UNIDO inaugurated its Global Alliance

on AI for Industry and Manufacturing Center of Excellence.³²

The environmental dimension of these emerging governance structures remains underdefined.

This paper — as a contribution from the Dynamic Coalition on Environment of the IGF — aims to help fill that gap. Section 8 analyses the roles and responsibilities of public sector, private sector, civil society, and international organizations in this evolving landscape.

In February 2025, at the AI Action Summit in Paris, France, UNEP, and ITU launched the Coalition for Environmentally Sustainable Artificial Intelligence, bringing together over 100 partners — including 37 technology companies, eleven countries, and five international organizations — for collaborative action on standardized measurement, life-cycle analysis frameworks, and prioritization of sustainable AI research.³³ In June 2025, UNEP released its Sustainable Procurement Guidelines for Data Centers and Servers, the first international framework designed to help governments reduce the energy and water consumption of data center operations.³⁴ These developments signal a growing institutional momentum — but also underscore that the environmental governance of AI is still in its formative stages, with significant gaps remaining in measurement, disclosure, and enforcement.

UNEP's 2024 issue note Artificial Intelligence End-to-End explicitly called for comprehensive assessment of AI's environmental impact across the full lifecycle — from material extraction and hardware manufacturing through operation and disposal — noting that “governments are racing to develop national AI strategies but rarely do they take the environment and sustainability into account.”³⁵

²⁶ BRICS Council, op. cit. See also: BRICS Leaders' Statement on the Global Governance of Artificial Intelligence, adopted following the 2025 Rio de Janeiro summit.

²⁷ EU AI Act, Regulation (EU) 2024/1689. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

²⁸ Government of the Russian Federation. National Strategy for AI Development through 2030 (updated 2024). <http://government.ru/en/>

²⁹ Internet Society. WSIS+20 Reaffirms Multistakeholder Governance and a Lasting IGF. December 2025. <https://www.internetsociety.org/blog/2025/12/ws20-reaffirms-multistakeholder-governance-and-a-lasting-igf/>

³⁰ UN General Assembly. Resolution A/RES/79/325. <https://docs.un.org/en/A/RES/79/325>

³¹ UN. Global Dialogue on AI Governance. <https://www.un.org/global-dialogue-ai-governance/en>

³² UNIDO. “UNIDO, Shanghai Municipality and MIIT launch AIM Global Centre of Excellence at World AI Conference 2024.” July 2024. <https://www.unido.org/news/unido-shanghai-municipality-and-miit-launch-aimglobal-centre-excellence-world-ai-conference-2024>

³³ UNEP. “New Coalition aims to put Artificial Intelligence on a more sustainable path.” 11 February 2025. <https://www.unep.org/news-and-stories/press-release/new-coalition-aims-put-artificial-intelligence-moresustainable-path>

³⁴ UNEP. “UNEP releases guidelines to curb the environmental impact of data centres.” 12 June 2025. <https://www.unep.org/technical-highlight/unep-releases-guidelines-curb-environmental-impact-data-centres>

³⁵ UNEP. “Artificial Intelligence (AI) end-to-end: The Environmental Impact of the Full AI Lifecycle Needs to be Comprehensively Assessed.” September 2024. <https://www.unep.org/resources/report/artificial-intelligenceai-end-end-environmental-impact-full-ai-lifecycle-needs-be>

This paper responds to that call by providing, for the first time, an integrated analysis that spans AI's full lifecycle (Section 2) and its applications across four environmental domains (Sections 3–6).

The International Telecommunication Union (ITU), through its Green Digital Action initiative, has launched a dedicated Sustainable AI working group and is developing technical standards for environmental sustainability in AI and data center operations.³⁶ ITU reports examine AI and the environment alongside associated technical standards development, establishing the groundwork for internationally harmonized measurement frameworks.³⁷ ISO has published TR 20226:2025, which incorporates sustainability considerations across the full AI lifecycle, including computational load, resource use, carbon impacts, pollution, disposal, and operational characteristics.³⁸

Scope, Objectives, and Structure

The paper is organized into nine sections, each serving a distinct analytical purpose.

This introduction (**Section 1**) establishes the rationale and scope. **Section 2** (*AI Models*) assesses the full life-cycle environmental impact of AI — covering energy consumption, carbon emissions, water usage, and sustainability metrics and tools.³⁹

Sections 3 through 6 each address a major environmental domain — climate change, biological diversity, desertification and land degradation, and water resources — aligned with the corresponding multilateral environmental agreements (UNFCCC, CBD, UNCCD, SDG 6).

Section 7 (*Policy and Regulatory Considerations*) provides a comparative analysis across six jurisdictions: the BRICS grouping, China, the European Union, India, Russia, and the United States. **Section 8** (*Stakeholder Roles*) examines how public sector institutions, international organizations, private sector actors, civil society organizations — including women, indigenous peoples, and local communities — can drive implementation. **Section 9** (*Next Steps*) identifies challenges, emerging trends, and concrete recommendations.

Each thematic section (Sections 3–6) follows a consistent three-part structure: why the environmental challenge matters; how AI is currently contributing or could contribute to addressing it; and what pathways exist for aligning innovation with sustainability goals, including corporate commitments and case studies. Where AI-enabled environmental benefits are discussed, they should be interpreted alongside the environmental costs associated with the AI systems required to deliver them.

Therefore, the **objectives** of the white paper are the following:

- (1) Analyze the full life-cycle environmental impact of AI models and data center infrastructure, including energy consumption, carbon footprint, and water usage.
- (2) Research how AI is being applied — and could be further applied — to address climate change, biodiversity conservation, land degradation, and water resource management.
- (3) Map the policy and regulatory landscape across key jurisdictions (BRICS, China, India, Russia).
- (4) Identify the roles of public sector, private sector, civil society, and international organizations in driving sustainable AI.
- (5) Outline challenges, emerging trends, and recommendations for the path forward.

Timeliness

This paper is designed for presentation at the **AI for Good Global Summit**⁴⁰ and the inaugural **Global Dialogue on AI Governance** in Geneva, 7–10 July 2026. It is also intended as a contribution to the **World Artificial Intelligence Conference (WAIC)** in Shanghai. These events in 2026 mark a pivotal moment. The international community is establishing the institutional architecture for AI governance — from the UN Scientific Panel on AI to the BRICS+ AI Alliance, from the EU AI Act to China's AI governance regulations. The environmental dimension of this architecture remains underdefined, and this paper aims to

³⁶ ITU. Green Digital Action. <https://www.itu.int/initiatives/green-digital-action/>

³⁷ ITU. AI and the Environment - International Standards for AI and the Environment. 2024 Report https://www.itu.int/dms_pub/itu-t/opb/env/T-ENV-ENV-2024-1-PDF-E.pdf

³⁸ ISO/IEC TR 20226:2025. Information technology — Artificial intelligence — Environmental sustainability aspects of AI systems. <https://www.iso.org/standard/86177.html>

³⁹ IEA, Energy and AI, op. cit.; Schneider Electric 2023 estimates (80% inference / 20% training split).

⁴⁰ ITU. AI for Good Global Summit 2026. <https://aiforgood.itu.int/summit26/>

contribute substantive, evidence-based content to fill that gap.

Critically, this is being done in a context where the global AI landscape is no longer dominated by a single center of gravity. China leads the BRICS in AI development, with the realized economic potential of generative AI across BRICS nations projected to reach \$600 billion by 2030.⁴¹ India, with the world's largest population and a rapidly growing AI ecosystem, hosted the first major AI governance summit in the Global South in February 2026. Russia, as the country with one of the most used languages on the web and a developed domestic AI ecosystem, has initiated the BRICS+ AI Alliance and contributes uniquely to multilingual AI development — a significant consideration given that approximately 90% of AI training data is in English, despite English being native to only about 18% of the world's population.^{42,43}

The effective governance of AI's environmental impact cannot be designed from any single national or cultural vantage point. It requires the kind of inclusive, multistakeholder dialogue that the IGF was created to foster and that the Global Dialogue on AI Governance is mandated to advance. This paper is offered as a contribution to that dialogue.

A Note on Approach

This paper is written in the spirit of the IGF's founding mandate: to provide a space for open, inclusive, and evidence-based dialogue on digital governance issues. It does not advocate for a single policy position, but seeks to present the best available evidence, highlight areas of consensus and uncertainty, and identify options for action adaptable to diverse national contexts.

The authors recognize that the relationship between AI and the environment is characterized by rapid change, significant data gaps, and genuine disagreements among experts. The IEA itself notes that there are no comprehensive global datasets on data center electricity consumption or emissions, and that even historical estimates can be widely divergent. Where data is uncertain, we say so. Where competing perspectives exist — such as whether economic constraints will naturally limit AI's energy growth, or whether external regulation is required — we present them fairly.

We invite readers — policymakers, researchers, industry leaders, and civil society advocates — to engage with this paper critically and constructively, and to contribute to the ongoing dialogue that will shape how humanity governs the relationship between its most powerful technology and its most essential resource: the natural environment.

⁴¹ Yakov and Partners study, cited in BRICS Council, op. cit.

⁴² W3Techs, op. cit.

⁴³ Stanford HAI, AI Index Report 2024. <https://aiindex.stanford.edu/report/>

Section 2.

AI Models

2.1	Life Cycle Assessment Overview for AI Models	22
2.2	Carbon Footprint of AI Models and Data Centers	27
2.3	Water Footprint of AI Models and Data Centers	30
2.4	Metrics and Tools for Sustainable Assessment	32



2.1 Life Cycle Assessment Overview for AI Models

Why Life Cycle Assessment Matters for AI and Who Should Conduct It

A life-cycle view is essential for understanding AI's environmental impact. The footprint of AI systems begins before model training and continues after deployment: extraction and processing of critical minerals, chip fabrication, server manufacturing, data-center construction, model training, model inference, cooling, water use, equipment replacement, and electronic waste. A narrow focus on training emissions risks understating the environmental costs of large-scale AI because inference, hardware refresh cycles, supply-chain emissions and water consumption can become dominant as AI applications are widely deployed.

Life cycle assessment (LCA) is a standardized methodology for evaluating the environmental impacts of a product, system, or process across its entire lifespan — from raw material extraction through manufacturing, operation, and end-of-life disposal. Codified in ISO 14040 and ISO 14044, LCA has been applied for decades in manufacturing, construction, and energy, providing the scientific basis for environmental labeling, regulatory compliance, and design optimisation.⁴⁴

Applying LCA to AI systems is both urgent and methodologically challenging. As AI becomes embedded in critical infrastructure — energy grids, healthcare, transportation, financial systems — the environmental consequences of its deployment accumulate across supply chains and geographies. Yet unlike a physical product such as a building or

a vehicle, an AI system is a composite of hardware, software, data, and services, each with its own lifecycle and each shared across multiple users and applications. This makes defining system boundaries, allocating impacts, and establishing functional units significantly more complex than in traditional LCA practice.⁴⁵

LCA of AI should ideally be conducted by three categories of actors. **AI developers** — companies that train and deploy models — are best positioned to measure operational impacts (energy, water, compute hours) because they have direct access to training logs and infrastructure data. **Hardware manufacturers** — producers of GPUs, CPUs, servers, and cooling systems — hold the data on embodied carbon, material composition, and supply chain impacts that are largely invisible to downstream users. **Independent researchers and standards bodies** — organizations like the ITU, ISO, OECD, and academic institutions — play the critical role of developing harmonized methodologies, validating industry claims, and filling data gaps through transparent, reproducible studies.⁴⁶

The results of AI LCA should serve multiple audiences and purposes: informing model selection by developers and enterprise buyers (choosing the least impactful model for a given task); enabling regulatory compliance under emerging frameworks such as the EU Energy Efficiency Directive and the ITU-T L.1801 standard; supporting corporate sustainability reporting under frameworks such as the GHG Protocol and CSRD; and providing the evidentiary base for policy interventions — from carbon pricing to green procurement requirements.⁴⁷

⁴⁴ ISO 14040:2006, ISO 14044:2006. <https://www.iso.org/standard/37456.html>

⁴⁵ Dokic, D. et al. "Towards Sustainability of AI: A Systematic Review of Existing LCA Approaches." HICSS-57, 2024. <https://aisel.aisnet.org/hicss-57/in/sustainability/2/>

⁴⁶ ITU, "Measuring What Matters: How to Assess AI's Environmental Impact," July 2025. <https://www.itu.int/hub/publication/s-gen-gda-001-2025/>; OECD, Measuring the Environmental Impacts of AI Compute and Applications, 2022. <https://www.oecd.org/en/publications/7babf571-en.html>

⁴⁷ ITU-T L.1801, February 2026. <https://www.itu.int/epublications/publication/itu-t-l-1801-2026-02>

Three Phases of AI's Environmental Lifecycle

The environmental impact of AI spans **three distinct phases**, each generating different types of environmental pressure and requiring different measurement approaches.

The upstream phase encompasses the extraction of raw materials (rare earth metals, lithium, cobalt, gallium, copper, gold), the manufacturing of semiconductors, GPUs, CPUs, memory, storage devices, and servers, the construction of data center buildings and cooling infrastructure, and the transportation of components across global supply chains. UNEP estimates that producing a single 2-kilogram computer requires approximately 800 kilograms of raw materials.⁴⁸ China controls the majority of global refined gallium supply, creating both environmental and geopolitical dependencies. The mining and smelting of these materials involve habitat destruction, water pollution, and long-term toxic effects on aquatic ecosystems, with impacts disproportionately concentrated in the Global South.⁴⁹

The semiconductor industry provides a particularly illustrative example of how these upstream environmental pressures are materializing within emerging manufacturing hubs. India is positioning itself as a strategic node in global semiconductor and electronics supply chains through the SEMICON India programme, supported by investments equivalent to \$7.9 billion. Announced investments across semiconductor design, fabrication, and electronics manufacturing ecosystems already exceed \$15 billion.⁵⁰ While such expansion strengthens supply-chain resilience and supports the growing demand for AI hardware, it also carries significant environmental implications. Semiconductor fabrication is highly resource-intensive: a typical fabrication facility may require more than five million gallons of ultra-pure water per day, which in turn necessitates at least eight million gallons of municipal water input.⁵¹ As

AI deployment accelerates, demand for advanced chips, GPUs, rare earth minerals, copper, lithium, cobalt, and related infrastructure is expected to rise substantially.⁵² Consequently, the environmental footprint of semiconductor manufacturing extends well beyond electricity consumption, encompassing intensive water use, hazardous chemical handling, mining-related ecological degradation, and industrial waste generation.

The operational phase includes all energy and water consumed during the use of AI systems: model training, inference, data storage, networking, and cooling. This phase also includes the idle energy consumption of data center infrastructure — the energy drawn by servers, cooling systems, and power distribution even when not actively processing AI workloads. The BLOOM lifecycle study found that this idle consumption represented 29% of total training-phase emissions, a figure often overlooked in assessments focused solely on GPU compute.⁵³

The downstream phase encompasses hardware obsolescence, electronic waste generation, and end-of-life disposal or recycling. AI hardware refresh cycles in data centers are typically 3–5 years, driven by rapid improvements in GPU performance. Global e-waste reached 62 million tons in 2022 and continues to grow. AI-specific e-waste is not tracked separately in any national or international reporting system, creating a significant gap in understanding the full environmental consequences of AI deployment.⁵⁴

India's e-waste sector illustrates how these downstream challenges intersect with both environmental management and social equity concerns. As one of the world's largest generators of electronic waste, India has developed a recycling market with an estimated value exceeding \$1 billion in 2021. However, a substantial share of e-waste processing continues to occur within the informal sector, where women are often concentrated in the lowest-paid and most hazardous activities,

⁴⁸ UNEP, AI has an environmental problem September 2024.

<https://www.unep.org/news-and-stories/story/ai-has-environmental-problem-heres-what-world-can-do-about>

⁴⁹ The ecology of AI, Carbon Neutral Systems, Springer, December 2025. <https://link.springer.com/article/10.1007/s44438-025-00018-8>

⁵⁰ Government of India. India Semiconductor Mission Programme Documents. New Delhi: Ministry of Electronics and Information Technology, Government of India; 2024.

⁵¹ Kumar MJ. Is India going to be a major hub of semiconductor chip manufacturing? IETE Technical Review. 2021;38(3):279-281. doi:10.1080/02564602.2021.1916166

⁵² International Energy Agency. Critical Minerals Outlook 2024. Paris: International Energy Agency; 2024.

⁵³ Wu, C. et al. "Sustainable AI," MLSys 2022. <https://arxiv.org/abs/2111.00364>;

Luccioni, A.S. et al. "Estimating the Carbon Footprint of BLOOM," JMLR 24(253), 2023. <https://jmlr.org/papers/v24/23-0069.html>

⁵⁴ Dokic et al., HICSS-57, 2024, *ibid*.

including waste collection and wire stripping.⁵⁵ As AI adoption accelerates and hardware replacement cycles shorten, growing volumes of discarded servers, GPUs, batteries, cooling equipment, and consumer devices are likely to intensify existing waste management pressures. Without significant strengthening of formal recycling systems and circular economy practices, the environmental and social burdens associated with AI-related e-waste are expected to increase.

The three-phase lifecycle can also be organized analytically into **four layers**, which is useful for assigning measurement responsibility and for connecting AI's footprint to its downstream effects:

- 1) compute infrastructure, including chips, servers, storage, and networking equipment;
- 2) data-center operations, including electricity, cooling, and water;
- 3) model development and use, including training, fine-tuning, and inference; and
- 4) downstream applications, including whether AI reduces or increases energy, material, and transport demand in other sectors.

This integrated framework allows us to assess both the environmental footprint of AI and the environmental benefits (analyzed in Sections 3 through 6) that AI may generate in climate, biodiversity, land and water governance.

The State of LCA Practice for AI

Systematic life cycle assessment of AI technology is still in its infancy. A 2025 review in *Procedia CIRP* concluded that LCA of AI “should not be limited to calculating the CO₂-equivalent, but take into account different relevant environmental impact indicators from resource use depletion to human toxicity to water consumption,” and that a dedicated LCA methodology for AI has yet to be developed.⁵⁶

The 2024 systematic review by Dokic, Groen In'T Woud, and Maass (HICSS-57) similarly found that

existing approaches for conducting LCA on AI applications are fragmented and incomplete. While the energy side of training and inference can be roughly estimated using tools such as CodeCarbon⁵⁷, a realistic estimate of AI's total environmental impact — including critical raw materials, hardware manufacturing, maintenance, and disposal does not seem feasible today.

Despite these challenges, a small but growing number of actors are conducting and publishing LCA results:

A group of researchers conducted the first comprehensive lifecycle carbon assessment of a large language model (**BLOOM**, 176B parameters) in 2023, covering dynamic power consumption during training, embodied hardware emissions, idle data center consumption, and real-time deployment tracking via CodeCarbon over 18 days. The total lifecycle footprint was approximately 50.5 tons CO₂eq — roughly double the training-only figure of 24.7 tons. The breakdown revealed that just 49% of emissions came from active GPU computation; 29% came from idle data center infrastructure, and 22% from hardware manufacturing.⁵⁸

Mistral AI, in collaboration with Carbone 4 and the French ecological transition agency (ADEME), published in July 2025 what it described as the first comprehensive lifecycle analysis of an AI model compliant with both the GHG Protocol Product Standard and ISO 14040/44. The study assessed Mistral Large 2 across three impact categories: greenhouse gas emissions (20.4 kt CO₂eq over 18 months of use), water consumption (281,000 m³), and resource depletion (660 kg Sb eq). Crucially, the study also published marginal inference impacts: a single 400-token response via the Le Chat assistant generates 1.14 g CO₂e, consumes 45 mL of water, and uses 0.16 mg Sb eq. This is the first time a commercial AI provider has published inference-level environmental data in a standards-compliant format.⁵⁹

Microsoft and WSP published in *Nature* (2025) the first cradle-to-grave LCA of data center cooling technologies, introducing the annualized virtual core (Vcore) as a functional unit for comparing different

⁵⁵ United Nations Institute for Training and Research. Global E-waste Monitor 2024. Geneva: United Nations Institute for Training and Research; 2024.

⁵⁶ “LCA of AI Applications: Research Gaps and Opportunities,” *Procedia CIRP*, 2025. <https://www.sciencedirect.com/science/article/pii/S2212827125003749>

⁵⁷ <https://codecarbon.io/>

⁵⁸ Luccioni et al., *JMLR*, 2023, *ibid*.

⁵⁹ Mistral AI, “Our contribution to a global environmental standard for AI,” July 2025. <https://mistral.ai/news/our-contribution-to-a-global-environmental-standard-for-ai>

data center configurations. The study evaluated four cooling techniques across their full lifecycle, finding that advanced cooling (cold plates, immersion) can substantially reduce GHG emissions, energy demand, and water consumption compared to traditional air cooling. Microsoft has made the LCA tool available as a free, open-source Excel-based instrument for any data center operator.⁶⁰

Sber and AIRI (Russia) developed the Eco2AI open-source library, which monitors CO₂-equivalent emissions during code execution, incorporating regional emission coefficients. While not a full LCA, Eco2AI represents one of the few tools originating from the BRICS countries and has been validated on real-world applications, including waste detection neural networks developed in collaboration with ITMO and Sber AI Lab.⁶¹

China provides an important case of **infrastructure-based life-cycle governance**. The national “Eastern Data, Western Computing” strategy establishes eight national computing-hub nodes — Beijing-Tianjin-Hebei, the Yangtze River Delta, the Guangdong-Hong Kong-Macao Greater Bay Area, Chengdu-Chongqing, Inner Mongolia, Guizhou, Gansu and Ningxia — supported by ten national data-centre clusters. The strategy guides computing demand from China’s eastern regions towards western computing resources in an orderly manner while coordinating energy availability, latency requirements and regional development objectives. This approach is particularly relevant to AI because the environmental effects of inference depend not only on the model itself, but also on where workloads are processed, whether renewable electricity is available, how cooling is managed and whether computing resources are efficiently utilized.

This case shows that sustainable AI cannot be reduced to model efficiency alone. It also requires spatial planning, power-system coordination, infrastructure standards and workload scheduling. For China, the key governance challenge is to balance the rapid growth of AI computing demand with green electricity access, utilisation rates, water stress and regional development priorities.

Methodological Challenges, Disagreements, and Data Deficits

LCA for AI is a young field, and significant methodological questions remain unresolved. These are not merely academic — they directly affect the comparability, reliability, and policy relevance of any environmental claims made about AI systems.

Defining the functional unit. In traditional LCA, the functional unit is the reference to which all environmental impacts are normalized (e.g., “1 kWh of electricity delivered”). For AI, there is no consensus on the appropriate functional unit. Options include: per model trained, per inference query, per token generated, per virtual core-year (as in the Microsoft/WSP study), or per unit of useful output. The choice of functional unit fundamentally changes what is being measured and compared, making cross-study comparisons unreliable.⁶²

System boundary definition. There is no agreement on where the boundary of an AI system ends. Should the LCA of a cloud-hosted AI model include the user’s terminal device? The network transmission? The data collection and curation pipeline? The BLOOM study included hardware manufacturing and idle infrastructure but excluded user terminals; the Mistral study explicitly excluded terminals. Ligozat et al. (2022) proposed a dual framework — software lifecycle (data acquisition, training, inference) mapped onto hardware lifecycle (manufacturing, use, disposal) — but this has not been widely adopted.⁶³

The GPU data gap. Perhaps the most critical data deficit is the absence of reliable lifecycle inventory data for modern GPUs. Mistral AI noted explicitly that “a reliable life-cycle inventory of GPUs is yet to be made, as their embodied impacts had to be approximated but account for a significant portion of total impacts.” GPU manufacturers — primarily NVIDIA, AMD, and Intel — do not publish detailed environmental product declarations (EPDs) for their AI accelerators. This means that the embodied carbon of the single most important component in

⁶⁰ Alissa, H. et al. “Using LCA to drive innovation for sustainable cool clouds,” Nature, 2025.

<https://www.nature.com/articles/s41586-025-08832-3>; WSP insight: <https://www.wsp.com/en-us/insights/2025-lca-data-centers>

⁶¹ Budenny, S. et al. “Eco2AI,” 2022. <https://arxiv.org/abs/2208.00406>; <https://github.com/sb-ai-lab/Eco2AI>

⁶² Dokic et al., HICSS-57, 2024, *ibid.*; Alissa et al., Nature, 2025, *ibid.*

⁶³ Ligozat, A.-L. et al. “Unraveling the hidden environmental impacts of AI solutions,” Sustainability, 14(9), 5172, 2022. <https://www.mdpi.com/2071-1050/14/9/5172>

AI infrastructure cannot be precisely quantified by anyone outside the manufacturer.⁶⁴

Scope 3 opacity. For AI systems, Scope 3 emissions — encompassing hardware manufacturing, logistics, and end-of-life processing — are often the largest category but the least transparent. No major tech company currently provides AI-specific Scope 3 data. Corporate sustainability reports aggregate all data center operations without distinguishing AI from non-AI workloads.⁶⁵

Beyond carbon: neglected impact categories. The Procedia CIRP review (2025) and the HICSS-57 study (2024) both emphasize that existing AI environmental assessments focus almost exclusively on greenhouse gas emissions (GWP), while other LCA impact categories — abiotic resource depletion, human toxicity, acidification, eutrophication, water consumption, land use, and biodiversity loss — are rarely assessed. The Mistral study represents a step forward by including water consumption and resource depletion alongside carbon, but this remains the exception.⁶⁶

Inference measurement. The ITU report *Measuring What Matters* (July 2025) identified inference as a critical gap: most measurement focuses on training, while inference — which accounts for the majority of lifecycle emissions in deployed systems — is poorly and inconsistently measured.⁶⁷ The Mistral study's publication of per-query inference impacts is the first standards-compliant data point of its kind, but it stands alone.

Towards Standardization

Despite these gaps, a convergence of standards and institutional initiatives is beginning to take shape. The ITU published Recommendation ITU-T L.1801 in February 2026 — the first international standard specifically for assessing the environmental impact of AI systems. It builds on existing ITU methodologies for ICT lifecycle assessment (L.1410) and enabling effects (L.1480), adapted to AI's specific characteristics.⁶⁸

ISO/IEC published TR 20226:2025, which incorporates sustainability considerations across the full AI lifecycle, including computational load, resource use, carbon impacts, pollution, disposal, and operational characteristics.⁶⁹

The Coalition for Environmentally Sustainable AI, launched in February 2025 by France, UNEP, and ITU with over 100 partners, has prioritized the development of standardized methods and metrics and comprehensive LCA frameworks.⁷⁰ France has published the Frugal AI methodology through AFNOR, which Mistral's LCA followed and which is being proposed as a model for international harmonization.⁷¹ The United Nations Environment Assembly adopted Resolution 7/9 on the environmental sustainability of AI systems in December 2025, providing the highest-level multilateral mandate for this work.⁷²

These developments suggest that within the next two to three years, AI-specific LCA standards will mature to the point where mandatory environmental disclosure becomes technically feasible. The question is whether political will and industry cooperation will match the pace of standards development.

⁶⁴ Mistral AI, July 2025, *ibid.*

⁶⁵ De Vries, A. "The carbon and water footprints of data centers," *Patterns*, December 2025. [https://www.cell.com/patterns/fulltext/S2666-3899\(25\)00278-8](https://www.cell.com/patterns/fulltext/S2666-3899(25)00278-8)

⁶⁶ Procedia CIRP, 2025, *ibid.*; Mistral AI, 2025, *ibid.*

⁶⁷ ITU, "Measuring What Matters," July 2025. <https://www.itu.int/hub/publication/s-gen-gda-001-2025/>

⁶⁸ ITU-T L.1801, February 2026. <https://www.itu.int/epublications/publication/itu-t-l-1801-2026-02-guidelinesfor-assessing-the-environmental-impact-of-artificial-intelligence-systems>

⁶⁹ ISO/IEC TR 20226:2025. <https://cdn.standards.iteh.ai/samples/iso/iso-iec-tr-20226-2025/0a8d6239489c40aeb3f7c734bd77c91b/iso-iec-tr-20226-2025.pdf>

⁷⁰ Coalition for Sustainable AI, February 2025.

<https://www.unep.org/news-and-stories/press-release/newcoalition-aims-put-artificial-intelligence-more-sustainable-path>

⁷¹ AFNOR. <https://www.afnor.org/en/news/artificial-intelligence/reference-framework-reduce-environmental-impact-ai/>

⁷² UNEA Resolution 7/9, December 2025. <https://docs.un.org/ru/UNEP/EA.7/Res.9>

2.2 Carbon Footprint of AI Models and Data Centers

The carbon footprint of AI is driven by model size, training frequency, inference volume, hardware efficiency, data-center utilisation, cooling systems and the carbon intensity of electricity. As generative AI moves from experimental use to large-scale deployment, inference is likely to account for a growing share of AI-related energy demand. Carbon management therefore needs to combine model-level efficiency with system-level measures such as poweraware scheduling, data-center heat management, renewable procurement and regional grid decarbonisation.

The Growth Trajectory

The carbon footprint of training frontier AI models has grown at an extraordinary pace. Training emissions have risen from 0.01 tons of CO₂ for AlexNet (2012) to 588 tons for GPT-3 (2020), 5,184 tons for GPT-4 (2023), and 8,930 tons for Meta’s Llama 3.1 405B (2024). Training compute is doubling approximately every five months, dataset sizes for large language models every eight months, and the power required for training annually.⁷³

At the system level, the most recent estimates suggest that AI systems could collectively be responsible for 32.6–79.7 million tons of CO₂ emissions in 2025 — a footprint comparable to that of New York City.⁷⁴ The IEA projects total data center electricity consumption to more than double from 415 TWh in 2024 to 945 TWh by 2030, with data center emissions potentially reaching 300 Mt CO₂ by 2035 in the base case, or up to 500 Mt in a high-growth scenario. The IEA describes AI as the

most important driver of this growth. In the United States, data centers are expected to consume more electricity by 2030 than the production of aluminum, steel, cement, chemicals, and all other energy-intensive goods combined.⁷⁵

India’s installed data-center capacity exceeded approximately 1 GW in 2025 with edge computing requirements projected to double by 2028.⁷⁶ Data centers currently account for approximately 0.5% of India’s total electricity consumption, a share that could rise to nearly 3% by 2030 as capacity expands.⁷⁷ In rapidly digitizing economies such as India, this growth intersects with already rising electricity demand associated with industrialization, cooling requirements, urbanization, and electrification.

Carbon Footprint Is a Design Variable

A central finding of the research literature is that the carbon footprint of AI is not a simple function of model size. It depends heavily on three variables that are within human control: model architecture and optimisation, hardware efficiency, and the carbon intensity of the energy grid.

Model architecture matters more than scale.

Du et al. (Google Research, 2021) demonstrated that GLaM — a model with 1.2 trillion parameters — consumed only one-third of the energy used to train GPT-3 (175 billion parameters) while achieving better overall benchmark performance. BLOOM (176B parameters), trained on a French nuclear-powered supercomputer, produced approximately 24.7 tons of CO₂eq, compared with 552 tons for

⁷³ Stanford HAI, AI Index Report 2025. <https://hai.stanford.edu/ai-index/2025-ai-index-report>

⁷⁴ De Vries, Patterns, December 2025, *ibid*.

⁷⁵ IEA, Energy and AI, April 2025. <https://www.iea.org/reports/energy-and-ai>; Carbon Brief, September 2025. <https://www.carbonbrief.org/ai-five-charts-that-put-data-centre-energy-use-and-emissions-into-context/>

⁷⁶ Data Center Dynamics. India Data Centre Market Reports. London: Data Center Dynamics; 2025.

⁷⁷ Office of the Principal Scientific Adviser to the Government of India. Democratising Access to AI Infrastructure. Version 3.0. New Delhi: Office of the Principal Scientific Adviser to the Government of India; 2025 Dec 29.

GPT-3, which was trained on older and less efficient hardware in a region with a more carbonintensive grid.⁷⁸ More recently, DeepSeek demonstrated that innovative optimisation techniques can reduce training costs by a factor of 20 while maintaining comparable performance.⁷⁹

The Applied Energy study (December 2025), which performed full lifecycle assessments comparing GPT-4o, Llama, and DeepSeek across multiple data center types and geographic locations, confirmed that model choice, geographic location, and grid carbon intensity are the three dominant variables determining generative AI's environmental impact.⁸⁰

Hardware efficiency is improving. Machine learning hardware energy efficiency has improved by approximately 40% per year, and inference costs have fallen over 280-fold since late 2022. At the same time, however, the total demand for compute is growing faster than efficiency gains, resulting in a net increase in absolute energy consumption and emissions. This is a classic rebound effect: efficiency improvements lower the per-unit cost of AI, which stimulates greater demand and total resource consumption.⁸¹

Regional energy sources are decisive. The carbon intensity of the electricity grid where a model is trained or deployed can vary by an order of magnitude. AWS data centers in Sweden or France emit a fraction of the CO₂ of comparable facilities in South Africa, Hong Kong, or Australia. This has direct policy implications: the geographic siting of AI infrastructure is an environmental decision.⁸²

Inference: The Larger and Overlooked Share

A critical but often overlooked finding is that **inference — not training — is the dominant source of operational emissions** for deployed AI systems. A Facebook AI study (2022), based on production data from Facebook data centers, found that inference accounts for approximately 65% of total AI lifecycle emissions, compared with 35% for training. Schneider Electric estimated an 80/20 split for 2023 workloads. As AI systems are deployed at scale — serving billions of queries daily — this ratio can only increase.⁸³

The Mistral AI LCA (July 2025) provided the first standards-compliant per-query inference data: a single 400-token response generates 1.14 grams of CO_{2e}, consumes 45 millilitres of water, and uses 0.16 milligrams of antimony equivalent. While individually small, these marginal impacts accumulate across billions of daily queries worldwide.⁸⁴

The ACM FAccT study by Luccioni, Jernite, and Strubell (2024) systematically compared energy consumption across different AI task types, finding that image generation is approximately 1,500 times more energy-intensive than text classification. This granularity is critical for policy: not all AI use carries the same environmental weight.⁸⁵

⁷⁸ Du, N. et al. "GLaM: Efficient Scaling of Language Models with Mixture-of-Experts," 2021. <https://arxiv.org/abs/2112.06905>; Patterson, D. et al. "Carbon Emissions and Large Neural Network Training," 2021. <https://arxiv.org/abs/2104.10350>; Luccioni et al., JMLR, 2023, *ibid*.

⁷⁹ Wang, Kantarcioglu. A Review of DeepSeek Models' Key Innovative Techniques, March 2025. <https://arxiv.org/html/2503.11486v1>

⁸⁰ Generative AI impact assessment through LCA of multiple data center typologies, Applied Energy, December 2025. <https://www.sciencedirect.com/science/article/abs/pii/S0306261925020185>

⁸¹ Stanford HAI, AI Index Report 2025, *ibid*.; IEEE Spectrum, "The State of AI 2025," May 2025. <https://spectrum.ieee.org/ai-index-2025>

⁸² IEA, Energy and AI, 2025, *ibid*.

⁸³ Wu, C. et al. "Sustainable AI," MLSys 2022. <https://arxiv.org/abs/2111.00364>

⁸⁴ Mistral AI, July 2025, *ibid*.

⁸⁵ Luccioni, A.S. et al. "Power Hungry Processing," ACM FAccT 2024. <https://dl.acm.org/doi/10.1145/3630106.3658542>

The Transparency Deficit

Perhaps the most consequential finding across the literature is the inadequacy of corporate environmental disclosure. No major technology company currently separates AI-specific emissions from general data center operations in its sustainability reporting. The EU AI Act, which entered application in 2025, dropped proposed requirements for CO₂ emissions transparency in its final text. The EU Energy Efficiency Directive introduced mandatory reporting for data centers with 500 kW or more, but does not require AI-specific disaggregation.⁸⁶

Without such transparency, neither markets nor regulators can effectively steer the industry towards lower-impact alternatives. The Nature Sustainability study (Ren et al., 2025) concluded bluntly that the AI server industry in the United States is “unlikely to meet its netzero aspirations by 2030 without substantial reliance on highly uncertain carbon offset and water restoration mechanisms.”⁸⁷

Carbon management for AI therefore needs to combine model-level efficiency with system-level measures such as power-aware scheduling, data-center heat management, renewable procurement, and regional grid decarbonisation. China’s data-center policy illustrates this system-level approach by linking computing-power expansion to explicit green targets. In 2024, Chinese authorities issued a green-development action plan for data centers that aims to reduce the average power usage effectiveness (PUE) of data centers below 1.5 by 2025, raise renewable-energy utilisation by 10 percent annually, and reach internationally advanced levels in PUE, energy efficiency, and carbon efficiency per unit of computing power by 2030.⁸⁸

The scale of this infrastructure underscores why embedding efficiency requirements into computing expansion matters: China’s total computing power reached 280 EFLOPS in 2024, and the country had built more than 4.25 million 5G base stations by the end of that year.⁸⁹

⁸⁶ De Vries, Patterns, 2025, *ibid.*; EU Energy Efficiency Directive recast, May 2024.

https://energy.ec.europa.eu/news/focus-data-centres-energy-hungry-challenge-2025-11-17_en

⁸⁷ Ren, S. et al. Nature Sustainability, November 2025. <https://www.nature.com/articles/s41893-025-01681-y>

⁸⁸ State Council of the People’s Republic of China. China sets green targets for data centers. 24 July 2024.

https://english.www.gov.cn/news/202407/24/content_WS66a0b167c6d0868f4e8e96ba.html

⁸⁹ National Data Administration / China Daily. Nation’s digital push gaining speed, *edge*. 19 May 2025.

https://www.nda.gov.cn/sjj/ywpd/sjzg/0519/20250519222423223703950_pc.html

2.3 Water Footprint of AI Models and Data Centers

The Scale of the Problem

Water is the least studied and least transparently reported dimension of AI's environmental footprint. Yet the numbers are significant. De Vries (Patterns, December 2025) estimated that AI systems' total water footprint could reach between 312.5 and 764.6 billion liters in 2025 — a range comparable to the global annual consumption of bottled water. The OECD has estimated that AI-related water consumption could reach 6.6 billion cubic meters by 2027. The Nature Sustainability study (Ren et al., 2025) projected an annual water footprint for US AI servers alone of 731–1,125 million cubic meters through 2030.⁹⁰

A scenario-based forecast published in the Journal of Cleaner Production (September 2025) modelled AI's water footprint under different growth and efficiency assumptions, confirming that even under optimistic efficiency scenarios, absolute water consumption continues to rise due to the sheer scale of infrastructure expansion.⁹¹

Two Types of Water Consumption

AI's water footprint has two distinct components. **Direct water use** occurs on-site in data centers, primarily for cooling. Evaporative cooling towers, adiabatic systems, and other cooling technologies consume water directly to maintain server temperatures within optimal ranges (typically 16–23°C). Cooling systems account for approximately 40% of total data center energy consumption, and

the water they require can be substantial: UNEP has noted that AI-related water use could be equivalent to half of the United Kingdom's total water consumption by 2027.⁹²

Indirect water use is embedded in electricity generation. Thermal and nuclear power plants consume significant quantities of water for cooling during electricity production. The water intensity of the grid varies enormously depending on the energy mix: De Vries found that the water intensity associated with US data center locations ranges from 0.68 to 11.98 liters per kilowatt-hour, a nearly 18-fold variation that is largely invisible to AI operators and users.⁹³

Geographic Concentration and Water Stress

Data centers are often concentrated in regions that already experiencing water stress. Internationally, rapid data center construction in India, the Middle East, and parts of China introduces AI water demand into regions where water allocation is already contested among agricultural, industrial, and domestic users.⁹⁴ Furthermore, for China, for instance, water-stress considerations are particularly important because computing hubs, renewable-energy resources and industrial clusters are unevenly distributed across regions. This pattern of spatial mismatch between digital infrastructure and hydrological capacity is especially evident in India, which serves as a critical case study of compounding water stress. India,

⁹⁰ De Vries, Patterns, 2025, *ibid.*; OECD estimate cited in UNFCCC TEC Technical Paper, 2025. <https://unfccc.int/tclear/tec/AI4climate.html>; Ren et al., Nature Sustainability, 2025, *ibid.*

⁹¹ Sustainable AI infrastructure: A scenario-based forecast of water footprint, J. Cleaner Production, September 2025. <https://www.sciencedirect.com/science/article/pii/S0959652625018785>

⁹² UNEP, September 2024, *ibid.*

⁹³ De Vries, Patterns, 2025, *ibid.*

⁹⁴ The ecology of AI, Carbon Neutral Systems, Springer, 2025, *ibid.*

home to 18% of the world's population, has access to only about 4% of the world's water resources.⁹⁵ According to NITI Aayog, approximately 600 million Indians already experience high to extreme water stress. Major technology and emerging AI/data centre hubs such as Chennai, Hyderabad, Bengaluru, and the Delhi NCR region are located in areas that face significant groundwater stress and urban water management challenges.⁹⁶ The 2019 Chennai water crisis further highlighted the vulnerability of rapidly urbanizing regions to acute water shortages under conditions of climate variability and infrastructure expansion.⁹⁷ These dynamics demonstrate how the scaling of AI infrastructure may intensify competition between industrial cooling needs, domestic water consumption, and agricultural demand in already climate-vulnerable regions.

This geographic concentration creates potential resource conflicts. Building water-intensive AI infrastructure in water-scarce regions is not merely an environmental risk but a social equity concern, particularly for communities that depend on the same water sources for agriculture and drinking water. The intersection of AI's water footprint with the environmental challenges addressed in Section 5 (desertification and land degradation) and Section 6 (water resources) of this paper is direct and consequential.

Innovation and Alternatives

Significant innovation in cooling technology is underway. Immersion cooling, in which servers are submerged in dielectric fluid, can achieve Power Usage Effectiveness (PUE) ratios as low as 1.05 — compared with 1.3–1.5 for conventional air-cooled

facilities — and can reduce cooling energy by up to 50%. **ByteDance** (TikTok) has deployed immersion cooling through its Volcano Engine platform. The Microsoft/WSP study published in *Nature* demonstrated that advanced cooling techniques (cold plates, one-phase and two-phase immersion) substantially reduce both water consumption and GHG emissions across the full data center lifecycle.⁹⁸

Closed-loop cooling systems, which recirculate water rather than evaporating it, and dry cooling systems, which use air instead of water, offer zero-water or near-zero-water alternatives, though often at the cost of higher energy consumption or reduced cooling capacity.

Data centers in Finland and Sweden demonstrate another approach: reusing waste heat from servers to heat up local buildings, turning a thermal liability into a community resource.⁹⁹

Measurement Gaps

Measurement of water footprint is even less mature than for carbon. Most technology companies report total water consumption in their sustainability reports but do not separate AI workloads from other operations, distinguish between water withdrawal and water consumption (water that is not returned to its source), or disaggregate direct cooling water from electricity-embedded water. Water Usage Effectiveness (WUE) is reported far less consistently than PUE. The UNEP procurement guidelines for data centers (June 2025) include WUE among recommended reporting metrics, but adoption remains voluntary in most jurisdictions.¹⁰⁰

⁹⁵ World Bank. India: Water security driving jobs, growth, and economic opportunity. Washington (DC): World Bank Group; 2026 Mar 19

⁹⁶ NITI Aayog. Composite Water Management Index. New Delhi: NITI Aayog, Government of India; 2019.

⁹⁷ World Resources Institute. How Does a Flood-prone City Run Out of Water? Inside Chennai's "Day Zero" Crisis. Washington (DC): World Resources Institute; 2019 Jun 25

⁹⁸ Alissa et al., *Nature*, 2025, *ibid.*

⁹⁹ <https://www.euroheat.org/dhc/knowledge-hub/datacentre-supplies-local-heating-in-maentsaelae-finland/>; <https://www.bloomberg.com/news/features/2025-05-14/finland-s-data-centers-are-heating-cities-too>

¹⁰⁰ UNEP, Sustainable Procurement Guidelines for Data Centres and Servers, June 2025.

<https://www.unep.org/technical-highlight/unep-releases-guidelines-curb-environmental-impact-data-centres>

2.4 Metrics and Tools for Sustainable Assessment

The Conceptual Foundation: Green AI

The intellectual framework for AI sustainability measurement was established in two seminal papers. Strubell, Ganesh, and McCallum (ACL, 2019) quantified for the first time the financial and environmental costs of training large NLP models, estimating that training a Transformer model with neural architecture search produced approximately 284 tons of CO₂ — roughly five times the lifetime emissions of an average American car. They proposed that energy costs should be reported alongside accuracy as a standard practice.¹⁰¹

Schwartz, Dodge, Smith, and Etzioni (“Green AI,” Communications of the ACM, 2020) formalized this proposal by introducing the distinction between “Red AI” — research that pursues accuracy through ever-larger compute budgets — and “Green AI” — research that treats efficiency as a primary evaluation metric. Their analysis of papers presented at major AI conferences in 2018–2019 found that the overwhelming majority reported only accuracy, with efficiency metrics largely absent. Five years later, this imbalance persists: the culture of AI research has not changed as fast as the measurement tools.¹⁰²

The Measurement Tool Ecosystem

A robust ecosystem of open-source and commercial tools now exists for tracking AI’s environmental impact during development and deployment. The most widely used include:

CodeCarbon (developed by Mila, BCG GAMMA, Haverford College, and Cohere), an open-source Python library that estimates CO₂ emissions from computing by tracking electricity consumption and applying regional carbon intensity factors. CodeCarbon was used by Hugging Face to monitor BLOOM’s real-time emissions over 18 days and has become the most cited tool in the research literature.¹⁰³

Eco2AI, developed jointly by Sber and the AI Research Institute (AIRI) in Russia, is another open-source Python library that monitors CO₂-equivalent emissions during code execution, with a specific focus on incorporating regional emission coefficients. Eco2AI has been validated on real-world applications, including the WaRP (Waste Recycling Plant) project — a collaboration between ITMO, Sber AI Lab, and AIRI that developed hierarchical waste detection neural networks with explicit sustainability tracking.¹⁰⁴

Eco4cast, also developed by AIRI with Sber’s support, goes beyond monitoring to provide predictive capabilities: it forecasts optimal time windows and geographic regions for computation with the lowest carbon intensity of the electricity grid. In certain scenarios, this temporal and spatial optimisation has reportedly achieved up to 90% reduction in CO₂ emissions compared with naïve scheduling. Eco4cast integrates a neural network that analyses emissions data and weather conditions across regions, and can be used both locally and in cloud environments.¹⁰⁵

Together, Eco2AI and Eco4cast illustrate that innovation in sustainable AI is not confined to

¹⁰¹ Strubell, E. et al. “Energy and Policy Considerations for Deep Learning in NLP,” ACL 2019. <https://aclanthology.org/P19-1355/>

¹⁰² Schwartz, R. et al. “Green AI,” Communications of the ACM, 2020. <https://arxiv.org/abs/1907.10597>

¹⁰³ Luccioni, A.S., JMLR, 2023, *ibid*.

¹⁰⁴ Budenny, S. et al. “Eco2AI,” 2022. <https://arxiv.org/abs/2208.00406>; <https://github.com/sb-ai-lab/Eco2AI>

¹⁰⁵ AIRI, Eco4cast. <https://github.com/AIRI-Institute/eco4cast>

Western institutions. Eco2AI's design — lightweight integration into existing Python workflows, a database of regional emission factors covering numerous countries, and validation on production research tasks such as the WaRP waste-sorting models — lowers the barrier to routine carbon accounting for researchers and developers. Eco4cast's predictive scheduling complements this by acting on the measurements: by shifting flexible workloads to greener time windows and regions, it converts emissions data into emissions reductions. For the multipolar, BRICS+ perspective that this paper advances, these tools are significant both as practical instruments and as evidence that the Global South and emerging economies are active contributors to — not merely consumers of — Green AI methods.¹⁰⁶

Additional tools include mlco2 (pre-training carbon estimation), LLMCarbon (carbon footprint prediction for LLM training), OpenCarbonEval (benchmarking platform), EnergyVis (interactive energy tracker by Basis Digital Energy), and Microsoft Sustainability Calculator (emissions forecasting for Azure).

Standardized Metrics and Frameworks

Several standardized metrics provide the building blocks for AI sustainability assessment. Power Usage Effectiveness (PUE), the ratio of total data center energy to IT equipment energy, is the most widely reported metric, with values ranging from approximately 1.05 for best-in-class immersion-cooled facilities to above 1.5 for older enterprise data centers. Water Usage Effectiveness (WUE) provides the equivalent metric for water, though it is reported far less consistently. Carbon Usage Effectiveness (CUE) measures direct carbon emissions per unit of IT energy.¹⁰⁷

The GHG Protocol, the most widely used international standard for carbon accounting, divides emissions into Scope 1 (direct on-site emissions), Scope 2 (indirect emissions from purchased electricity), and Scope 3 (all other indirect emissions across the value chain). For AI systems, Scope 3 is typically the most significant but least transparent category.

At the international level, a convergence of AI-specific standards is emerging. ITU-T L.1801 (February

2026) is the first international standard specifically designed for assessing the environmental impact of AI systems. ISO/IEC TR 20226:2025 incorporates sustainability across the full AI lifecycle. France's Frugal AI methodology, developed through AFNOR with input from over 100 organizations, provided the basis for Mistral AI's LCA.¹⁰⁸

Rather than relying on any single index, sustainable-AI assessment is best built around a **core basket of indicators**: the above-mentioned PUE, WUE, CUE, as well as renewable-electricity share, grid carbon intensity, server-utilisation rate, carbon intensity per unit of computing power, model energy per training run, inference energy per request, hardware-replacement cycle, and e-waste recovery rate. For AI applications in environmental governance, assessment should additionally capture avoided emissions, avoided water losses, monitoring accuracy, response time, cost reduction, and distributional effects. Crucially, such metrics become more effective when they are linked to planning and procurement — incorporated into project approval, government procurement, industrial standards, and public reporting — so that disclosure can gradually extend from data-center indicators to model- and service-level indicators for large models and high-volume AI services.

The Adoption Gap

Despite the availability of tools and the emergence of standards, actual adoption remains low. The ITU report *Measuring What Matters* (July 2025) found that most environmental measurement focuses on the training phase, while inference — the dominant source of operational emissions — is poorly and inconsistently measured. Supply chain (Scope 3) emissions are even less transparent. AI developers commonly rely on indirect estimates rather than direct measurement, and most corporate sustainability reports do not distinguish AI-specific impacts.¹⁰⁹

The OECD identified these measurement gaps as early as 2022, calling for a “data revolution” in AI environmental accounting. Three years later, the gaps largely persist. As the *Procedia CIRP* review (2025) observed, systematic LCA of AI is “still in its infancy” — technically feasible, but not yet standard practice.¹¹⁰

¹⁰⁶ Budenny, S. et al. “Eco2AI: carbon emissions tracking of machine learning models as the first step towards sustainable AI,” 2022. <https://arxiv.org/abs/2208.00406>; AIRI, Eco4cast. <https://github.com/AIRIInstitute/eco4cast>

¹⁰⁷ IEA, *Energy and AI*, 2025, *ibid.*

¹⁰⁸ ITU-T L.1801, February 2026, *ibid.*; ISO/IEC TR 20226:2025, *ibid.*; Mistral AI, 2025, *ibid.*

¹⁰⁹ ITU, “*Measuring What Matters*,” July 2025. <https://www.itu.int/hub/publication/s-gen-gda-001-2025/>

¹¹⁰ OECD, 2022, *ibid.*; *Procedia CIRP*, 2025, *ibid.*

Institutional Momentum

Several institutional developments are creating pressure for change. The United Nations Environment Assembly adopted Resolution 7/9 on the environmental sustainability of AI systems in December 2025, providing the highest-level multilateral mandate for standardized environmental assessment of AI.¹¹¹

The Coalition for Environmentally Sustainable AI, launched in February 2025 by France, UNEP, and ITU with over 100 partners, has prioritized standardized methods and comprehensive LCA frameworks. UNEP's Sustainable Procurement Guidelines for Data Centers and Servers (June

2025) provide the first international procurement framework with environmental performance criteria. The EU Energy Efficiency Directive's mandatory reporting requirements for large data centers, effective since May 2024, represent the first binding regulatory framework for data center environmental disclosure.¹¹²

These developments suggest that within the next two to three years, AI-specific environmental reporting will shift from voluntary to mandatory in at least some jurisdictions. The remaining barriers are not technical but political and institutional: harmonizing standards across jurisdictions, mandating disclosure, and creating incentives for adoption.

¹¹¹ UNEA Resolution 7/9, December 2025. <https://docs.un.org/r/UNEP/EA.7/Res.9>

¹¹² Coalition for Sustainable AI, February 2025, *ibid.*; UNEP Guidelines, June 2025, *ibid.*; EU EED, *ibid.*

3.1 Why Do We Care About Climate Change

Climate change is the most visible environmental domain in the AI debate because AI systems can both increase electricity demand and support mitigation and adaptation. The rapid growth of data centers makes AI part of the climate problem, while AI-enabled forecasting, grid optimisation, industrial process control and climate-risk modelling can make AI part of the solution. The policy challenge is to ensure that AI's enabling effects in decarbonisation exceed its direct and indirect emissions.

The International Framework

The United Nations Framework Convention on Climate Change (UNFCCC), adopted in 1992 and entered into force in 1994, provides the foundational legal framework for international climate action. Since the first Conference of the Parties (COP1) in Berlin in 1995, the UNFCCC process has produced two landmark agreements: the Kyoto Protocol (1997), which established binding emission reduction targets for developed countries, and the Paris Agreement (2015), which set the goal of holding global warming well below 2°C above pre-industrial levels and pursuing efforts to limit it to 1.5°C.¹¹³

The IPCC Sixth Assessment Synthesis Report (AR6, 2023) confirmed that human activities have unequivocally caused approximately 1.1°C of warming since pre-industrial times and that limiting warming to 1.5°C requires global CO₂ emissions to decline by approximately 43% by 2030 relative to 2019 levels and reach net zero by around 2050. Current trajectories fall far short of this pathway.¹¹⁴

The first Global Stocktake, completed at COP28 in Dubai (December 2023), assessed collective progress under the Paris Agreement and found significant gaps in both mitigation and adaptation. It called for transitioning away from fossil fuels in energy systems, tripling renewable energy capacity, and doubling rate of energy efficiency growth by 2030.

COP29 in Baku (November 2024) adopted the New Collective Quantified Goal (NCQG) on climate finance, setting a target of at least \$300 billion per year by 2035 for developing countries. It also finalized the rules for the Article 6.4 carbon market mechanism under the Paris Agreement, creating a framework for international carbon credit trading. However, COP29 produced no decisive language on fossil fuel phase-out, and developing countries expressed dissatisfaction with the finance goal as insufficient.¹¹⁵

COP30 in Belém, Brazil (November 2025) joined about 24,000 delegates for countries and 13,000 observers, 113 Parties submitted new or updated Nationally Determined Contributions (NDCs) covering nearly 80% of global emissions, with combined pledges projected to reduce emissions approximately 12% below 2019 levels by 2035 (with the total number of submitted NDCs now exceeding 150, representing near-universal global participation). COP30 adopted the Baku-to-Belém Roadmap to mobilize \$1.3 trillion annually for climate action in developing countries by 2035. It also tripled adaptation funding commitments and adopted operationalization of a Just Transition mechanism by COP31. However, the conference fell short on explicit fossil fuel phase-out language and forest protection measures, despite being held on the edge of the Amazon.¹¹⁶

¹¹³ UNFCCC. <https://unfccc.int/>; Paris Agreement, 2015. <https://unfccc.int/process-and-meetings/the-parisagreement>

¹¹⁴ IPCC AR6 Synthesis Report, 2023. <https://www.ipcc.ch/report/ar6/syr/>

¹¹⁵ UN Climate Change Conferences. <https://www.un.org/en/climatechange/un-climate-conferences>; WRI COP30 Outcomes. <https://www.wri.org/insights/cop30-outcomes-next-steps>

¹¹⁶ European Commission, "What did COP30 achieve?" December 2025. https://climate.ec.europa.eu/news/other-reads/news/what-did-cop30-achieve-2025-12-01_en; Earth.Org, "Did COP30 Succeed or Fail?" November 2025. <https://earth.org/did-cop30-succeed-or-fail/>

The AI–Climate Nexus

Artificial intelligence occupies a unique position within the climate challenge: it is both a contributor to and a potential mitigator of greenhouse gas emissions. As established in Section 2 of this paper, data centers consumed approximately 415 TWh of electricity in 2024 (around 1.5% of global demand), with AI as the primary growth driver. The IEA projects data center emissions could reach 300–500 Mt CO₂ by 2035.¹¹⁷

At the same time, the UNFCCC Technology Executive Committee has recognized AI as a potentially transformative tool for climate action. The #AI4ClimateAction Initiative, launched at COP28 in Dubai and operationalized through a workplan spanning 2024–2027, explores AI’s role as an enabler of climate solutions for mitigation and adaptation,

with particular attention to least developed countries (LDCs) and small island developing states (SIDS). At COP30 in Belém, the UNFCCC Technology Mechanism presented the AI for Climate Action Award, celebrating open-source AI-driven solutions from and for developing countries. The launch of the AI Climate Institute (AICI) was also announced at COP30, aimed at equipping people and institutions in developing countries with skills to harness AI for climate action.¹¹⁸

This dual character — AI as both problem and solution — demands that climate policy address both sides simultaneously. The remainder of this section examines AI’s contributions to climate action (Section 3.2) and the corporate and policy pathways for aligning AI development with climate goals (Section 3.3).

¹¹⁷ IEA, Energy and AI, April 2025. <https://www.iea.org/reports/energy-and-ai>

¹¹⁸ UNFCCC Technology Mechanism, AI for Climate Action Initiative. <https://unfccc.int/tclear/tec/AI4climate.html>; Climate Change AI @ COP30. <https://www.climatechange.ai/events/cop30>

3.2 How AI Is Contributing and Could Contribute to Address Climate Change

AI can support climate action in at least five ways: improving renewable-energy forecasting; optimizing grid dispatch and demand response; improving industrial energy efficiency; strengthening climate-risk modelling and disaster early warning; and improving carbon monitoring, reporting and verification. These applications are especially relevant in large emerging economies where energy demand is still growing and where decarbonisation must be compatible with industrial upgrading.

Climate Modelling and Weather Prediction

AI could transform the speed and accuracy of weather forecasting and climate modelling. Google DeepMind’s **GraphCast** model (2023) demonstrated that a machine learning-based approach can produce 10-day global weather forecasts more accurately than the leading physics-based system (ECMWF’s HRES) at a fraction of the computational cost — generating a forecast in under one minute on a single TPU, compared with hours on a supercomputer with thousands of processors.¹¹⁹

Microsoft’s **Aurora** foundation model for Earth system prediction, trained on over a million hours of diverse weather and climate simulations, has demonstrated the ability to help in forecasting of atmospheric chemistry, ocean dynamics, and extreme weather events at unprecedented resolution.¹²⁰ Huawei’s

Pangu-Weather has achieved similar progress in medium-range forecasting. These models represent a paradigm shift from purely physics-based climate simulation to hybrid AI-physics approaches that can explore climate scenarios orders of magnitude faster.¹²¹

In Russia, **Yandex** researchers, together with scientists from the Higher School of Economics, developed an AI-based model to predict El Niño, one of the most influential climate phenomena affecting global weather patterns. The model combines deep learning with physics-based climate simulations, using both historical ocean temperature observations and large volumes of synthetic climate data to identify long-term patterns in the Pacific Ocean. Trained on Yandex Cloud infrastructure, the system can forecast El Niño events up to 18 months in advance, helping improve preparedness for floods, droughts, wildfires, and other climate-related risks. This approach demonstrates how AI can enhance traditional climate modelling by extracting complex relationships from vast environmental datasets and extending the forecasting horizon for critical weather phenomena.¹²²

In India — among the countries most vulnerable to climate-related risks, including heatwaves, floods, cyclones, droughts, glacial melt, and sea-level rise — the India Meteorological Department (IMD) has adopted AI and machine-learning systems for monsoon and rainfall prediction, supporting

¹¹⁹ DeepMind, “GraphCast: AI model for faster and more accurate global weather forecasting.” <https://deepmind.google/discover/blog/graphcast-ai-model-for-weather-prediction/>

¹²⁰ Microsoft. Aurora: A Foundation Model for the Earth System. <https://microsoft.github.io/aurora/intro.html>

¹²¹ Huawei Cloud. Pangu Weather Model. <https://www.huaweicloud.com/intl/en-us/about/takecloudleap2024/ai-weather-prediction.html>

¹²² Yandex Cloud. What does El Niño lead to and how Yandex Cloud tools help predict it. April, 2023. <https://yandex.cloud/en/blog/posts/2023/04/el-nino>

disaster preparedness.¹²³ These advances also support agricultural resilience and adaptation planning in highly climate-sensitive regions where a large share of the population depends on agriculture and fisheries.

Beyond greenhouse-gas mitigation, AI is also increasingly applied to broader environmental health challenges such as air-quality monitoring. India experiences one of the world's highest disease burdens attributable to **air pollution**,¹²⁴ with many cities recording annual PM2.5 concentrations substantially above WHO guideline values.¹²⁵ AI and machine-learning systems are being used for pollution forecasting, source apportionment, satellite-based exposure assessment, low-cost sensor calibration, and urban emissions modelling — capabilities that are particularly valuable where conventional monitoring is limited and exposure patterns are highly dynamic due to traffic, industry, seasonal biomass burning, and meteorology.¹²⁶

Emissions Monitoring and Tracking

Climate TRACE (Tracking Real-Time Atmospheric Carbon Emissions), a coalition of organizations using AI and satellite data, provides independent, facility-level greenhouse gas emissions data for virtually every country and major emitting sector. As of 2025, Climate TRACE monitors over 350 million assets globally and has revealed that actual emissions from many sectors are significantly higher than those reported in national inventories. This independent verification capacity is particularly valuable for the Enhanced Transparency Framework under the Paris Agreement, which requires all Parties to submit biennial transparency reports from 2024 onwards.¹²⁷

UNEP's International Methane Emissions Observatory (IMEO), launched in 2021, uses AI-powered analysis of satellite data to detect methane emission hotspots globally and issue early warning signals to affected governments. Methane — a greenhouse gas approximately 80 times more potent

than CO₂ over a 20-year period (in the UNFCCC 100-year period is used and another coefficient for methane should be used) — is a priority target under the Global Methane Pledge, endorsed by over 150 countries.¹²⁸

Energy System Optimisation

AI is enabling more efficient integration of variable renewable energy sources into power grids.

Google DeepMind's AI system for data center cooling reduced energy costs by 40%, a finding that has been applied across Google's global data center fleet and has influenced industry-wide practices.¹²⁹

Siemens Gamesa Renewable Energy developed an AI system that analyses real-time data from wind turbines — including meteorological conditions and equipment performance — to predict wind generation, reducing maintenance costs and improving energy reliability.

In Russia, Sber's **Andromeda** solution uses a hybrid architecture that optimises task distribution based on energy consumption, scheduling computations during periods when lower-carbon energy is available.

AI-driven smart grids can balance supply and demand in real time, forecast consumption patterns, optimize storage, and reduce curtailment of renewable generation. A study published in *Nature Climate Change* (Kaack et al., 2022) provided a comprehensive framework for understanding how AI can be aligned with climate mitigation, identifying applications across electricity, transportation, buildings, industry, and land use. The authors also cautioned about “knock-on emissions” — indirect effects where AI-enabled efficiencies may increase overall consumption through rebound effects.¹³⁰

At the national level, AI–energy integration is increasingly framed as infrastructure and energy security strategy rather than only enterprise-

¹²³ Government of India, Press Information Bureau. New Delhi: Press Information Bureau. Lawrence Berkeley National Laboratory. Data Centre Water Use and Sustainability Reports. Berkeley (CA): Lawrence Berkeley National Laboratory; various years.

¹²⁴ Health Effects Institute. State of Global Air 2024. Boston (MA): Health Effects Institute; 2024.

¹²⁵ World Health Organization. WHO Global Air Quality Database. Geneva: World Health Organization; 2024.

¹²⁶ Council of Scientific and Industrial Research. Air Quality Forecasting and Monitoring Research. New Delhi: Council of Scientific and Industrial Research; various years.

¹²⁷ Climate TRACE. <https://climatetrace.org/>

¹²⁸ UNEP IMEO. <https://www.unep.org/topics/energy/methane/international-methane-emissions-observatoryimeo>

¹²⁹ Evans, R. “DeepMind AI Reduces Google Data Centre Cooling Bill by 40%.” 2016. <https://deepmind.google/discover/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/>

¹³⁰ Kaack, L.H. et al. “Aligning Artificial Intelligence with Climate Change Mitigation,” *Nature Climate Change*, 2022. <https://www.nature.com/articles/s41558-022-01377-7>

level efficiency. In September 2025, Chinese authorities issued a plan to accelerate the integration of AI with the energy sector, aiming to establish an innovation system for **AI–energy integration** by 2027 and to make China’s AI applications in the energy sector world-leading overall by 2030. The plan calls for wider use of AI in power grids, renewable energy, and nuclear power, and emphasizes coordination between computing power and electricity supply — with climate relevance lying in improved renewable integration, reduced curtailment, optimised dispatch, and more flexible power systems.¹³¹

Wildfire Prediction and Disaster Response

AI is increasingly deployed for predicting and managing wildfires, which are growing in frequency and intensity also due to climate change. AI models analyze historical data on weather conditions, topography, vegetation, and human activity to create fire probability maps and predict fire spread in real time.

In Russia, MTS deployed a smart video monitoring system in Primorye using AI and over 30 cameras to detect smoke and fire at distances up to 20 kilometers. Ufa State Petroleum Technical University (UGNTU) developed an AI program for

identifying dead standing trees from drone-based remote sensing data, with accuracy exceeding 90%. East Siberian State University (VSGUTU) created an AI system for forest fire prediction using satellite data with approximately 85% accuracy.

The Novosibirsk Center for Digital Transformation, together with Innogeotekh, launched a pilot monitoring project that detected 1,100 cases of illegal deforestation across 8,600 hectares using AI analysis of satellite imagery (2017–2024). The Yugra Forest Control project uses neural networks to detect illegal logging with 92% accuracy, reducing violations by half.

Carbon Capture and Materials Discovery

AI is accelerating the discovery of materials for carbon capture and low-carbon manufacturing.

AI-driven molecular simulation can screen millions of candidate compounds for CO₂ absorption capacity, dramatically reducing the time and cost of laboratory experimentation. AI has also been applied to optimizing concrete formulations that reduce the carbon intensity of construction — a sector responsible for approximately 8% of global CO₂ emissions.

¹³¹ State Council of the People’s Republic of China, “China unveils plan on AI-energy integration to drive green transition”, 8 September 2025. https://english.www.gov.cn/news/202509/08/content_WS68be8c3ec6d0868f4e8f566d.html

3.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments

Corporate Climate Targets in the AI Sector

Major technology companies have adopted ambitious climate targets, though progress has been uneven and, in some cases, contradicted by rising AI-related emissions.

Microsoft committed to becoming carbon-negative by 2030 and plans to offset its entire historical carbon footprint by 2050. The company is investing in renewable energy for all data centers, forest restoration, soil carbon sequestration, and has purchased nuclear power: in September 2024, Microsoft announced the restart of the Three Mile Island nuclear plant (837 MW) to power its data centers. However, Microsoft's own sustainability report showed that its carbon emissions increased by 29% between 2020 and 2024, primarily driven by AI infrastructure expansion.

Google has committed to operating on 24/7 carbon-free energy by 2030. In October 2024, Google signed an agreement with Kairos Power for six to seven small modular nuclear reactors (500 MW total), with the first expected online by 2030. Google's 2024 Environmental Report showed that its emissions rose 48% compared to 2019, largely due to data center growth. The company's DeepMind division has been a pioneer in applying AI to energy efficiency, but the scale of infrastructure growth has outpaced efficiency gains.

Amazon (AWS) pledged to reach net-zero emissions by 2040 under The Climate Pledge. In March 2024, Amazon purchased a \$650 million data center campus adjacent to the Susquehanna nuclear plant,

securing up to 960 MW of nuclear power. AWS has also invested in advanced cooling technologies, including liquid cooling for high-density AI chips, and developed custom energy-efficient processors (Graviton 3, Trainium, Inferentia).

Alibaba Cloud reported that its cloud migration services reduced customer CO₂ emissions by 6.863 million tons in 2023, according to a Carbon Trust assessment. ByteDance deployed immersion cooling through its Volcano Engine platform, reducing cooling energy by 50%.

While **OpenAI** has not publicly announced company-wide climate targets, its CEO, Sam Altman, has invested in Helion Energy, a nuclear fusion startup. The convergence of AI's massive, baseload power requirements and nuclear energy's carbon-free, weather-independent generation makes this an economically logical but socially and politically complex choice. The IEA's Energy and AI report (2025) projects that nuclear power will play an increasingly important role in meeting data center electricity demand towards the end of this decade and beyond.¹³²

Frameworks and Accountability Mechanisms

Several international frameworks provide accountability structures for corporate climate commitments. The **Science Based Targets initiative (SBTi)** validates corporate emission reduction targets against the Paris Agreement's 1.5°C pathway. As of 2025, over 7,000 companies have committed to SBTi targets. However, SBTi has

¹³² IEA, Energy and AI, 2025. <https://www.iea.org/reports/energy-and-ai>

faced criticism for its 2024 decision to allow the use of carbon offsets for Scope 3 emissions, which some argue weakens the standard.¹³³

RE100, led by the Climate Group, brings together over 400 companies committed to sourcing 100% renewable electricity. Several major AI companies are members, though the 24/7 carbonfree energy standard championed by Google goes beyond annual renewable energy matching to hour-by-hour accounting.¹³⁴

The UNFCCC's **Race to Zero** campaign, launched ahead of COP26, provides a framework for non-state actors — including companies, cities, regions, and investors — to commit to net-zero emissions by 2050 at the latest. The campaign has enrolled over 11,000 non-state actors — a modest fraction of the global business community, yet a meaningful pioneer cohort accumulating the methodologies, tools, and institutional experience that broader adoption will eventually depend on — though monitoring and enforcement of commitments remain challenges.¹³⁵

The **Task Force on Climate-related Financial Disclosures (TCFD)**, now integrated into the ISSB sustainability disclosure standards, requires companies to report on climate-related risks and opportunities. The EU's Corporate Sustainability Reporting Directive (CSRD), effective from 2024, introduces mandatory sustainability reporting for large EU companies, including climate metrics. These frameworks create market pressure for AI companies to disclose and manage their environmental impact.

Corporate commitments are most credible when they move beyond general carbon-neutrality pledges towards **measurable operational pathways**: publishing energy and water metrics, increasing renewable-electricity procurement, improving server utilisation, adopting lowcarbon cooling, designing smaller task-specific models, and reporting emissions from both training and inference. Where AI is deployed to cut emissions in sectors such as power, steel, chemicals, logistics, manufacturing, and buildings, those benefits are credible only if application-level savings are measured against

the additional computing demand required to produce them. This paper therefore recommends that companies disclose both AI's enabling benefits and AI's own footprint within a unified accounting framework.

Instruments from BRICS and the Global South

In Russia, the Sakhalin carbon neutrality experiment, with over 32 million carbon units in circulation as of December 2024, is developing the national carbon trading infrastructure. Sber offers green certificates confirming renewable energy use. AIRI's Eco4cast library enables temporal and geographic optimisation of AI computations for lower carbon intensity, reportedly achieving up to 90% CO₂ reduction in some scenarios. Rosatom is building wind farms, and Hevel is developing solar energy projects.¹³⁶

The BRICS+ AI Alliance, launched in December 2024, and the BRICS framework program for AI cooperation (adopted at the 2024 Kazan Summit) include climate forecasting among priority areas for joint AI research. The New Development Bank (NDB) launched a \$5 billion digital sovereignty fund in 2025 for AI infrastructure in BRICS countries.¹³⁷

The **Baku-to-Belém Roadmap**, presented by the COP29 and COP30 presidencies, recognises that “technological transformations — including advances in artificial intelligence, quantum computing, and space technologies — are expected to shape how societies function and interact” in the context of mobilizing \$1.3 trillion annually for climate action by 2035. This acknowledgement signals growing recognition within the UNFCCC process that AI governance and climate governance are converging.¹³⁸

The net effect of AI on climate change — whether its mitigation contributions ultimately outweigh its emissions — is not a technical question but a governance question. It depends on the choices made by policymakers, technology companies, and civil society in the next critical decades.

¹³³ SBTi. <https://sciencebasedtargets.org/>

¹³⁴ RE100. <https://www.there100.org/>

¹³⁵ UNFCCC Race to Zero. <https://unfccc.int/climate-action/race-to-zero-campaign>

¹³⁶ AIRI, Eco4cast. <https://github.com/AIRI-Institute/eco4cast>

¹³⁷ BRICS Council, “Joint BRICS Projects in AI.” <https://bricscouncil.ru/en/analytics/sovместnye-proekty-stranbriks-v-oblasti-iskusstvennogo-intellekta>

¹³⁸ UNFCCC, Baku to Belém Roadmap to 1.3T. https://unfccc.int/sites/default/files/resource/Relatorio_Roadmap_COP29_COP30_EN_final.pdf

Section 4.

Biological Diversity

4.1 Why Do We Care About Biological Diversity	44
4.2 How AI Is Contributing and Could Contribute to Address Conservation and Sustainable Use of Biological Diversity	46
4.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments	48



4.1 Why Do We Care About Biological Diversity

Biodiversity loss weakens ecosystem resilience, food security, public health and the stability of natural carbon sinks. Compared with climate change, biodiversity is more local and heterogeneous, which creates major monitoring challenges. Traditional field surveys are often costly, periodic and labor-intensive. AI can help by turning remote sensing, camera traps, acoustic devices and ecological databases into continuous monitoring systems.

The International Framework

The Convention on Biological Diversity (CBD), adopted at the 1992 Earth Summit in Rio de Janeiro and entered into force in December 1993, is the principal international legal instrument for the conservation of biological diversity, the sustainable use of its components, and the fair and equitable sharing of benefits arising from genetic resources. With 196 Parties, the CBD has near-universal participation.¹³⁹

The Kunming-Montreal Global Biodiversity Framework (GBF), adopted at COP15 in December 2022, sets out 4 goals for 2050 and 23 targets for 2030, establishing the most ambitious roadmap to date for halting and reversing biodiversity loss. Its headline goal is to protect at least 30% of the world's land and ocean areas by 2030 (the “30x30” target). The GBF builds on the lessons of its predecessor, the Aichi Biodiversity Targets (2011–2020), none of which were fully achieved.¹⁴⁰

COP16 of the CBD, held in Cali, Colombia, in October–November 2024, produced several landmark outcomes. Most significantly for the AI–biodiversity nexus, COP16 established the **Cali Fund** — a multilateral mechanism for the fair and equitable sharing of benefits from the use of Digital Sequence Information (DSI) on genetic resources. The Cali Fund requires companies that commercially benefit from genomic data — including those in the pharmaceutical, agricultural biotechnology, and AI sectors — to contribute financially to biodiversity conservation. At least 50% of the Cali Fund's resources are allocated to indigenous peoples and local communities, recognizing their role as custodians of biodiversity. The Cali Fund was formally launched in Rome in February 2025 at the resumed session of COP16.¹⁴¹

COP17 of the CBD, scheduled for 2026, will review progress on implementing the GBF targets and the Cali Fund mechanism, with a particular focus on integrating technological solutions — including AI — into national biodiversity strategies and action plans (NBSAPs).¹⁴²

¹³⁹ CBD. <https://www.cbd.int/convention/>

¹⁴⁰ CBD, Kunming-Montreal Global Biodiversity Framework. <https://www.cbd.int/gbf/>

¹⁴¹ UNEP, “How a groundbreaking agreement could raise billions to protect the web of life,” November 2024. <https://www.unep.org/news-and-stories/story/how-groundbreaking-agreement-could-raise-billions-protect-web-life>;

UN Sustainable Development, “The Cali Fund launches,” February 2025. <https://www.un.org/sustainabledevelopment/?p=115133>

¹⁴² Blue Marine Foundation, COP17 overview.

<https://www.bluemarinefoundation.com/our-events/cop17-convention-on-biological-diversity-2026/>

The Scale of the Crisis

The biodiversity crisis is accelerating. Approximately one million species are threatened with extinction, many within decades. The global rate of species extinction is at least tens to hundreds of times higher than the average rate over the past 10 million years. Currently, less than 7% of ecosystems are surveyed at high resolution, requiring enhanced data-sharing initiatives and new technologies to close the monitoring gap.¹⁴³

This creates both an urgent demand for and a significant opportunity for AI: the sheer scale and complexity of biodiversity monitoring — spanning millions of species across terrestrial, marine, and freshwater ecosystems — exceeds the capacity of traditional observation methods. AI offers the ability to process satellite imagery, acoustic recordings, camera trap images, and environmental DNA samples at scales that would be impossible for human analysts alone. As with the previous sections, the effectiveness of these applications depends not only on algorithmic performance but also on the quality of ecological monitoring systems, expert validation, and transparent environmental governance.

¹⁴³ IPBES Global Assessment Report, 2019. <https://www.ipbes.net/global-assessment>

4.2 How AI Is Contributing and Could Contribute to Address Conservation and Sustainable Use of Biological Diversity

AI can contribute to biodiversity governance through species recognition, habitat mapping, invasive species detection, wildlife movement analysis, illegal-activity detection and ecological-risk prediction. The main advantage is not simply automation; it is the ability to convert fragmented observation data into timely ecological signals for conservation decisions.

Monitoring and Species Identification

AI is transforming biodiversity monitoring through the analysis of vast datasets from multiple sensor types. **Satellite and drone imagery** analyzed by deep learning models can detect changes in habitat extent, forest cover, and land use at high spatial and temporal resolution.

Camera traps equipped with AI-powered image recognition can identify individual animals, track population dynamics, and detect invasive species — processing millions of images that would otherwise require years of manual analysis.

Bioacoustics — the use of AI to analyze sound recordings from ecosystems — enables detection and monitoring of species that are difficult to observe visually, including birds, bats, marine mammals, and insects. AI models can identify species from their vocalizations, detect changes in ecosystem

health through soundscape analysis, and monitor compliance with protected area regulations.

Environmental DNA (eDNA) — genetic material shed by organisms into their environment — can be collected from water, soil, or air samples and analyzed using AI to detect the presence of species without direct observation. This technique is particularly valuable for monitoring aquatic biodiversity and detecting rare or elusive species.

China offers a concrete example of how AI shifts biodiversity monitoring from periodic surveys towards **continuous observation**. Drawing on a Ministry of Ecology and Environment briefing, China has applied bird image recognition, acoustic recognition, and plant-species image recognition for year-round monitoring. In Jiangsu Province, 100 bird-acoustic recognition devices installed around 20 biodiversity observation stations had collected over 400,000 data entries covering more than 240 species since 2023, while substantially reducing labour costs. The case illustrates that AI does not replace ecological expertise but expands the frequency, scale, and comparability of observation — and that biodiversity applications depend on field devices, data standards, expert validation, and long-term public-sector support.¹⁴⁴ Independent verification and publicly available monitoring data remain essential to ensure that AI-supported biodiversity assessments are scientifically robust and reproducible.

¹⁴⁴ China Daily, "AI transforms China's conservation efforts", 31 March 2026.
<https://www.chinadaily.com.cn/a/202603/31/WS69cb2b88a310d6866eb40dda.html>

India, which contains approximately 7–8% of all recorded species on only about 2.4% of global land area¹⁴⁵ and includes four major biodiversity hotspots¹⁴⁶, is deploying AI to protect wildlife from poaching through camera- and drone-based monitoring¹⁴⁷, with applications also emerging for human–wildlife conflict management using geospatial analytics and movement prediction systems.

The Biodiversity Policy Analyzer

AI-driven policy tools are helping governments align their national strategies with global targets. The **Biodiversity Policy Analyzer**, developed in collaboration with CGIAR and other partners, assists governments in mapping their National Biodiversity Strategies and Action Plans (NBSAPs) against the 23 targets of the Kunming-Montreal GBF, identifying gaps and opportunities for alignment.¹⁴⁸

Scientific Assessment and Attribution

AI contributes to scientific understanding of biodiversity change by enabling researchers to detect trends across large datasets and attribute biodiversity changes to specific drivers — distinguishing, for example, between the effects of climate change, land-use change, pollution, and invasive species. Machine learning models can integrate data from multiple sources (field surveys, remote sensing, citizen science platforms) to produce more accurate and timely assessments than traditional methods.

AI for Biodiversity Finance

AI is emerging as a tool for mobilizing and directing biodiversity finance. The **LSE Grantham Research Institute** has analyzed AI's role in financing biodiversity conservation, highlighting applications in risk assessment, natural capital accounting, and monitoring compliance with biodiversity-related financial disclosures. As the Taskforce on Nature-related Financial Disclosures (TNFD) framework gains adoption, AI tools that can assess nature-related risks and dependencies across corporate supply chains will become increasingly valuable.¹⁴⁹

GPAI and UNDP Frameworks

The Global Partnership on AI (GPAI) published “Biodiversity and AI: Opportunities & Recommendations for Action” (November 2022), identifying priority areas for AI deployment in biodiversity conservation and providing recommendations for responsible AI use in ecological contexts.¹⁵⁰

UNDP's “People-Centric AI for Conserving Biodiversity” report (December 2025) emphasizes that AI systems should respect ethical standards, including the rights of indigenous peoples and local communities to control their traditional knowledge and biodiversity data. The report advocates for AI approaches that are inclusive, participatory, and designed to empower local actors rather than replace them.¹⁵¹

¹⁴⁵ International Union for Conservation of Nature (IUCN). India. Gland (Switzerland): IUCN. Wildlife Institute of India. AI-enabled Wildlife Monitoring.

¹⁴⁶ Government of India, Ministry of Environment and Forests. India's Fourth National Report to the Convention on Biological Diversity. New Delhi: Ministry of Environment and Forests; 2009.

¹⁴⁷ Paul H. “How India is using AI to stop poaching: 5 tech initiatives protecting wildlife.” The Better India. 2025 Jun 30.

¹⁴⁸ CGIAR, “Leveraging human-centred AI for biodiversity,” 2024. <https://events.cgiar.org/leveraginghumancenteredartific>

¹⁴⁹ LSE Grantham Institute, “What is AI's role in financing biodiversity conservation?”

¹⁵⁰ GPAI, Biodiversity and AI. <https://wp.oecd.ai/app/uploads/2025/05/biodiversity-and-AI-opportunitiesrecommendations-for-action.pdf>

¹⁵¹ UNDP, “People-centric AI for conserving biodiversity,” December 2025. <https://www.undp.org/sites/g/files/zskgke326/files/2025-12/people-centric-ai-for-conserving-biodiversity.pdf>

4.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments

Private companies can support biodiversity-related AI by developing low-power edge devices, open species-recognition datasets, privacy-preserving ecological data platforms and transparent model-validation procedures. Public agencies and research institutions should define ecological indicators and quality-control standards. A specific caution applies to biodiversity data: sensitive species-location information may require restricted access to prevent poaching or habitat disturbance, and should not be treated as a generic AI resource.

Ethical Principles for AI in Biodiversity

The International Research Center for AI Ethics and Governance at Tsinghua University has developed principles for AI in biodiversity conservation that advocate an ecocentric approach — recognizing that non-human life has intrinsic value and that AI systems should be designed to promote ecological prosperity, not merely to optimize human utility. These principles include: minimising AI's own energy consumption and environmental footprint; avoiding the destruction of habitats for AI infrastructure; ensuring transparency and explainability of algorithms used in conservation decisions; and complying with international environmental law, including the CBD.¹⁵²

Data Sovereignty and Indigenous Rights

A critical issue at the intersection of AI and biodiversity is data sovereignty — the right of indigenous peoples and local communities to control how their traditional knowledge and biodiversity data is collected, stored, analyzed, and used. The CBD's provisions on traditional knowledge (Article 8(j)) and the Cali Fund's allocation of at least 50% of resources to indigenous peoples and local communities reflect growing recognition that equitable AI governance must include these actors as rights-holders, not merely as data sources.

The UNDP report (2025) emphasizes that AI systems for conservation should respect indigenous data sovereignty and ensure that communities retain authority over the use and distribution of their knowledge. This principle has implications for how training datasets are assembled, how AI models are validated, and how the outputs of AI-driven conservation decisions are governed. This theme connects directly to Section 8 (Stakeholder Roles), where the role of civil society — including women, indigenous peoples, and local communities — is analyzed in detail.

¹⁵² International Research Center for AI Ethics and Governance, Tsinghua University.

<https://aiig.tsinghua.edu.cn/en/info/1026/1193.htm>;

AI Ethics and Governance Institute.

<https://ai-ethics-andgovernance.institute/2022/09/22/principles-on-artificial-intelligence-for-biodiversity-conservation/>

Capacity Building and the 2030 Timeline

The GBF's 23 targets for 2030 create an urgent timeline for the deployment of AI tools in biodiversity conservation. However, significant barriers remain: insufficient high-resolution ecological data, limited technical capacity in many developing countries, and the high cost of deploying and maintaining AI systems in remote or under-resourced environments. Closing these gaps requires targeted investment in capacity building, technology transfer, and South-

South cooperation — themes addressed in the UNFCCC, CBD, and UNCCD contexts and in Section 8 of this paper.

The AI for Good Discovery series, in partnership with the CBD, explores how AI can support the conservation and sustainable use of biodiversity, providing a platform for practitioners, researchers, and policymakers to share experiences and develop best practices.¹⁵³

¹⁵³ ITU AI for Good / CBD partnership.

<https://aiforgood.itu.int/new-ai-for-biodiversity-series-how-can-we-useai-to-monitor-biodiversity-and-support-conservation-actions/>

5.1 Why Do We Care About Desertification and Land Degradation

Desertification and land degradation affect food security, rural livelihoods, dust storms, biodiversity and carbon sinks. Their causes are complex: climate variability, water stress, overgrazing, unsustainable farming, urban expansion and industrial activity. AI is useful in this area because land degradation is spatially distributed and can be monitored through satellite imagery, meteorological data, soil data and field observations.

The International Framework

The United Nations Convention to Combat Desertification (UNCCD), adopted in 1994 and entered into force in 1996, is the sole legally binding international agreement linking environment and development to sustainable land management. With 197 Parties (196 countries and the European Union), the UNCCD addresses land degradation, drought, and desertification — processes that fuel poverty, conflict, food insecurity, and biodiversity loss.¹⁵⁴

The Scale of the Challenge

Up to 40% of the world's land is already classified as degraded, directly affecting over 3.2 billion people and causing estimated economic losses of \$878 billion annually. Three-quarters of Earth's land became permanently drier over the past three decades. Drylands now make up 40.6% of all land on Earth (excluding Antarctica), and the number of people living in drylands has doubled from 1.2 billion in 1990 to 2.3 billion in 2020.¹⁵⁵

Unless 1.5 billion hectares of land are restored by 2030, it may not be possible to achieve a land-degradation-neutral world. At least \$2.6 trillion in total investment is needed by 2030 — equivalent to approximately \$1 billion per day — to restore over one billion hectares of degraded land and build drought resilience.¹⁵⁶

National assessments illustrate the scale in major economies. According to the **Indian Space Research Organisation**, approximately 96 million hectares — nearly 30% of India's geographical area — are affected by land degradation or desertification.¹⁵⁷ AI and remote sensing technologies are increasingly applied for vegetation stress monitoring, drought assessment, land-use change analysis, crop monitoring, and soil degradation mapping. These tools are especially important in dryland regions where climate variability directly affects agricultural productivity, water availability, food security, and rural livelihoods.

UNCCD COP16, held in Riyadh, Saudi Arabia, in December 2024 — the first UNCCD COP in the Middle East and North Africa region — brought together 197 Parties to negotiate a global drought regime and strengthen commitments to land restoration. The conference produced major financial pledges and advanced negotiations on drought resilience, though a fully agreed global drought protocol was deferred to COP17 in Mongolia in 2026.

¹⁵⁴ UNCCD. <https://www.unccd.int/convention/about-convention>

¹⁵⁵ UNCCD, Global Land Outlook 2. <https://www.unccd.int/resources/global-land-outlook/glo2>; UNCCD Desertification and Drought Day 2024. <https://www.unccd.int/events/desertification-drought-day/2024>; UNCCD press release, “Three-quarters of Earth's land became permanently drier,” December 2024. <https://www.unccd.int/news-stories/press-releases/three-quarters-earths-land-became-permanently-drier-lastthree-decades>

¹⁵⁶ UNCCD Financial Needs Assessment, COP16, December 2024. <https://www.unccd.int/sites/default/files/2024-12/FNA%20press%20release%20EN.pdf>

¹⁵⁷ Indian Space Research Organisation, Space Applications Centre, Desertification and Land Degradation Atlas of India, 2021. ISBN 978-93-82760-39-9

5.2 How AI Is Contributing and Could Contribute to Address Desertification and Land Degradation: Challenges and Mitigation Issues

AI can support land-degradation governance by classifying land-cover change, detecting vegetation stress, forecasting drought risk, optimizing irrigation and grazing plans, and evaluating ecological restoration outcomes. The strongest applications combine remote sensing with local knowledge and policy instruments. AI can identify risk patterns, but restoration decisions still depend on land tenure, water rights, local livelihoods and ecological suitability.

The International Drought Resilience Observatory (IDRO)

The centerpiece AI initiative in this domain is the **International Drought Resilience Observatory (IDRO)** — the first global AI-driven platform for proactive drought management, developed under the International Drought Resilience Alliance (IDRA) and the UNCCD. The IDRO prototype was unveiled at UNCCD COP16 in Riyadh in December 2024.

It harnesses remote-sensing data and artificial intelligence to analyze and visualize key social and environmental drought resilience indicators, providing actionable information for authorities, land managers, and water managers.¹⁵⁸

UNCCD Executive Secretary Ibrahim Thiaw stated at COP16: “We can only achieve meaningful results by joining forces; IDRO is a perfect example of the collaborative mindset we require to protect people from impending and future droughts.” The full version of IDRO, with all functionalities, is expected to launch at UNCCD COP17 in Mongolia in 2026, marking a shift towards proactive drought management worldwide.¹⁵⁹

Land Degradation Monitoring

AI tools combining machine learning with satellite imagery analysis are used to monitor land degradation in real time, enabling identification of degradation trends, tracking reforestation and restoration efforts, and supporting sustainable land management decisions. The UNCCD promotes the use of open, public datasets — including aerial imagery from Landsat, Sentinel, and other satellite programs — for training AI models to assess aridity trends and their impacts.

In Saudi Arabia, a large-scale counter-desertification program launched in August 2023 uses AI algorithms to analyze satellite images, monitoring changes in land use, vegetation cover, and soil moisture levels. The deployment of AI is helping detect and assess desertification trends in

¹⁵⁸ UNCCD, “What you need to know about the first AI-driven tool for drought resilience.” ; UNCCD, “Global response to drought takes center stage at UN land conference in Riyadh,” December 2024.

<https://www.unccd.int/news-stories/press-releases/global-response-drought-takes-center-stage-un-landconference-riyadh>

¹⁵⁹ UNCCD/IDRA. <https://idralliance.global/solutions/what-you-need-know-about-first-ai-driven-tool-droughtresilience>

vulnerable areas and supporting decisions about intervention priorities.¹⁶⁰

China's long-term desertification-control experience provides a substantial non-Western empirical base. In 2024, China completed a green belt around the Taklamakan Desert as part of the decades-long Three-North Shelterbelt program; reporting noted that more than 30 million hectares of trees had been planted under the program and that China's forest coverage had risen above 25 percent by the end of 2023.¹⁶¹ For AI governance, the principal significance lies not in the tree-planting programme itself, but in how AI and remote sensing support long-term ecological restoration by selecting suitable species, monitoring vegetation survival, detecting illegal land-use change, forecasting dust-storm risks, and evaluating restoration outcomes over time.

Sand and Dust Storm Monitoring

The UNCCD is actively building capacity for monitoring sand and dust storms — a growing hazard linked to desertification that affects agriculture, health, and infrastructure across Africa, Asia, and the Middle East. AI-powered analysis of satellite and ground-based sensor data can improve early warning systems, predict storm trajectories, and assess impacts on affected populations.

¹⁶⁰ ORF, "High and dry? AI and the fight against desertification," June 2024.

<https://www.orfonline.org/expertspeak/high-and-dry-ai-and-the-fight-against-desertification>

¹⁶¹ Reuters, "China completes 3,000-km green belt around its biggest desert, state media says", November 29, 2024.

<https://www.reuters.com/world/china/china-completes-3000-km-green-belt-around-its-biggest-desertstate-media-says-2024-11-29/>

5.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments — Challenges and Mitigation Issues

Companies involved in agriculture, mining, infrastructure and renewable-energy projects can use AI to monitor land disturbance, plan ecological restoration and reduce soil and vegetation damage. However, corporate commitments should be tied to independent verification. Independent validation should accompany AI-generated assessments wherever they inform ecological restoration, regulatory compliance or environmental reporting. For developing countries, open satellite data and shared analytical tools are especially important because desertification often affects regions with limited monitoring capacity.

Persistent Challenges

The application of AI to combat desertification and land degradation faces more significant barriers than in other environmental domains covered in this paper. Three challenges stand out:

Data scarcity. High-resolution, labelled datasets for training AI models on land degradation processes are limited, particularly in the regions most affected — the Sahel, Central Asia, the Horn of Africa, and parts of South Asia. Unlike climate modelling, which benefits from decades of standardized atmospheric data, land degradation monitoring often relies on fragmented, inconsistent, and low-resolution data sources. The UNCCD has called for enhanced open-data initiatives and standardized data-sharing protocols.

Limited technical capacity. Many of the countries most affected by desertification lack the technical infrastructure, computational resources, and trained personnel to develop, deploy, and maintain AI systems. This capacity gap risks creating a dependency on external expertise and technology — a concern that connects to the broader equity arguments developed in Section 1 and Section 8 of this paper.

Funding constraints. The UNCCD's Financial Needs Assessment (COP16, 2024) estimated that \$2.6 trillion in total investment is needed by 2030, yet current investment levels are a fraction of this amount. AI deployment at scale requires sustained funding for both technology development and operational costs, including data infrastructure, local training programs, and long-term maintenance.

Pathways Forward

Despite these challenges, several developments offer reasons for cautious optimism. The IDRO represents a multilateral, AI-driven solution that is designed for accessibility — users need not be specialists to draw value from the tool, which supports interventions across global, regional, national, and local scales. The World Bank has signaled support for IDRO, noting that drought resilience is critical to its vision of eradicating extreme poverty on a livable planet.

The intersection of desertification with water resources (Section 6) and climate change (Section 3) means that AI solutions developed for these adjacent domains — such as climate forecasting models, water management systems, and agricultural optimisation tools — can potentially be adapted for land degradation contexts. Cross-domain AI applications could accelerate progress while reducing the need for bespoke development.

The policy frameworks needed to support AI deployment in this domain — including mandatory data-sharing, capacity building mechanisms, and funding commitments — are addressed in Section 7 (Policy and Regulatory Considerations) and Section 9 (Next Steps).

6.1 Why Do We Care About Water Resources

Water is both a direct resource input for AI infrastructure and an environmental domain where AI can generate significant benefits. Water scarcity, pollution, flooding and inefficient irrigation all require timely information. AI can help improve forecasting, allocation, leakage detection, water-quality monitoring and flood response. At the same time, the water footprint of data centers means that AI infrastructure must be evaluated against local water stress.

The International Framework

Sustainable Development Goal 6 (SDG 6) seeks to ensure the availability and sustainable management of water and sanitation for all by 2030. Progress towards all SDG 6 targets is currently off track — some severely. Approximately 4 billion people, or half the world’s population, experience severe water scarcity for at least part of the year. An estimated 2.2 billion people — 27% of the global population — were without access to safely managed drinking water in 2022, with four out of five of those living in rural areas. The situation concerning sanitation is worse: 3.5 billion people worldwide lack access to safely managed sanitation.¹⁶²

The water crisis is not static. Global freshwater withdrawals continue to grow at an average rate of 0.7% per year, driven primarily by cities, countries, and regions undergoing rapid economic development. Renewable water availability per person has declined by a further 7% over the past decade, while pressure on already scarce freshwater resources is intensifying in several regions. A 40% global shortfall in freshwater resources is foreseen by 2030 if current trends continue.¹⁶³

Climate change is accelerating seasonal variability and uncertainty in water availability across most regions. Mountain glaciers — the world’s “water towers” — are retreating at unprecedented rates, threatening the freshwater supply of billions of people downstream. The UN World Water Development Report 2025 emphasizes that the urgent need to drastically reduce carbon emissions bears directly on water security: as glaciers disappear, the natural storage systems that regulate river flow throughout the year are lost, increasing both flood risk during melt seasons and drought risk during dry seasons. Pollution, land and ecosystem degradation, and natural hazards further compromise the availability of water resources.

Agriculture is the dominant consumer of global freshwater, accounting for approximately 70% of all withdrawals. Industry accounts for roughly 20%, and domestic use approximately 10%. However, AI-driven data centers represent a rapidly growing and geographically concentrated new category of industrial water use that is not well captured in traditional water accounting frameworks — a point developed in detail below.

The Water–AI Nexus: A Dual Relationship

As established in Section 2.3 of this paper, AI infrastructure itself is a significant and growing consumer of water. Data centers consumed between 312.5 and 764.6 billion liters of water in 2025, a range comparable to the global annual consumption of bottled water. The OECD estimates

¹⁶² UNESCO/UN-Water, World Water Development Report 2025: Mountains and Glaciers — Water Towers, March 2025. <https://www.unwater.org/publications/un-world-water-development-report-2025>

¹⁶³ FAO, 2025 AQUASTAT Water Data Snapshot, December 2025. <https://www.fao.org/newsroom/detail/renewable-water-availability-per-person-plunges-7-percent-in-a-decade-as-global-scarcity-deepens-fao-data-shows/en>

that AI-related water consumption could reach 6.6 billion cubic meters by 2027.

UNEP has noted that this could be equivalent to half of the United Kingdom's total water consumption.¹⁶⁴

This creates a direct and consequential tension. AI is simultaneously one of the most promising tools for managing water resources — through flood prediction, leak detection, irrigation optimisation, and water quality monitoring — and a contributor to water stress, particularly in regions where data centers are concentrated alongside agricultural and domestic users. In the United States, Virginia hosts approximately 26% of total data center capacity; parts of Arizona, Texas, and northern California are also major data center hubs in water-scarce environments.

In India, 50% of data centers are located in extremely water-stressed regions (Bengaluru, Mumbai). Internationally, rapid data center construction in the Middle East and parts of China introduces AI water demand into regions where allocation is already contested.

This dual character — mirroring the integrated framework established in Section 1 — requires that water policy for AI address both sides simultaneously: reducing AI's own water consumption (through efficiency standards, cooling innovation, and siting decisions) and maximising AI's contributions to water security (through the applications described in Section 6.2). The net effect depends on governance choices, not on technology alone.

¹⁶⁴ De Vries, A. "The carbon and water footprints of data centers," *Patterns*, December 2025.

[https://www.cell.com/patterns/fulltext/S2666-3899\(25\)00278-8](https://www.cell.com/patterns/fulltext/S2666-3899(25)00278-8); UNEP, "AI has an environmental problem," September 2024.
<https://www.unep.org/news-and-stories/story/ai-has-environmental-problem-heres-whatworld-can-do-about>

6.2 How AI Is Contributing and Could Contribute to Address Water Resources Issues

AI can support water governance through hydrological forecasting, flood simulation, water quality anomaly detection, irrigation optimisation, reservoir scheduling and digital-twin basin management. The value of AI lies in integrating heterogeneous data streams, including weather, river flow, reservoirs, water gates, irrigation demand, satellite imagery and sensor networks.

Flood Prediction and Early Warning

AI is transforming flood forecasting — one of the most consequential applications of machine learning for human welfare. Floods have a greater impact on people than any other type of natural disaster globally, and the acceleration of flood-related events due to climate change underscores the urgency of effective early warning systems, especially in low- and middle-income countries where 90% of vulnerable populations reside.

Google's **Flood Hub** platform uses AI to predict riverine flooding up to 7 days in advance across over 80 countries, providing street-level forecasts of flood extent and depth to governments, aid organizations, and communities. The system covers approximately 460 million people and provides forecasts freely and without registration barriers. A study published in *Nature* (2024) demonstrated that Google's AI model achieves reliability in predicting extreme riverine events in ungauged watersheds at up to a 5-day lead time that is comparable to or better than the reliability of nowcasts (same-day predictions) from

the Copernicus Emergency Management Service's Global Flood Awareness System — the current state of the art in physics-based global flood modelling.¹⁶⁵

Flood Hub was initially deployed in India and Bangladesh — two countries where monsoon flooding displaces millions annually — and was expanded across Africa in 2022 and to South America in 2024. The World Bank estimates that upgrading flood early warning systems in developing countries to developed-country standards could save an average of 23,000 lives per year. In April 2026, Google extended Flood Hub with a new AI model for flash-flood prediction, capable of forecasting rapid flooding events up to 24 hours ahead. Unlike traditional river forecasting, this tool focuses on rapid flooding driven by intense rainfall — events that are among the deadliest and hardest to predict because they can strike far from riverbanks.¹⁶⁶

These advances demonstrate a model for how AI can be applied to humanitarian challenges at global scale: taking a problem that affects hundreds of millions of people, applying machine learning to publicly available data, and delivering results for free through accessible platforms.

However, the system's accuracy depends on the quality of input data, and in regions with sparse river gauge networks or unreliable weather forecasts, predictions can be less precise. There is also an inverse correlation between the amount of publicly available streamflow data in a country and national GDP — meaning that the countries with the least data are often the most vulnerable.

¹⁶⁵ Nearing, G. et al. "AI Increases Global Access to Reliable Flood Forecasts," *Nature*, 2024.

<https://arxiv.org/pdf/2307.16104>; Google Research blog: <https://research.google/blog/using-ai-to-expand-global-access-to-reliable-flood-forecasts/>

¹⁶⁶ Google blog, "How Google uses AI to improve global flood forecasting," January 2026.

<https://blog.google/technology/ai/google-ai-global-flood-forecasting/>

Urban Water Infrastructure and Leak Detection

One of the most economically significant applications of AI in water management is the detection and prevention of water losses in urban distribution networks. Globally, an estimated 30–40% of treated water is lost to leaks before reaching consumers — a staggering waste of a scarce resource that also represents lost revenue for water utilities and increased energy consumption for water treatment and pumping.

AI-based leak detection systems analyze pressure, flow, and acoustic data from sensors embedded in pipe networks to identify leaks before they become visible. A systematic review of 53 studies published between 2018 and 2025 found that deep learning architectures — convolutional neural networks, LSTM networks, and autoencoders — dominate tasks requiring extraction of complex spatiotemporal patterns from pipeline sensor data. Machine learning approaches have been shown to detect leaks with high accuracy, including small leaks that would be undetectable by traditional methods.¹⁶⁷

Microsoft has invested in AI-driven leak detection through partnerships with FIDO Tech, deploying acoustic sensors and machine learning algorithms across water networks in several countries. The integration of AI with the Internet of Things (AIoT) — combining AI analytics with networked sensor hardware — enables real-time, continuous monitoring of water infrastructure at scales that were previously economically prohibitive. Research published in *Water Resources Management* (2024) demonstrated that hydrophone-based AIoT systems can detect and classify leaks in real-world pipe networks with high reliability, including in noisy urban environments.¹⁶⁸

Precision Agriculture and Irrigation Optimisation

Agriculture consumes approximately 70% of global freshwater withdrawals, making irrigation optimisation one of the highest-impact applications of AI for water conservation. AI-powered precision irrigation systems integrate data from multiple sources — soil moisture sensors, weather forecasts, satellite imagery, and crop growth models — to

determine the optimal timing, quantity, and location of water application for each field or even each plant.

These systems can reduce water use by 20–40% while maintaining or improving crop yields. The benefits are particularly significant in water-stressed regions of South Asia, the Middle East, and sub-Saharan Africa, where agriculture is both the dominant water user and the primary source of livelihoods. Smart irrigation scheduling — where AI algorithms adjust water delivery based on real-time soil and atmospheric conditions — represents one of the most mature and commercially deployed applications of AI for environmental benefit.

However, the adoption of precision irrigation in developing countries remains limited by the cost of sensor infrastructure, the availability of reliable internet connectivity, and the need for training and technical support. These barriers echo the capacity-building challenges identified in Sections 4 and 5 and underscore the need for international cooperation and technology transfer.

India is the world's largest user of groundwater, withdrawing an estimated 230 cubic kilometers per year — about a quarter of the global total — with more than 85% of drinkingwater supplies and 60% of irrigated agriculture dependent on it; on current trends, around 60% of India's aquifers could reach a critical condition within two decades.¹⁶⁹ AI-assisted hydrological systems and satellite monitoring technologies are increasingly being used for aquifer assessment, irrigation optimisation, groundwater forecasting, flood prediction, and leakage detection. AI-enabled systems may therefore contribute both to climate adaptation and to improved wateruse efficiency in highly water-stressed regions. This is particularly significant since groundwater supports a major proportion of India's irrigation and drinking water systems.

Water Quality Monitoring and Prediction

AI models analyze data from satellite imagery, in-situ sensors, and laboratory analyses to detect contamination, predict algal blooms, and monitor compliance with water quality standards in real time. Machine learning can identify pollution sources, forecast water quality changes, and classify water conditions faster and at lower cost than traditional

¹⁶⁷ "AI in Water Distribution Networks: A Systematic Review," *Smart Cities*, March 2026. <https://www.mdpi.com/2624-6511/9/3/45>

¹⁶⁸ "AIoT-Driven Leak Detection in Real Water Networks Using Hydrophones," *Water Resources Management*, December 2024. <https://link.springer.com/article/10.1007/s11269-024-04077-3>

¹⁶⁹ World Bank. *India Groundwater: A Valuable but Diminishing Resource*. Washington (DC): World Bank; 2012 Mar 6.

laboratory-based monitoring methods. AI-enhanced wastewater treatment systems use real-time monitoring and fault detection to optimize treatment processes, reducing both water pollution and the energy consumed by treatment plants.¹⁷⁰

The ability to monitor water quality at scale is particularly important for developing countries, where laboratory capacity is limited and contamination events may go undetected for extended periods. Satellite-based AI monitoring can provide early warnings of eutrophication, sedimentation, and chemical pollution in rivers, lakes, and coastal waters — enabling earlier intervention and reducing the health burden of waterborne diseases.

Groundwater and Watershed Management

AI is increasingly used to model groundwater recharge rates, predict aquifer depletion, and optimize extraction schedules — critical capabilities as groundwater, which accounts for approximately 99% of all liquid freshwater on Earth, faces growing pressure from over-extraction and contamination. Integrated watershed management systems use AI to balance the competing demands of agriculture, industry, energy production, urban supply, and ecosystems across entire river basins, incorporating climate projections and population growth scenarios.

Desalination — the conversion of seawater or brackish water to freshwater — is another area where AI is improving efficiency. Machine learning optimizes the energy-intensive reverse osmosis process, reducing both energy consumption and the environmental impact of brine discharge. As water scarcity intensifies in coastal regions, AI-optimized desalination could become an increasingly important component of water supply portfolios.

China's water-resources sector illustrates how AI, digital twins, and sensor networks can be combined with public-sector institutions for basin-scale governance. In the Shule River basin in Gansu, a virtual model of river channels, canals, reservoirs, and water gates built with geographic information systems and building information modelling uses large-model simulation to optimize water scheduling, automate gate operations, and support precise water release.

China has reportedly completed 94 construction tasks for digital-twin water-resource management, with work underway on the Yangtze and Yellow Rivers, the national backbone water network, and the South-to-North Water Diversion Project's central route; a national ecological-flow monitoring and early-warning platform gathers data from 283 ecological-flow sections and supports real-time warning for 235 control sections across 165 key rivers and lakes.

The case also shows the implementation requirements for AI in environmental governance: long-term monitoring networks, interoperable data, professional hydrological models, scenario simulation, and the institutional authority to act on warnings.¹⁷¹

The IDRO Connection

The International Drought Resilience Observatory (IDRO), discussed in detail in Section 5, is directly relevant to water resources management. Drought is fundamentally a water scarcity event, and the IDRO's AI-driven analysis of social and environmental drought resilience indicators provides actionable information for water managers — not just land managers. The connection between Sections 5 and 6 underscores a broader point: environmental challenges are interconnected, and AI tools developed for one domain can often be adapted for adjacent applications.

¹⁷⁰ "AI-driven approaches for responsible urban water management," Discover Water, Springer, February 2026. <https://link.springer.com/article/10.1007/s43832-026-00365-8>

¹⁷¹ Qiushi / Xinhua, "China steps up efforts to build digital twin systems in water conservancy sector," 16 January 2025. https://en.qstheory.cn/2025-01/16/c_1064820.htm

6.3 Sustainable Solutions Supporting Economic Solutions: Paths to Companies' Commitments

For companies, water-related AI commitments should include both footprint and application dimensions. Data-center operators should disclose direct and indirect water use, cooling technologies, local water-stress conditions and water-reuse measures. Firms applying AI in water management should disclose model performance, uncertainty, data quality and public-interest safeguards. Where AI is used in water allocation or flood response, public accountability is essential because algorithmic errors can have serious social consequences.

Corporate Water Stewardship in the AI Sector

Recognizing the growing scrutiny of data center water consumption, several major technology companies have adopted ambitious water stewardship commitments.

Google committed in 2021 to becoming “water-positive” by 2030, pledging to replenish 120% of the freshwater it consumes across its offices and data centers. By 2025, Google reported having replenished more than 7 billion gallons of water and supported 165 replenishment projects across 97 watersheds globally. The company estimates that these projects will replenish more than 19 billion gallons by 2030 once fully implemented. Google is also investing in freshwater alternatives for cooling, including seawater and reclaimed wastewater.¹⁷²

Microsoft made a similar pledge in September 2020, committing to replenish more water than its operations consume by 2030. By 2024, Microsoft reported having more than 80 water replenishment and access projects worldwide, and launched a new data center design in August 2024 that consumes zero water for cooling — a significant engineering milestone. The company has reduced its water intensity (water consumed per kilowatt-hour) by over 80% from its first generation of data centers to its current generation. Microsoft also met its 2030 target early, providing 1.5 million people with access to clean water and sanitation services through partnerships with Water.org.¹⁷³

Amazon (AWS) has pledged to return more water to communities than it uses by 2030. **Meta** (Facebook) announced a similar water-positive goal in 2021.

These commitments represent meaningful steps, but they face important limitations. As established in Section 2.3, corporate water reporting typically does not distinguish between AI and non-AI workloads, conflating the water footprint of general cloud services with that of AI-specific computation. The concepts of “water-positive” and “replenishment” — while well-intentioned — involve returning water to stressed watersheds near operations, which may not directly offset the water consumed in other locations. And as data center growth accelerates, the absolute volume of water consumed continues to increase even as water intensity per unit of compute declines.

¹⁷² Google, “2026 Water Stewardship Portfolio,” March 2026. <https://blog.google/company-news/outreach-and-initiatives/sustainability/2026-water-stewardship-portfolio/>

¹⁷³ Microsoft, “Sustainable by design: Transforming datacenter water efficiency,” July 2024. <https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/07/25/sustainable-by-design-transforming-datacenter-water-efficiency/>; Microsoft blog, “The journey to water positive,” March 2023. <https://blogs.microsoft.com/on-the-issues/2023/03/22/water-positive-climate-resilience-open-call/>

International Accountability Frameworks

Several international frameworks provide structure for corporate water stewardship, though none currently includes AI-specific requirements:

The **CEO Water Mandate**, a UN Global Compact initiative launched in 2007, commits signatory companies to advancing water stewardship, sanitation, and sustainable development goals. It provides guidance on water risk assessment, stakeholder engagement, and public disclosure.

The **CDP Water Security** program enables companies to disclose their water risks, management practices, and targets to investors and the public. CDP data allows stakeholders to compare corporate water performance across sectors and geographies.

The **Alliance for Water Stewardship (AWS)** International Water Stewardship Standard provides site-level certification for responsible water use, including criteria for good water governance, sustainable water balance, good water quality status, and important water-related areas.

The **Taskforce on Nature-related Financial Disclosures (TNFD)**, which builds on the climate-focused TCFD framework, includes freshwater ecosystems among its priority areas — creating a potential pathway for integrating water impact disclosure into mainstream financial reporting.

However, none of these frameworks currently requires or enables AI-specific water disclosure.

As AI becomes a significant category of industrial water use, the absence of AI-specific metrics and reporting requirements represents a growing governance gap.

The Regulatory Landscape

Water governance for data centers is significantly less developed than energy governance.

While the EU Energy Efficiency Directive (2023) mandates energy reporting for large data centers, no equivalent mandatory water reporting framework exists in most jurisdictions. The UNEP Sustainable Procurement Guidelines for Data Centers and Servers (June 2025) include Water Usage Effectiveness (WUE) among recommended reporting metrics, but compliance is voluntary.¹⁷⁴

Oregon (USA) mandated water reporting for large cooling users in June 2025 — the first US state to do so — but this remains an exception rather than the rule. No jurisdiction currently requires data centers to conduct water impact assessments before construction, assess the cumulative water demand of multiple data centers in the same watershed, or demonstrate that water consumption will not adversely affect other users or ecosystems.

This regulatory gap is particularly concerning in regions facing water stress. Building waterintensive AI infrastructure in water-scarce regions is not merely an environmental risk but a social equity concern — particularly for communities that depend on the same water sources for agriculture and drinking water. The intersection of AI's water footprint with the environmental challenges addressed in Section 5 (desertification and land degradation) is direct: drought is a water scarcity event, and AI infrastructure that consumes water in droughtprone regions exacerbates the very problem that AI drought-management tools (like the IDRO) are designed to address.

The policy frameworks needed to close this gap — including mandatory water disclosure, AI-specific disaggregation, cumulative impact assessment, and integration of water considerations into data center siting decisions — are addressed in Section 7. The broader question of how to balance AI's water costs against its water management benefits is a governance question that must be resolved through the kind of inclusive, multistakeholder dialogue that the IGF Dynamic Coalition on Environment is designed to facilitate.

¹⁷⁴ UNEP, Sustainable Procurement Guidelines, June 2025.

<https://www.unep.org/technical-highlight/unepreleases-guidelines-curb-environmental-impact-data-centres>

Section 7.

Policy and Regulatory Considerations

7.1 Comparative Overview	64
7.2 Jurisdiction Profiles	67
7.3 Key Regulatory Gaps	73



7.1 Comparative Overview

The governance of AI's environmental impact is developing unevenly across jurisdictions. No country has yet adopted a comprehensive regulatory framework that fully integrates AI governance with environmental governance. However, several jurisdictions are advancing partial solutions across different dimensions of the problem, creating a patchwork that is instructive — both for what it achieves and for what it leaves unaddressed.

This section analyses the policy landscape across six jurisdictions — BRICS, China, the European Union, India, Russia, and the United States — organized first around four crosscutting dimensions that enable direct comparison, then through jurisdiction-specific profiles that capture the institutional context and trajectory of each.

The four cross-cutting dimensions are: environmental disclosure requirements for AI and data centers; renewable energy and efficiency standards; data center environmental constraints (including siting restrictions); and the degree of linkage between AI governance frameworks and environmental governance frameworks. The analysis draws on the findings of Sections 2–6, which established the evidentiary base, and represents the shift from description to comparative analysis focused on practical governance parameters.

Cross-Cutting Comparison

Environmental disclosure. The European Union leads globally with the most comprehensive mandatory data center environmental reporting framework. The recast Energy Efficiency Directive (EU/2023/1791) introduced reporting obligations for data centers with an IT power demand of at least 500 kW, effective from May 2024. A European database, administered by the European Commission, now collects and

publishes data on energy performance and water footprint of qualifying data centers. The Commission published its first analytical report based on 2024 reporting data in July 2025. However, AI-specific disaggregation is not required — data centers report aggregate energy use without distinguishing between AI and non-AI workloads.¹⁷⁵

China's March 2025 guidelines from the National Development and Reform Commission (NDRC) push data center hubs towards higher renewable electricity shares, with mandatory green power credit purchases for operators in designated computing regions. While not framed as “environmental disclosure” per se, these requirements impose substantive environmental obligations on data center operators.¹⁷⁶

The United States relies primarily on voluntary corporate disclosure, supplemented by emerging state-level requirements. A January 2025 Executive Order directed the Department of Energy to draft reporting requirements for AI data centers covering their entire lifecycle, including embodied carbon, water usage, and waste heat — but these have not yet been codified as binding regulation. Oregon mandated water reporting for large cooling users in June 2025, becoming the first US state to impose such a requirement.

India, Russia, and BRICS frameworks as collective bodies have not yet introduced mandatory AI or data center environmental disclosure requirements, though individual BRICS members (China) have begun doing so.

Renewable energy and efficiency standards. China has implemented the most direct government-mandated linkage between data centers and renewable energy. The NDRC guidelines require

¹⁷⁵ European Commission, Energy Performance of Data Centres.

https://energy.ec.europa.eu/topics/energyefficiency/energy-efficiency-directive/energy-performance-data-centres_en

¹⁷⁶ ChinaUSFocus, “Environmental AI Governance: U.S. and China Have Different Roads,” January 2026.

<https://www.chinausfocus.com/energy-environment/environmental-ai-governance-us-and-china-have-differentroads-to-developing-green-ai-systems>

data centers in designated computing regions to purchase green power credits, and provinces such as Inner Mongolia offer incentives for data centers to pair directly with local wind and solar farms. Between 2023 and 2024, over 500 data center infrastructure projects were announced in China. In the first half of 2025, China installed 357 GW of new wind and solar capacity — more than India’s entire installed power capacity — providing a rapidly expanding base of renewable energy for AI infrastructure.¹⁷⁷

Germany’s transposition of the EU Energy Efficiency Directive — the Energy Efficiency Act (Energieeffizienzgesetz, EnEFG) — goes beyond the EU minimum by mandating a renewable electricity share of 50% for data centers (rising to 100% from January 2027), energy reuse obligations for new data centers from July 2026, and binding PUE compliance from July 2027. This represents the most stringent set of data center energy requirements in any major economy.¹⁷⁸

International PUE benchmarks vary: Singapore requires 1.3 or lower for new data centers; Australia requires 1.4 or lower; California and China require 1.5 or lower. The EU has not yet set binding PUE targets at the Union level, though this is expected in the Data Center Energy Efficiency Package planned for Q2 2026.¹⁷⁹

In the US, major technology companies have made voluntary commitments (Google’s 24/7 carbon-free energy, Amazon’s Climate Pledge), and Virginia proposed PUE efficiency targets tied to tax incentives in 2024. India targets 500 GW of non-fossil electricity capacity by 2030 but has not linked this specifically to AI or data center requirements.

Data center environmental constraints. Ireland paused new data center construction near Dublin until 2028 due to energy grid concerns — the most restrictive measure taken by any jurisdiction, reflecting data centers’ consumption of over 20% of Ireland’s total electricity.

Singapore imposed a temporary moratorium on new data centers from 2019 to 2022, lifting it with stringent efficiency conditions.

China’s “East Data, West Computing” program, launched in 2022, directs new data center construction to eight national computing-hub nodes in western China (Guizhou, Gansu, Inner Mongolia, and others), reducing pressure on energy- and waterstressed eastern regions while leveraging western China’s abundant renewable energy and cooler climates. The programme establishes eight national computing hub-nodes — Beijing-Tianjin-Hebei, the Yangtze River Delta, the Guangdong-Hong Kong-Macao Greater Bay area, Chengdu-Chongqing, Inner Mongolia, Guizhou, Gansu, and Ningxia — supported by ten national data-centre clusters. It guides computing demand from eastern China towards western computing resources in an orderly manner rather than concentrating all hubs in western China. No other major jurisdiction guides computing demand from eastern regions towards designated national computing-hub nodes while restricting the construction of large new data centres outside those designated hubs.

AI–environment governance linkage. This is the most significant gap across all jurisdictions. AI governance frameworks and environmental governance frameworks have been developed largely independently, and no country has yet achieved a comprehensive integration of the two.

The **EU AI Act** (Regulation 2024/1689) is the world’s most comprehensive AI regulation, but its environmental provisions were significantly weakened during trilogue negotiations. The European Parliament’s 2023 position included requirements for providers of foundation models to adhere to standards for reducing energy and resource use, to enable measurement and logging of environmental impact across the full lifecycle, and for high-risk AI systems to include energy consumption in logging capabilities. However, the final text reduced these to: a requirement for GPAI model providers to disclose “known or estimated energy consumption of the model” (Annex XI); a mandate for the Commission to request standardization bodies to develop energy efficiency standards; and voluntary codes of conduct on environmental sustainability (Article 69). The Heinrich Böll Foundation described this as a “missed opportunity,” noting that the final

¹⁷⁷ Brookings Institution, “Global energy demands within the AI regulatory landscape,” April 2026. <https://www.brookings.edu/articles/global-energy-demands-within-the-ai-regulatory-landscape/>

¹⁷⁸ White & Case, “Energy efficiency requirements under the EU AI Act,” 2025. <https://www.whitecase.com/insight-alert/energy-efficiency-requirements-under-eu-ai-act>

¹⁷⁹ ACM Europe Policy Brief, “Powering Europe’s Digital Transformation,” June 2025. https://europe.acm.org/binaries/content/assets/public-policy/europe-tpc/acm_data_center_final.pdf

provisions apply only to GPAI models (not all AI systems), focus exclusively on energy (ignoring water, minerals, and e-waste), and do not require continuous monitoring after deployment — leaving the inference phase undocumented.¹⁸⁰

China's **Global AI Governance Action Plan**, presented at the 2025 World Artificial Intelligence Conference (WAIC) in Shanghai in July 2025, represents the most explicit integration of AI governance and environmental sustainability in any major policy document.

The Action Plan advocates for “jointly establishing AI energy and water efficiency standards” and promoting “green computing technologies such

as low-power chips and efficient algorithms.” It also prioritizes digital infrastructure support for developing countries, including clean electricity and data centers, and calls for a unified computing power standards system. However, the Action Plan is a policy proposal rather than binding legislation.¹⁸¹

The European Commission has signaled that further integration is forthcoming. A Data Center Energy Efficiency Package, a Strategic Roadmap on Digitalization and AI for the Energy Sector, and a Cloud and AI Development Act are all planned for 2026, with the latter aiming to triple EU data center processing capacity while requiring compliance with energy efficiency, water efficiency, and circularity standards.¹⁸²

¹⁸⁰ Böll Foundation, “The EU AI Act and environmental protection: the case for a missed opportunity,” April 2024.

<https://eu.boell.org/en/2024/04/08/eu-ai-act-missed-opportunity>;

Ebert, K. et al. “AI, Climate, and Regulation,” ACM 2025. <https://arxiv.org/html/2410.06681v2>

¹⁸¹ Global AI Governance Action Plan, WAIC 2025. https://www.fmprc.gov.cn/mfa_eng/xw/zyxw/202507/t20250729_11679232.html

¹⁸² European Commission, White & Case, October 2025.

<https://www.whitecase.com/insight-alert/data-centresand-energy-consumption-evolving-eu-regulatory-landscape-and-outlook-2026>

7.2 Jurisdiction Profiles

BRICS

The BRICS+ framework for AI cooperation, adopted at the 2024 Kazan Summit, includes climate forecasting among priority areas for joint AI research. The BRICS+ AI Alliance, launched in December 2024, aims to coordinate AI development across member states. The New Development Bank (NDB) launched a \$5 billion digital sovereignty fund in 2025 for AI infrastructure in BRICS countries. However, BRICS AI cooperation frameworks do not yet include specific environmental governance provisions or shared environmental standards for AI infrastructure. This represents an opportunity: as BRICS members collectively account for a large and growing share of global data center capacity, the adoption of shared environmental benchmarks — even voluntary ones — could have significant impact.¹⁸³

The BRICS dimension is important because AI and environmental governance are increasingly shaped by multipolar infrastructure, standards and development priorities. BRICS cooperation can help broaden the AI-environment agenda beyond the US-EU regulatory frame by adding issues of development, infrastructure access, energy security, digital sovereignty, open technology and capacity building. Therefore, BRICS is not merely a geographic grouping but a platform for discussing how emerging economies can govern AI's environmental footprint while using AI to address development-related environmental challenges.

The policy question for BRICS is how to create interoperable but development-sensitive governance tools. Potential areas include common reporting templates for data-center energy and water use, shared datasets for climate and biodiversity applications, joint work on AI for renewable integration, green-computing standards, and

technology cooperation for Global South capacity building.

China

China's AI-environment governance can be understood through the interaction of **several policy streams described below**, demonstrating that AI sustainability is not confined to AI regulation. It is also embedded in energy planning, data-center approval, industrial policy, ecological monitoring and public-sector digital transformation.

China's current AI governance framework is anchored by the State Council's Opinions on Deepening the Implementation of the «AI Plus» Initiative (Guo Fa [2025] No. 11), issued in August 2025. As China's first comprehensive top-level policy framework for AI development, it positions AI-enabled ecological governance as one of the country's strategic priorities under the Beautiful China initiative while also promoting international cooperation on AI governance. The more specific measures on green data centres, computing-power coordination and ecological monitoring discussed below should therefore be understood within this broader strategic framework.

(1) **International norm-setting and sustainable development.** On the level of global principles, the 2024 Shanghai Declaration on Global AI Governance explicitly calls for AI to be used in industrial innovation, environmental protection, resource utilisation, energy management, and biodiversity, while emphasizing global cooperation, the role of the United Nations, North-South and South-South cooperation, and increased representation and voice for developing countries.¹⁸⁴ The 2025 Global AI Governance Action Plan further frames AI as an international public good and calls for digital infrastructure, intelligent computing power, data centers, and clean power to

¹⁸³ BRICS Council, "Joint BRICS Projects in AI."

<https://bricscouncil.ru/en/analytics/sovlestnye-proekty-stranbriks-v-oblasti-iskusstvennogo-intellekta>

¹⁸⁴ Ministry of Foreign Affairs of the People's Republic of China, "Shanghai Declaration on Global AI Governance," 4 July 2024.

https://www.mfa.gov.cn/eng/xw/zyxw/202407/t20240704_11448351.html

be developed in a more inclusive and interoperable manner, in support of the 2030 Agenda.¹⁸⁵

(2) Green standards for digital infrastructure.

At the level of data-center infrastructure, China's Special Action Plan for the **Green and Low-Carbon Development of Data Centers** attaches environmental requirements directly to computing-power expansion: reducing average data-center PUE below 1.5 by 2025, increasing renewable-energy utilisation by 10% annually, and reaching internationally advanced levels in PUE, energy efficiency, and carbon efficiency per unit of computing power by 2030.¹⁸⁶ According to official reporting, the average PUE across the eight national computing-hub clusters is approximately 1.3, with the most advanced facilities achieving values as low as 1.04.

(3) AI–energy system integration. At the level of the power system, China's 2025 AI-energy integration plan sets 2027 and 2030 targets that provide a concrete timeline for using AI in energy security, power-grid operation, renewable energy, nuclear power and the coordination of computing power with electricity supply.¹⁸⁷ This stream reflects a more operational governance model, in which AI is treated simultaneously as a tool for improving power-system efficiency and as a new source of electricity demand that must be coordinated with grid planning.

(4) AI for ecological monitoring and enforcement.

At the level of environmental administration, China's Ministry of Ecology and Environment has promoted the use of AI, big data, and cloud computing to modernize environmental governance capacity. Current applications include biodiversity monitoring, forest scanning, off-site environmental law enforcement, vehicle-emissions fraud detection, and the construction of integrated sky-ground-sea intelligent sensing networks.¹⁸⁸

(5) AI governance with environmental provisions.

At the level of AI regulation itself, China operates one of the world's most extensive frameworks, combining national legislation, administrative measures, and provincial policies, and increasingly attaching environmental conditions to AI and data-center activity. The overarching frame is set by China's dual carbon goals (carbon peaking before 2030 and carbon neutrality before 2060) into which AI infrastructure is

progressively being integrated. Three mechanisms link AI regulation to environmental outcomes most directly.

First, the National Development and Reform Commission's March 2025 guidelines mandate renewable-energy procurement for data centers in designated computing regions, embedding green-power requirements into project approval.

Second, the "East Data, West Computing" program strategically redistributes new data-center construction towards eight western computing hubs, leveraging renewable energy and cooler climates while easing pressure on energy- and water-stressed eastern regions.

Third, provincial governments have developed complementary measures: Inner Mongolia offers incentives for data centers to pair directly with wind and solar farms, and Shanghai has partnered with China Telecom on green data-center initiatives. The scale these instruments govern is substantial: between 2023 and 2024, more than 500 data-center infrastructure projects were announced in China, with at least 150 operational by the end of 2024.

Technical and self-governance instruments complete the picture — the National Guidelines for the Development of a Comprehensive Standardization System for the AI Industry (2024) provide the standards framework, and, since 2017, an integrated approach combining government oversight with industry self-governance has developed, as reflected in the Generative AI Industry Self-Discipline Initiative (2024). Taken together with the 2025 Global AI Governance Action Plan — the most significant policy statement to date linking AI governance to environmental sustainability at the international level, issued alongside China's proposal to establish a global AI cooperation organization that could potentially include environmental dimensions — these instruments show AI governance and environmental governance beginning to converge within a single national framework. Taken together, these initiatives illustrate a layered governance architecture in which top-level strategic planning (AI Plus), infrastructure planning (East Data, West Computing), sector-specific regulation, and ecological monitoring operate as complementary

¹⁸⁵ Ministry of Foreign Affairs of the People's Republic of China, "Global AI Governance Action Plan," 26 July 2025. https://www.fmprc.gov.cn/mfa_eng/xw/zyxw/202507/t20250729_11679232.html

¹⁸⁶ State Council of the People's Republic of China. "China sets green targets for data centers." July 24, 2024.

¹⁸⁷ State Council of the People's Republic of China. "China unveils plan on AI-energy integration to drive green transition." September 8, 2025. https://english.www.gov.cn/news/202509/08/content_WS68be8c3ec6d0868f4e8f566d.html

¹⁸⁸ China Daily. "AI transforms China's conservation efforts." March 31, 2026. <https://www.chinadaily.com.cn/a/202603/31/WS69cb2b88a310d6866eb40dda.html>

components of China's AI-environment governance framework.

European Union

The EU has the most layered regulatory framework touching AI environmental impacts, though it remains fragmented across multiple instruments that were developed independently. The key instruments are:

The **EU AI Act** (Regulation 2024/1689): the world's first comprehensive AI law, entered into force August 2024 with phased enforcement over several years. Its environmental provisions, while weakened from the Parliament's original position, still require GPAI model providers to disclose energy consumption data and mandate the development of energy efficiency standards by standardization bodies. The AI Office can demand technical documentation on energy consumption from providers without prior notice. Providers of GPAI models launched before August 2025 have a two-year compliance grace period; those launched after face immediate obligations.

The **Energy Efficiency Directive** (EU/2023/1791): introduced mandatory reporting for data centers with 500 kW or more, including energy performance and water footprint indicators.

The first European database report was published in July 2025. The Commission is preparing a Data Center Energy Efficiency Package for Q2 2026 that will include a rating scheme and the first EU minimum performance standards for data centers.

The **Corporate Sustainability Reporting Directive** (CSRD): effective from 2024, requires large EU companies to report on sustainability under European Sustainability Reporting Standards, including energy consumption, GHG emissions, and water use.

The **EU Taxonomy** Climate Delegated Act: classifies data center activities against climate mitigation criteria, building on the European Code of Conduct for Energy Efficiency in Data Centers.

The gap between AI-specific governance (AI Act) and environmental governance (EED, CSRD, Green Deal) remains the most significant structural weakness. The AI Act does not address water, materials, or e-waste; the EED does not distinguish AI from other data center workloads. The planned 2026 legislative package — Data Center Energy

Efficiency Package, Cloud and AI Development Act, Strategic Roadmap on Digitalization and AI for the Energy Sector — represents the most ambitious attempt to bridge this gap.

The European Union provides a regulatory model contrasting to that of China. It is centered on risk-based AI regulation, data governance, product safety, environmental policy and corporate sustainability reporting. The EU AI Act establishes a broad legal framework for AI safety and accountability, while environmental disclosure and energy-efficiency requirements may be addressed through other instruments. This shows that environmental considerations can be distributed across AI law, energy law, product regulation and sustainability reporting.

India

Despite rapid AI infrastructure expansion, India currently lacks comprehensive environmental disclosure frameworks specifically governing hyperscale AI infrastructure and data centers.

Existing regulations do not yet comprehensively require reporting related to lifecycle carbon emissions, water consumption, embodied emissions, cooling efficiency, or AI-specific sustainability metrics.

At the same time, Indian states are actively competing to attract AI and data-center investments through infrastructure incentives and industrial policies. This creates an important governance challenge involving the balance between digital infrastructure growth, climate commitments, water sustainability, and environmental protection. India therefore represents a globally important policy environment for understanding how rapidly developing economies may integrate AI expansion with environmental resilience and sustainable development objectives.

The scale of that expansion underscores the stakes. India AI Mission (2024) and the development of homegrown AI models such as BharatGen (launched June 2025) prioritize digital sovereignty and economic competitiveness; the Mission has been allocated approximately ₹10,300 crore (about \$1.2 billion) over five years, encompassing 38,000 GPUs, and is paired with a semiconductor push (SEMICON India, with announced investments exceeding \$15 billion). India's AI adoption rate stands at 59%, and data-center capacity is projected to reach 2,073 MW by 2027 — an 85% increase from 2025 levels.¹⁸⁹

¹⁸⁹ Insights on India, "India's Focus on AI and Its Environmental Impact," January 2026. <https://www.insightsonindia.com/2026/01/14/indias-focus-on-ai-and-its-environmental-impact/>

Against this backdrop, three structural gaps illustrate the governance challenge concretely. First, India's Environmental Impact Assessment (EIA) system is designed for manufacturing plants, mines, and heavy industry, not for cloud-based AI operations, so GPU clusters and data centers currently operate without environmental clearance. Several major data centre hubs in India, including Bengaluru and Mumbai, are located in regions experiencing significant water stress and that AI hardware refresh cycles generate e-waste that flows into informal recycling systems lacking advanced mineral-recovery capabilities. Third, although India's target of 500 GW of non-fossil electricity capacity by 2030 is among the world's most ambitious renewable-energy commitments, there is no explicit policy linking renewable-energy deployment to data-center requirements, and the National Data Center Policy, still in development, has not yet incorporated environmental performance standards.

Russia

Russia's distinctive contribution to the AI-environment debate lies in the convergence of two national priorities: a large-scale, low-carbon energy build-out positioned explicitly as the foundation for AI compute, and a broad environmental-policy agenda anchored in long-term carbon neutrality and a dedicated national development goal.

Russia's AI governance framework itself is centralized and strategy-driven, resting on Presidential Decree No. 490 (2019) on the development of artificial intelligence and the National Strategy for the Development of Artificial Intelligence through 2030, updated by Presidential Decree No. 124 (February 2024) with new targets for AI adoption across priority economic sectors.¹⁹⁰

The most distinctive feature of the Russian approach is its framing of energy as the enabling advantage for AI. Speaking at the AI Journey 2025 conference in Moscow (November 2025), President Vladimir Putin emphasized that the primary focus of global AI investment is “on expanding computing power and generating additional gigawatts of energy,” noting that electricity consumption by data centers “will more than triple in this decade alone.” He linked the siting of AI data centers directly to the development of national energy infrastructure — including clean sources such as nuclear power — and stated that the

growing potential of domestic nuclear energy would allow Russia to consistently expand its computing capacity for AI. Russia announced plans to build 38 nuclear power units within two decades, to move towards serial production of small modular reactors, and to continue constructing data centers at major nuclear plants.¹⁹¹ This represents an energy-led, supply-side pathway to lower-carbon AI compute that complements the demand-side efficiency focus dominant in other jurisdictions — though, as with all nuclear expansion, it carries its own water-consumption and lifecycle considerations of the kind discussed in Sections 2 and 3.

On the environmental side, Russia has committed to achieving carbon neutrality by 2060 and has established “Environmental Well-being” as one of its national development goals under Presidential Decree No. 309 (May 2024). The goal sets concrete 2030 targets for waste recycling, reducing hazardous emissions, restoring water bodies, and protecting ecosystems.¹⁹²

As in most jurisdictions, AI governance and environmental governance in Russia are not yet formally joined — there is no specific regulatory framework requiring environmental performance disclosure for data centers or AI systems. What makes the Russian case instructive is a gradual evolution of environmental legislation towards the operational accounting tools on which any future AI-environment framework would depend. This trajectory runs from extended producer responsibility and waste regulation, through the federal climate law (Federal Law No. 296-FZ of 2021 “On Limiting Greenhouse Gas Emissions,” which introduced mandatory emissions reporting for large emitters), to the Sakhalin carbon-neutrality experiment, which is piloting a regional emissions-trading and carbon-accounting system.¹⁹³

Each of these instruments builds elements of the measurement and accounting infrastructure that linking AI operations to environmental governance would ultimately require.

Beyond governance, Russia's contribution is especially relevant in three areas aligned with this paper's multipolar approach. The first is multilingual AI: at AI Journey 2025, President Putin stressed the priority of developing national language models —

¹⁹⁰ Decree of the President of the Russian Federation No. 490 of 10 October 2019, “On the Development of Artificial Intelligence in the Russian Federation”; Decree No. 124 of 15 February 2024.

¹⁹¹ President of Russia, “AI Journey international conference,” 19 November 2025. <http://en.kremlin.ru/events/president/news/78498>

¹⁹² Decree of the President of the Russian Federation No. 309 of 7 May 2024, “On the National Development Goals of the Russian Federation for the period up to 2030 and for the future up to 2036”; Government of the Russian Federation, Unified Plan for Achieving the National Development Goals. <http://static.government.ru/media/files/ZsnFICpxWknEXeTfQdmcFHNei2FhcR0A.pdf>

¹⁹³ Federal Law No. 296-FZ of 2 July 2021, “On Limiting Greenhouse Gas Emissions”; Federal Law No. 34-FZ of 6 March 2022 on the Sakhalin experiment.

“both fundamental and smaller, industry-specific ones” — fully trained and overseen domestically.¹⁹⁴ Russian-language foundation models such as Sber’s GigaChat and Yandex’s models, combined with one of the world’s largest non-English digital ecosystems, help broaden AI beyond its predominant English-language base — addressing the linguistic-bias concern raised in Section 9, where approximately 90 percent of AI training data is in English while only a minority of the world’s population are native English speakers. The second is cold-region environmental monitoring: as documented in Sections 3 and 5, Russian institutions apply AI to boreal forests, wildfire detection, illegal-logging identification, and permafrost — high-latitude ecosystems that are critical to the global carbon balance yet underrepresented in global environmental datasets. The third is the BRICS+ AI cooperation agenda: Russia actively advances multilateral AI cooperation, including a dedicated BRICS+ AI forum convened at AI Journey 2025 and proposals to harmonize AI-related legislation among partner states, with President Putin describing a “joint contour” in AI to be built with partners within BRICS, the SCO, and other formats.¹⁹⁵

Finally, Russian institutions contribute internationally significant open-source tools for measuring and reducing the carbon footprint of computation. As discussed in Section 2.4, Eco2AI and Eco4cast — developed by Sber and the AI Research Institute (AIRI) — enable carbon tracking and temporal-geographic optimisation of AI workloads, reinforcing this paper’s argument that Green AI methods and tools are emerging well beyond the United States and the European Union.

United States

The United States remains central to the AI-environment debate because many of the world’s frontier AI developers and hyperscale cloud providers are based there. It is the dominant AI developer globally, with 40 notable models produced in 2024, \$252.3 billion in corporate AI investment, and over 33% of the world’s data centers. The strain on the power system is already visible: data centers have exceeded 10% of electricity consumption in at least five US states,¹⁹⁶ and the tension between rapid AI infrastructure expansion and grid capacity has become one of the defining features of the American case.

Unlike China’s infrastructure-planning-oriented model or the European Union’s law-centered model, the US approach is led primarily by corporate action and state-level variation rather than national coordination. Federal governance operates through executive orders and voluntary frameworks rather than comprehensive legislation. A January 2025 Executive Order directed the Department of Energy to draft data-center lifecycle reporting requirements covering embodied carbon, water usage, and waste heat — the most detailed federal proposal to date, though not yet enacted. The NIST AI Risk Management Framework provides voluntary guidance but does not address environmental performance. At the congressional level, the Artificial Intelligence Environmental Impacts Act, introduced in early 2024, would task the Environmental Protection Agency with studying AI’s environmental footprint and developing measurement standards, but it has not yet been enacted.

In the absence of mandatory federal standards, corporate disclosure and clean-power procurement have become the principal drivers of environmental performance. The major hyperscale developers have made the most ambitious voluntary commitments: Google has committed to 24/7 carbon-free energy by 2030, Microsoft to becoming carbon-negative, and Amazon to net-zero by 2040 under The Climate Pledge. These companies have also become the largest corporate purchasers of clean power, and their procurement increasingly extends to nuclear energy to secure round-the-clock, low-carbon baseload for AI workloads — including Microsoft’s agreement to restart the Three Mile Island plant, Google’s contract with Kairos Power for small modular reactors, and Amazon’s acquisition of a data-center campus adjacent to the Susquehanna nuclear plant. As established in Section 3, however, these commitments coexist with rising absolute emissions (Microsoft’s emissions rose 29% and Google’s 48% in recent years), as infrastructure growth outpaces efficiency gains.

US-based institutions are also leading much of the world’s model-efficiency and Green AI research — the foundational work on measuring and reducing AI’s energy footprint (Strubell et al.; Schwartz et al.) and key open-source tools originated largely in US and US-linked research environments, and frontier developers continue to drive advances in model architecture, inference optimisation, and custom

¹⁹⁴ President of Russia, 19 November 2025, *ibid.*

¹⁹⁵ President of Russia, 19 November 2025, *ibid.*; RIA Novosti, “Putin instructed the wide use of AI technologies”, 19 November 2025, in Russian. <https://ria.ru/20251119/putin-2056109778.html>

¹⁹⁶ Stanford HAI, AI Index Report 2025. <https://hai.stanford.edu/ai-index/2025-ai-index-report>; IEA, Energy and AI, 2025

energy-efficient hardware (such as Amazon's Graviton, Trainium, and Inferentia chips).

On data-center siting, governance is fragmented across states. Virginia — host to the world's largest concentration of data centers — proposed PUE efficiency targets tied to tax incentives in 2024; Oregon mandated water reporting for large cooling users in June 2025. At the same time, several states offer tax breaks or expedited permitting for data centers, sometimes without environmental conditions.

This combination of corporate leadership and state-level variation creates both opportunities and risks.

On one hand, well-resourced companies can adopt ambitious voluntary standards and finance clean-power and efficiency innovation faster than regulation would require. On the other, there is no national floor of mandatory environmental performance, and states may compete on regulatory laxity to attract data-center investment. The American model thus stands in instructive contrast to China's planning-led approach and the EU's law-led approach: where China embeds environmental requirements into infrastructure approval and the EU distributes them across binding legislation, the United States relies on the voluntary ambition of a small number of dominant firms — effective at the frontier, but uneven across the system as a whole.

7.3 Key Regulatory Gaps

Building on the jurisdiction-specific analysis presented in Section 7, this subsection identifies six regulatory gaps that continue to hinder the development of environmentally sustainable AI.

These gaps span the entire AI lifecycle and cut across technical, institutional, and governance domains. Addressing them will require coordinated action by governments, industry, standards bodies, researchers, and international organizations.

Disclosure Gap

Many AI developers, cloud service providers, and data center operators do not disclose environmental information at a level that enables meaningful comparison across organizations or technologies. Existing sustainability reports typically aggregate emissions, energy consumption, and water use at the corporate or facility level, while AI-specific impacts remain largely invisible. Although standards such as ITU-T L.1801 and ISO TR 20226:2025 provide a foundation for more consistent reporting, adoption remains voluntary in most jurisdictions.

The absence of harmonized disclosure requirements limits transparency, hinders benchmarking, and constrains evidence-based policymaking.

Inference Gap

Public debate, research, and regulation have focused predominantly on the environmental impacts of training large AI models. However, as AI systems move into widespread deployment, inference may become the dominant source of operational energy consumption, water use, and associated emissions. Billions of daily interactions with foundation models can collectively generate environmental impacts that exceed those associated with model training.

Despite this shift, inference-related resource consumption remains poorly measured, inconsistently reported, and largely absent from regulatory frameworks. This creates a growing mismatch between the focus of governance efforts and the evolving environmental footprint of AI systems.

Water Gap

Water consumption remains one of the least visible dimensions of AI sustainability.

Discussions of AI environmental impacts often focus on electricity demand and carbon emissions, while overlooking the substantial water requirements associated with semiconductor manufacturing, electricity generation, and data center cooling. Even where water-use data are reported, information on local water stress conditions is often unavailable.

As a result, identical levels of water consumption may have very different environmental consequences depending on location. Integrating water-use reporting and watershed-level risk assessment into AI governance frameworks remains an important unresolved challenge.

Geographic Equity Gap

The environmental costs and benefits of AI are often distributed unevenly across regions. AI workloads may be processed in locations where electricity generation, cooling systems, and water resources bear the environmental burden, while the economic and productivity benefits accrue elsewhere. Similar asymmetries exist across the AI supply chain, where mining, manufacturing, and electronic waste management are frequently concentrated in regions with lower regulatory capacity and higher environmental vulnerability. Current governance

frameworks rarely address these transboundary impacts, creating a gap between local environmental costs and global economic gains.

Application Accounting Gap

While growing attention is being paid to measuring the environmental footprint of AI systems, far less progress has been made in measuring their environmental benefits. Claims regarding avoided emissions, energy savings, improved resource efficiency, biodiversity protection, or climate adaptation outcomes are often based on project-specific methodologies that are difficult to compare or independently verify. Unlike carbon accounting, there is currently no widely adopted framework for assessing the environmental performance of AI applications in a consistent manner. Without standardized methodologies, it remains difficult to evaluate whether the environmental benefits generated by AI outweigh the environmental costs associated with its development and deployment.

Public Accountability Gap

AI is increasingly being used in environmental monitoring, regulatory enforcement, disaster management, and resource allocation. These applications may influence decisions affecting land use, water access, environmental compliance, and public investment. However, governance frameworks often provide limited safeguards regarding transparency, explainability, contestability, and human oversight. Individuals, communities, and organizations affected by AI-assisted environmental decisions may have little visibility into how decisions are made or how errors can be challenged. Ensuring public accountability therefore requires not only technical accuracy but also procedural safeguards, including explainability requirements, appeal mechanisms, independent audits, and meaningful human review.

Together, these six gaps highlight the limitations of current governance approaches and define priority areas for future policy development. The recommendations presented in Section 9 are designed to address these shortcomings through coordinated action across stakeholders and governance levels.

Section 8.

Role of Public Sector, Private Sector, Civil Society and International Organizations

8.1 Role of Public Sector	76
8.2 Role of International Organizations	78
8.3 Role of Private Sector	80
8.4 Role of Civil Society	82



The environmental governance of AI is not the responsibility of any single category of actor. It requires coordinated action across governments, international organizations, the private sector, and civil society — each contributing capabilities that the others lack. This section analyses what each category of actor is best positioned to do, where complementarities exist, and where gaps in coordination persist. It responds to the Fudan University comments calling for a more operational approach: not merely describing what actors do, but specifying what they should do and how their roles reinforce one another.

The responsibilities outlined below should be understood as complementary rather than hierarchical. Effective AI-environment governance requires coordinated action across governments, international organisations, academia, civil society and the private sector.

8.1 Role of Public Sector

Governments play three essential and irreplaceable roles in aligning AI development with environmental sustainability: establishing regulatory frameworks, exercising procurement leverage, and directing public research investment.

Public authorities should set measurement standards, require environmental disclosure for data centers and high-impact AI systems, coordinate computing-power expansion with energy and water planning, fund AI for environmental public goods, and ensure that AI-based environmental decisions remain accountable. Public-sector capacity is essential when AI is used for biodiversity monitoring, water management and environmental enforcement, as shown in cases of China.

Regulatory Frameworks

As Section 7 demonstrated, the regulatory landscape for AI's environmental impact is evolving but remains fragmented and incomplete. Governments are the only actors with the authority to mandate environmental disclosure, set binding efficiency standards, impose siting restrictions, and enforce compliance. The most impactful near-term regulatory action available to governments is mandatory environmental disclosure for data centers and AI systems — including energy consumption, carbon emissions, water use, and e-waste — disaggregated by AI workload.

The EU Energy Efficiency Directive provides the most advanced model, but even it falls short of AI-specific disaggregation. Germany's Energy Efficiency Act (EnEfG) goes further by mandating a 50% renewable electricity share for data centers (rising to 100% by 2027), energy reuse obligations, and binding PUE targets — demonstrating that stringent regulation is both feasible and compatible with continued data center investment. China's NDRC guidelines requiring green power credit purchases for data centers in designated computing regions represent another approach — integrating energy policy with industrial planning for AI infrastructure.¹⁹⁷

Governments must also address the structural siloing of AI governance and environmental governance identified in Section 7. This requires either embedding environmental provisions within AI legislation (as the European Parliament originally proposed but the final EU AI Act largely did not deliver) or creating explicit coordination mechanisms between AI regulators and environmental agencies.

Public Procurement

When governments purchase AI services — for healthcare diagnostics, weather forecasting, tax administration, defense logistics, urban planning, or climate modelling — they can impose environmental performance criteria on suppliers. This creates

¹⁹⁷ See Section 7 sources.

market incentives for sustainable AI without requiring new legislation.

The UNEP Sustainable Procurement Guidelines for Data Centers and Servers (June 2025) provide the first international template for environmentally responsible procurement of digital infrastructure. Governments can go further by requiring suppliers to disclose the lifecycle environmental impact of the AI services they provide, use renewable energy for computation, and meet minimum efficiency standards (PUE, WUE). Public procurement represents a powerful lever because government spending on AI is large and growing — yet environmental criteria are rarely included in procurement specifications.¹⁹⁸

Public Research Investment

Government-funded research can accelerate Green AI by directing investment towards energy-efficient model architectures, sustainable hardware design, and AI applications for environmental monitoring. Public research institutions have produced many of the foundational studies on AI's environmental impact (Strubell et al., 2019; Schwartz et al., 2020; Luccioni et al., 2023) and host key open-source measurement tools. Sustained public funding for this research ecosystem is essential — particularly as

the research itself generates insights that the private sector, which has economic incentives to downplay environmental costs, may be reluctant to produce independently.

Governments can also impose environmental conditions on public grants for AI research — requiring funded projects to report energy consumption and carbon emissions, as the Green AI literature has advocated since 2019.

Direct Environmental Constraints

In exceptional cases, governments have imposed direct constraints on AI infrastructure.

Ireland's decision to pause new data center construction near Dublin until 2028 — driven by data centers' consumption of over 20% of national electricity — demonstrates that governments are willing to prioritize energy and water security over data center growth when the two conflict. Singapore's temporary moratorium (2019–2022) and China's "East Data, West Computing" geographic redistribution represents variations on this approach. These precedents are instructive: they show that environmental limits on AI infrastructure are politically feasible and can be designed to redirect rather than simply block investment.

¹⁹⁸ UNEP, Sustainable Procurement Guidelines, June 2025.

<https://www.unep.org/technical-highlight/unepreleases-guidelines-curb-environmental-impact-data-centres>

8.2 Role of International Organizations

International organizations provide the normative frameworks, technical standards, convening power, and capacity-building resources needed for global coordination. No single organization covers the full AI–environment nexus, but their collective contribution is creating the architecture for international governance.

ITU (International Telecommunication Union)

The ITU plays a central role through three complementary mechanisms. The **AI for Good** platform provides the primary international venue for multistakeholder dialogue on AI applications for sustainable development, including environmental applications. The **Green Digital Action** initiative and its Sustainable AI working group address the environmental footprint of digital infrastructure. And **ITU-T L.1801** (February 2026) — the first international standard specifically for assessing the environmental impact of AI systems — provides the technical foundation for harmonized measurement. The ITU also co-founded the Coalition for Environmentally Sustainable AI in February 2025, together with UNEP and France, bringing together over 100 partners.¹⁹⁹

UNEP (United Nations Environment Programme)

UNEP leads on environmental policy integration for AI. Its AI End-to-End issue note (September 2024) provided the first comprehensive UN assessment of AI’s full lifecycle environmental impact. The Sustainable Procurement Guidelines for Data Centers and Servers (June 2025) provide

the first international procurement framework with environmental performance criteria. UNEA Resolution 7/9 on the “Environmental sustainability of artificial intelligence systems” (December 2025) provides the highest-level multilateral environmental mandate for the work described in this paper.²⁰⁰

UNFCCC

The UNFCCC Technology Mechanism has directly connected AI development to climate action through the **#AI4ClimateAction Initiative** (launched COP28, 2023), which explores AI’s role as an enabler of climate solutions with particular attention to least developed countries and small island developing states. At COP30 in Belém (November 2025), the UNFCCC presented the AI for Climate Action Award and announced the launch of the **AI Climate Institute (AICI)**, designed to equip developing-country institutions with skills to harness AI for climate action. The **Baku-to-Belém Roadmap** explicitly recognizes AI among the technological transformations shaping climate action.²⁰¹

CBD (Convention on Biological Diversity)

The CBD Secretariat is integrating AI into biodiversity governance, including through the AI for Biodiversity series with ITU’s AI for Good platform. The establishment of the Cali Fund for Digital Sequence Information at COP16 (2024) — requiring companies using genomic data to contribute financially to biodiversity conservation — represents the most concrete regulatory linkage between AI-era technologies and biodiversity governance to date. COP17 (2026) will review implementation and potentially strengthen AI-related provisions.

¹⁹⁹ ITU-T L.1801. <https://www.itu.int/epublications/publication/itu-t-l-1801-2026-02>; Coalition for Sustainable AI. <https://www.unep.org/news-and-stories/press-release/new-coalition-aims-put-artificial-intelligence-moresustainable-path>

²⁰⁰ UNEA Resolution 7/9. <https://docs.un.org/ru/UNEP/EA.7/Res.9>

²⁰¹ UNFCCC Technology Mechanism. <https://unfccc.int/tclear/tec/AI4climate.html>

UNCCD

The UNCCD is developing the International Drought Resilience Observatory (IDRO) — the first global AI-driven platform for proactive drought management — with a full launch expected at COP17 in Mongolia in 2026. As discussed in Section 5, the IDRO represents a concrete, multilaterally developed AI tool designed for accessibility across development contexts.

UNESCO

UNESCO adopted the Recommendation on the Ethics of Artificial Intelligence in November 2021 — the first global standard-setting instrument for AI ethics, endorsed by all 193 Member States. The Recommendation establishes four foundational values, including “Environment and Ecosystem Flourishing,” and calls for AI systems to be assessed against sustainability goals including the SDGs. AI technologies should be assessed against their impact on sustainability, and all actors involved in the AI lifecycle must comply with laws and standards designed for environmental and ecosystem protection. While non-binding, the Recommendation carries significant political weight and has influenced national AI strategies in multiple jurisdictions.²⁰²

OECD and GPAI

The OECD published the first intergovernmental measurement framework for AI environmental impacts (2022), calling for a “data revolution” in AI environmental accounting.

The OECD AI Principles provide a widely endorsed governance framework. The Global Partnership on AI (GPAI), hosted by the OECD, produced the “Biodiversity and AI: Opportunities & Recommendations for Action” report (November 2022), identifying priority areas for responsible AI deployment in ecological contexts.²⁰³

IGF Dynamic Coalition on Environment

The Internet Governance Forum’s Dynamic Coalition on Environment — the institutional home of this white paper — provides a unique multistakeholder platform for sustained dialogue on AI and environmental governance. Unlike specialized agencies that address either AI or environment, the Dynamic Coalition operates at the intersection, connecting civil society, governments, the private sector, and international organizations in ongoing deliberation. The DC-Environment’s contribution to the Global Dialogue on AI Governance and the AI for Good Global Summit represents this integrative role in practice.

Across these bodies, international organizations can coordinate standards, build capacity, support Global South access to sustainable AI infrastructure, and convene governments, industry, and civil society around shared measurement frameworks.

The Global Dialogue on AI Governance, AI for Good, WSIS+20, the IGF Dynamic Coalition on Environment, and WAIC can jointly help place AI’s environmental dimension on the global governance agenda.

²⁰² UNESCO, Recommendation on the Ethics of AI. <https://www.unesco.org/en/artificialintelligence/recommendation-ethics>

²⁰³ OECD, 2022. <https://www.oecd.org/en/publications/7babf571-en.html>;

GPAI. <https://wp.oecd.ai/app/uploads/2025/05/biodiversity-and-AI-opportunities-recommendations-for-action.pdf>

8.3 Role of Private Sector

The private sector — encompassing AI developers, hardware manufacturers, cloud infrastructure providers, and enterprise AI users — is the primary driver of AI's environmental footprint and therefore the primary actor whose behavior must change. The evidence reviewed in Sections 2 and 3 demonstrates that corporate decisions on model architecture, hardware procurement, energy sourcing, cooling technology, and geographic siting determine AI's environmental impact more directly than any other factor.

Private-sector actors should disclose AI-related energy and water use, improve model and infrastructure efficiency, expand renewable electricity use, support open environmental datasets, and publish credible assessments of AI-enabled environmental benefits. Corporate AI-for-environment claims should be measurable and independently verifiable. A meaningful distinction must be drawn between AI applications that genuinely reduce environmental pressure and applications that merely optimize business operations while shifting environmental costs elsewhere.

AI Developers

Companies that train and deploy AI models — including Google, Microsoft, Meta, OpenAI, Anthropic, Baidu, Alibaba, Sber, Yandex, and others — control the design decisions that determine the energy and carbon intensity of each model. Section 2.2 demonstrated that these choices can vary AI's carbon footprint by factors of 3x–20x: GLaM consumed only one-third the energy of GPT-3 while achieving better performance; DeepSeek reduced training costs by a factor of 20; BLOOM, trained on French nuclear power, produced a fraction of the emissions it would have on a coal-dependent grid.

Developers also control transparency. The decision to publish lifecycle environmental data — as Hugging Face did for BLOOM (JMLR, 2023) and Mistral AI did for Mistral Large 2 (July 2025, first ISO 14040-compliant commercial LCA) — or not (as

most companies currently choose) is a corporate decision with direct governance implications. Transparency enables comparison, benchmarking, and accountability; its absence prevents all three.

The shift from training-dominated to inference-dominated resource consumption (65% inference vs 35% training per Meta's telemetry) means that developers' choices about how models are deployed, optimised, and served matter more for total environmental impact than how they are trained. The Mistral study's publication of per-query inference data (1.14 g CO₂e per 400-token response) is the first data point of its kind — it should become standard practice, not an exception.

Hardware Manufacturers

Hardware manufacturers — primarily NVIDIA, AMD, Intel, and Huawei for AI accelerators — hold the keys to Scope 3 (supply chain) transparency. As Section 2.1 documented, no major GPU manufacturer currently publishes environmental product declarations (EPDs) for AI accelerators. Mistral AI noted explicitly that “a reliable life-cycle inventory of GPUs is yet to be made, as their embodied impacts had to be approximated but account for a significant portion of total impacts.”

The Google TPU LCA (Schneider et al., 2025), published as a peer-reviewed study covering five generations of TPU hardware with first-party manufacturing data, demonstrates that cradle-to-grave lifecycle assessment of AI accelerators is technically feasible. The barrier is willingness, not capability. Hardware manufacturers are uniquely positioned — and arguably uniquely obligated — to close this data gap, because no other actor has access to the manufacturing, materials, and supply chain data needed for accurate embodied carbon and resource depletion calculations.²⁰⁴

Hardware design choices also determine the longevity and recyclability of AI infrastructure.

²⁰⁴ Schneider, I. et al. “Life-Cycle Emissions of AI Hardware,” Google, 2025. <https://arxiv.org/pdf/2502.01671>

AI hardware refresh cycles in data centers are typically 3–5 years, driven by rapid performance improvements. Designing for repairability, component reuse, and material recovery could significantly reduce AI's e-waste footprint — an area where regulatory incentives (such as extended producer responsibility schemes) may be needed to complement voluntary corporate action.

Cloud Infrastructure Providers

Cloud providers — AWS, Azure, Google Cloud, Alibaba Cloud — operate the physical infrastructure and therefore control energy sourcing, cooling technology, facility design, and geographic siting. Their corporate sustainability commitments have set important benchmarks: Google's 24/7 carbon-free energy target, Microsoft's carbon-negative and water-positive goals, Amazon's Climate Pledge, Alibaba Cloud's 100% clean energy by 2030.

However, Section 3.3 documented a persistent gap between commitment and outcome.

Microsoft's emissions rose 29% between 2020 and 2024; Google's rose 48% compared to 2019 — in both cases primarily driven by AI infrastructure expansion. Microsoft's launch of a zerowater cooling data center design in August 2024 and its 80%

reduction in water intensity over successive data center generations demonstrate that technological solutions exist. The challenge is that infrastructure growth is outpacing efficiency gains — the rebound effect identified by Kaack et al. (Nature Climate Change, 2022).

Sustainability Innovators

Companies across the AI value chain are demonstrating that environmental innovation is possible and can create competitive advantage. ByteDance deployed immersion cooling through its Volcano Engine platform, achieving PUE of 1.05 and cooling energy savings of up to 50%. Yandex's data center in Mäntsälä, Finland, reuses waste heat to heat up to 50% of local buildings. Apple's Daisy robot disassembles iPhones for component recycling. Rosatom has developed rare metal recycling capabilities. Sber and AIRI developed Eco2AI and Eco4cast — open-source tools for carbon monitoring and optimisation that demonstrate BRICS-originated innovation in sustainable AI.

These innovations are valuable but remain scattered and voluntary. The policy challenge — addressed in Section 7 — is to create frameworks that make environmental innovation the norm rather than the exception.

8.4 Role of Civil Society

Civil society — encompassing academic researchers, indigenous peoples and local communities, advocacy organizations, and citizens — plays roles that cannot be substituted by governments or corporations: independent knowledge production, accountability monitoring, and ensuring that governance reflects diverse perspectives and protects the rights of affected communities.

Civil society organizations can help monitor environmental justice, data rights, community impacts and the distribution of AI's environmental costs and benefits. Local communities, including indigenous peoples and vulnerable groups, should have channels to participate when AI systems affect land, biodiversity, water allocation or environmental enforcement. AI-based environmental monitoring should not become a purely technocratic process detached from public participation.

Academic and Research Institutions

The foundational evidence base for AI environmental governance was produced overwhelmingly by academic researchers and publicly funded institutions. The seminal papers that launched the field — Strubell, Ganesh, and McCallum (ACL, 2019) on energy costs of NLP training; Schwartz, Dodge, Smith, and Etzioni (Communications of the ACM, 2020) on the Red AI/Green AI distinction; and Luccioni, Viguier, and Ligozat (JMLR, 2023) on the BLOOM lifecycle assessment — were academic contributions, not corporate publications.

Academic institutions also developed and maintain key open-source measurement tools: CodeCarbon (Mila, BCG GAMMA, Haverford College), Eco2AI (Sber/AIRI), and the broader research infrastructure that enables independent verification of corporate environmental claims.

This independence is critical. Corporate AI companies have economic incentives to present their environmental impact in the most favorable light; academic researchers, while not immune to bias, operate under norms of reproducibility, peer review, and transparency that provide important checks. Sustained public funding for independent AI environmental research — including lifecycle assessment, measurement methodology development, and policy analysis — is essential for maintaining the evidence base on which governance decisions depend.

Indigenous Peoples and Local Communities

Indigenous peoples and local communities are both affected by AI infrastructure and holders of knowledge relevant to biodiversity conservation and land management. Their role in AI governance extends beyond consultation to rights-holding.

The Convention on Biological Diversity's provisions on traditional knowledge (Article 8(j)) and access and benefit-sharing establish a legal framework for recognizing indigenous rights over biodiversity data. The Cali Fund, established at CBD COP16 in November 2024, allocates at least 50% of its resources to indigenous peoples and local communities, recognizing their role as custodians of biodiversity and their right to benefit from the use of genetic resources — including genomic data analyzed by AI.²⁰⁵

The UNDP's "People-Centric AI for Conserving Biodiversity" report (December 2025) provides a framework for ensuring that AI systems for conservation respect indigenous data sovereignty — the right to control how traditional knowledge and biodiversity data is collected, stored, analyzed,

²⁰⁵ UNEP, "How a groundbreaking agreement could raise billions," November 2024.

<https://www.unep.org/news-and-stories/story/how-groundbreaking-agreement-could-raise-billions-protect-web-life>

and used. This principle has implications for how training datasets are assembled, how AI models are validated in ecological contexts, and how the outputs of AI-driven conservation decisions are governed.²⁰⁶

These principles connect directly to the equity dimension of sustainable AI: the communities least responsible for AI's environmental footprint are often those most affected by its consequences — through mining of raw materials, water diversion for cooling, and land conversion for data center construction.

Communities Affected by AI Infrastructure

In several jurisdictions, communities located near data centers or data center construction sites have raised concerns about energy consumption, water use, noise pollution, visual impact, and the

strain on local infrastructure. In Virginia (USA), which hosts approximately 26% of US data center capacity, community groups have challenged the environmental impact of continued expansion. In the Netherlands, community opposition has influenced data center approval decisions. In Ireland, concerns about data centers' electricity consumption — exceeding 20% of national demand — contributed to the government's decision to pause new construction.

These community voices are essential for ensuring that AI infrastructure development does not impose disproportionate environmental burdens on specific localities. Effective governance requires mechanisms for meaningful community participation in data center siting and approval decisions — including environmental impact assessments, public consultation periods, and the ability to challenge approvals through administrative or judicial review.

²⁰⁶ UNDP, December 2025.

<https://www.undp.org/sites/g/files/zskgke326/files/2025-12/people-centric-ai-forconserving-biodiversity.pdf>

Section 9. Next Steps

9.1 Challenges	85
9.2 Emerging Trends and Future Directions	87
9.3 Recommendations	89



9.1 Challenges

This white paper identifies **six cross-cutting challenges**.

Challenge 1. Data on AI's energy, carbon and water footprint remain incomplete and often non-comparable.

Despite growing attention to AI sustainability, data on AI's energy consumption, greenhouse gas emissions and water use remain incomplete, inconsistent, and often non-comparable across organizations. Reporting practices vary widely, with most companies disclosing aggregated data center metrics rather than AI-specific impacts. While standards such as ITU-T L.1801 and ISO TR 20226:2025 represent important advances in environmental accounting, implementation has not yet started and global reporting obligations are still limited.

A similar measurement challenge exists on the benefits side. There is currently no internationally recognized framework for validating claims regarding AI-enabled environmental outcomes, such as emissions reductions, water savings, or improvements in climate mitigation and adaptation. As a result, environmental benefits are frequently based on self-reported estimates that cannot be systematically compared across jurisdictions or applications. This asymmetry means that the environmental costs of AI are becoming increasingly measurable, while its environmental benefits remain difficult to verify.

Challenge 2. Environmental impacts are geographically uneven and can create equity problems.

The environmental costs and benefits of AI are distributed unevenly across regions and social groups. The extraction of critical minerals required for semiconductors and AI hardware is concentrated in a limited number of countries, while associated environmental degradation often affects local communities with limited participation in AI value creation. Likewise, water-intensive data centers are

increasingly located in regions already experiencing water stress, including parts of the United States, India, China, and the Middle East.

At the same time, AI-generated electronic waste is frequently processed through informal recycling systems in developing countries, exposing vulnerable populations to environmental and health risks. The benefits of AI innovation, by contrast, remain concentrated among a relatively small number of companies and countries possessing advanced computing infrastructure. These asymmetries raise broader questions of environmental justice and equitable distribution of both risks and opportunities associated with AI development.

Equity concerns also extend to data and knowledge representation. The dominance of English-language datasets and the limited inclusion of indigenous and local knowledge systems may reduce the effectiveness of AI-based environmental solutions in some regions while reproducing existing patterns of digital inequality.

Challenge 3. AI infrastructure is expanding faster than regulatory and measurement systems.

The rapid growth of AI infrastructure is occurring faster than the development of governance frameworks capable of measuring and managing its environmental consequences. Data centers, semiconductor fabrication facilities, and supporting energy infrastructure involve long-lived investments whose environmental impacts may persist for decades. Decisions regarding facility location, energy sourcing, cooling technologies, and hardware procurement are being made today, often before comprehensive sustainability requirements have been established.

This challenge is compounded by governance fragmentation. Responsibility for AI governance, environmental regulation, energy policy, and digital infrastructure development is distributed across multiple national and international institutions. As a result, environmental considerations are often

absent from AI governance frameworks, while AI-related impacts remain insufficiently addressed within environmental policy regimes. The pace mismatch between technological deployment and regulatory adaptation creates a narrowing window for proactive governance.

Challenge 4. AI-for-environment applications may generate rebound effects if efficiency gains lower costs and increase consumption.

Many environmental benefits associated with AI arise from efficiency improvements, including optimised energy systems, more efficient industrial processes, precision agriculture, and intelligent resource management. However, such gains do not automatically translate into lower overall resource consumption.

As AI systems become more efficient and computational costs decline, demand for AI services often increases. This phenomenon, commonly known as the rebound effect or Jevons paradox, may partially or completely offset environmental gains achieved through efficiency improvements. Recent trends illustrate this dynamic: hardware efficiency continues to improve and inference costs have fallen dramatically, yet total electricity demand from data centers continues to increase. Managing rebound effects therefore requires governance approaches that focus not only on relative efficiency improvements but also on absolute resource consumption and system-wide environmental outcomes.

Challenge 5. Environmental AI systems require high-quality data and long-term institutional maintenance.

The effectiveness of AI for environmental applications depends fundamentally on the availability of high-quality data, monitoring infrastructure, and institutional capacity. Yet many regions facing the greatest environmental risks also suffer from severe data limitations.

Hydrological monitoring networks, biodiversity observations, meteorological records, and land-

use datasets are often sparse or incomplete in developing countries, reducing the accuracy and reliability of AI-based environmental systems.

These limitations are further amplified by climate change itself. AI models trained on historical observations may struggle to predict unprecedented environmental conditions, including extreme heatwaves, droughts, floods, and ecosystem transformations. Addressing this challenge requires sustained investment not only in data collection and environmental monitoring but also in physics-informed AI approaches, scientific validation frameworks, and long-term institutional maintenance. Environmental AI should be understood as a continuous public infrastructure commitment rather than a one-time technological deployment.

Challenge 6. International cooperation is complicated by digital sovereignty, export controls and uneven access to computing power.

AI development is increasingly shaped by geopolitical competition. Access to advanced semiconductors, cloud infrastructure, foundation models, and high-performance computing resources is concentrated in a limited number of countries and firms. Export controls, technology restrictions, and competing approaches to digital sovereignty complicate international cooperation precisely when global environmental challenges require collective action.

These constraints are particularly significant for developing countries, many of which face the greatest exposure to climate and environmental risks while lacking access to the computational resources necessary to develop and deploy advanced AI systems. The resulting imbalance risks creating a new digital-environmental divide, in which countries most in need of AI-enabled environmental solutions have the least capacity to build or adapt them. Expanding international cooperation through technology transfer, capacity-building initiatives, open scientific collaboration, and multilateral financing mechanisms therefore remains essential for ensuring that the environmental benefits of AI are broadly shared.

9.2 Emerging Trends and Future Directions

Despite the challenges in measuring and governing AI's environmental impacts, a number of emerging technological, institutional, and policy developments suggest that there is the potential for a more sustainable AI ecosystem, provided that governance frameworks evolve in parallel with technological progress. The following section outlines key future directions that are likely to shape the environmental footprint of AI over the coming decade.

Standardized AI Environmental Reporting and Accountability

A foundational trend is the transition from voluntary sustainability disclosures towards standardized environmental reporting for AI systems and infrastructure. International standards such as ITU-T L.1801 and ISO TR 20226:2025 provide the first common frameworks for measuring AI-related energy consumption, greenhouse gas emissions, and resource use. At the same time, open-source measurement tools — including Eco2AI, Eco4cast and many others — have made environmental assessment increasingly accessible. The next phase is likely to involve mandatory disclosure requirements and AI-specific reporting obligations, allowing policymakers and stakeholders to compare environmental performance across organizations and jurisdictions. A key priority will be the development of reporting systems that jointly track energy, water, carbon, and material impacts rather than treating them as separate metrics.

Life-Cycle Assessment for Foundation Models

Environmental assessment is gradually expanding beyond operational electricity consumption towards full life-cycle analysis of AI systems. Recent studies, including Mistral AI's life-cycle assessment, demonstrate that emissions can be measured at the model and query level, while growing attention

is being paid to embodied impacts associated with semiconductor manufacturing, data center construction, cooling infrastructure, and electronic waste generation. Future governance frameworks are likely to require lifecycle assessments for foundation models and major AI infrastructure projects, enabling a more comprehensive understanding of environmental trade-offs across the entire AI value chain.

Carbon- and Water-Aware Workload Scheduling

As AI computing become increasingly resource-intensive, greater attention is being directed towards the spatial and temporal optimisation of computing resources. Carbon-aware scheduling shifts computational tasks towards periods and locations with lower carbon intensity electricity, while emerging water-aware approaches seek to reduce computing activity during periods of local water stress. Such approaches recognize that environmental impacts depend not only on how much computing is performed, but also on where and when it occurs. The integration of environmental indicators into cloud and data center management systems may become a core feature of sustainable AI infrastructure.

Model Compression and Task-Specific Small Models

A growing body of evidence suggests that improvements in AI performance do not necessarily require proportional increases in computational scale. Models such as DeepSeek and GLaM have demonstrated that architectural innovation can significantly reduce energy and training requirements while maintaining competitive performance. Advances in quantization, distillation, pruning, and model compression are making it possible to deploy increasingly capable systems with lower resource demands.

At the same time, the expansion of edge AI—deploying models directly on devices or local infrastructure rather than centralized hyperscale data centers—offers opportunities to reduce data transmission, latency, and infrastructure requirements. Over the longer term, emerging technologies such as neuromorphic computing may further alter the energy profile of AI systems by fundamentally redesigning computing architectures around principles of biological efficiency.

Clean-Power and Computing-Power Coordination

The rapid expansion of AI is increasingly linking digital policy with energy policy. Major technology companies are investing directly in low-carbon electricity generation, including renewable energy, nuclear power, and emerging clean-energy technologies. The challenge is evolving from securing sufficient electricity to coordinating computing growth with broader energy-system planning.

Future approaches are likely to involve closer integration between AI infrastructure development and national energy strategies, including grid planning, renewable deployment, storage capacity, and demand management. China's experience illustrates this trend, where the location of computing hubs is increasingly coordinated with renewable energy resources and broader industrial policy objectives. Such coordination will become increasingly important as AI emerges as a major source of electricity demand.

AI and Digital Twins for Environmental Monitoring and Governance

AI is increasingly being deployed not only as a subject of governance but also as a tool for governance itself. Applications such as biodiversity monitoring platforms, flood forecasting systems, drought resilience tools, and nature-risk assessment frameworks demonstrate how AI can strengthen environmental decision-making. Satellite imagery, remote sensing, and machine learning are enabling near-real-time monitoring of emissions, deforestation, biodiversity loss, land degradation, and water resources at unprecedented scales.

The emergence of digital twins — virtual representations of environmental systems that integrate real-time data streams with predictive modelling — may further transform environmental governance. Such systems could support scenario analysis, policy evaluation, and adaptive management across climate, biodiversity, land, and water domains.

Open Environmental Data and Capacity Building for the Global South

The effectiveness of environmental AI depends fundamentally on access to high-quality data.

Yet many regions most vulnerable to climate change, biodiversity loss, and land degradation continue to face significant data gaps. Expanding open environmental datasets, improving monitoring infrastructure, and strengthening technical capacity in developing countries are therefore emerging as critical priorities.

Without targeted investment, AI risks reinforcing existing inequalities between data-rich and data-poor regions. Future international cooperation efforts are likely to focus increasingly on environmental data sharing, scientific collaboration, digital public goods, and equitable access to computing resources. Ensuring that countries in the Global South can participate not only as users but also as developers of environmental AI systems will be essential for achieving globally inclusive sustainability outcomes.

Taken together, these trends suggest that sustainable AI governance is evolving beyond narrow concerns about model performance or carbon emissions. Future governance frameworks are likely to integrate AI regulation, environmental reporting, infrastructure planning, energy policy, ecological monitoring, and resource management into a more holistic governance model. Experiences from China and other jurisdictions indicate that the most effective approaches will increasingly treat AI systems as components of broader socio-technical and ecological systems rather than as isolated digital technologies.

9.3 Recommendations

Based on the analysis presented across all sections of this paper, the Dynamic Coalition on Environment offers the following recommendations, organized by actor and designed to address the specific gaps identified in Sections 7 and 8.

For Governments

Adopt mandatory environmental disclosure for AI systems and data centers, including energy consumption, carbon emissions, water use, and e-waste, disaggregated by AI workload. The EU Energy Efficiency Directive provides a starting point; Germany’s Energy Efficiency Act provides a more ambitious model. Disclosure should extend to both training and inference phases and cover the full lifecycle.

Integrate environmental criteria into AI governance frameworks. The current siloing of AI policy and environmental policy — exemplified by the EU AI Act’s dilution of environmental provisions during trilogue negotiations — must be addressed. China’s Global AI Governance Action Plan, which explicitly advocates joint AI energy and water efficiency standards, offers a model for this integration. AI legislation should require environmental impact assessment for high-resource AI systems, and environmental legislation should include AI-specific provisions.

Include environmental performance criteria in public procurement of AI services, using the UNEP Sustainable Procurement Guidelines (June 2025) as a template. When governments are major purchasers of AI, procurement requirements can drive market-wide change without requiring new regulation.

Establish or strengthen data center siting policies that account for regional energy mix, water stress, ecological sensitivity, and cumulative impact. China’s “East Data, West Computing” program demonstrates that strategic spatial coordination of computing demand and infrastructure development can improve energy efficiency while reducing pressure on resource-constrained regions.

Invest in Green AI research and capacity building, including funding for energy-efficient model architectures, sustainable hardware, AI applications for environmental monitoring, and training programs for developing-country researchers and practitioners. The UNFCCC AI Climate Institute (AICI) provides a model.

For AI Developers and Technology Companies

Publish full lifecycle environmental assessments of AI models, following the methodologies demonstrated by Hugging Face (BLOOM, JMLR 2023) and Mistral AI (Mistral Large 2, July 2025, first ISO 14040-compliant commercial LCA). Include training, inference, and infrastructure impacts across multiple impact categories (carbon, water, resource depletion). Transparency is a precondition for accountability.

Adopt and report against standardized environmental metrics (PUE, WUE, CUE, Scope 1/2/3 emissions) with AI-specific disaggregation. Metrics should be assessed jointly to provide a holistic view of environmental performance and avoid burden shifting between energy, water, and carbon impacts. Follow emerging standards including ITU-T L.1801 and ISO TR 20226:2025.

Invest in efficiency-oriented research and deployment — including model compression, architecture optimisation, inference efficiency, and temporal/geographic compute scheduling (as demonstrated by AIRI’s Eco4cast). Treat environmental performance as a primary evaluation criterion alongside accuracy and cost.

Implement water stewardship programs that go beyond aggregate “water-positive” commitments to include site-specific impact assessment, AI workload disaggregation, and engagement with local communities in water-stressed regions.

For Hardware Manufacturers

Publish environmental product declarations (EPDs) for AI accelerators, including embodied carbon, water, critical raw materials, and material composition data.

Design for longevity, repairability, and recyclability to reduce the e-waste impact of rapid hardware refresh cycles. Consider extended producer responsibility models and participate in circular economy initiatives.

For International Organizations

Accelerate the adoption of AI-specific environmental standards (ITU-T L.1801, ISO TR 20226:2025) in national regulatory frameworks, providing technical assistance and model legislation to facilitate adoption.

Establish a dedicated international forum — or strengthen existing ones such as the IGF Dynamic Coalition on Environment — where AI governance and environmental governance systematically converge. The current fragmentation across ITU, UNEP, OECD, UNESCO, and national bodies creates gaps and inconsistencies that a coordinating mechanism could address.

Support capacity building in developing countries for both AI development and environmental monitoring. Ensure that AI tools for environmental management are accessible, affordable, and relevant to developing-country contexts — not designed

exclusively for data-rich, resource-rich environments. The UNFCCC AI Climate Institute and the IDRO provide models for this approach.

Strengthen the evidence base through continued support for independent, peer-reviewed research on AI's environmental impacts and benefits, including lifecycle assessment, measurement methodology, and comparative policy analysis.

For Civil Society and Research Institutions

Continue developing and maintaining open-source measurement tools and transparent datasets for AI environmental assessment. The independence and accessibility of tools like Eco2AI, Eco4cast and others, are essential for accountability.

Ensure that indigenous peoples and local communities are included as rights-holders in AI governance processes, with meaningful participation in decisions about data collection, AI deployment in conservation contexts, and the governance of digital sequence information. The CBD's Article 8(j), the Cali Fund, and the UNDP "People-Centric AI" framework provide the normative basis; implementation requires sustained civil society engagement.

Monitor and publicly evaluate corporate environmental commitments, providing independent accountability. Academic research that verifies or challenges corporate sustainability claims should be supported and protected.

naturepeople.ru



fudan.edu.cn



pahléindia.org



intgovforum.org

