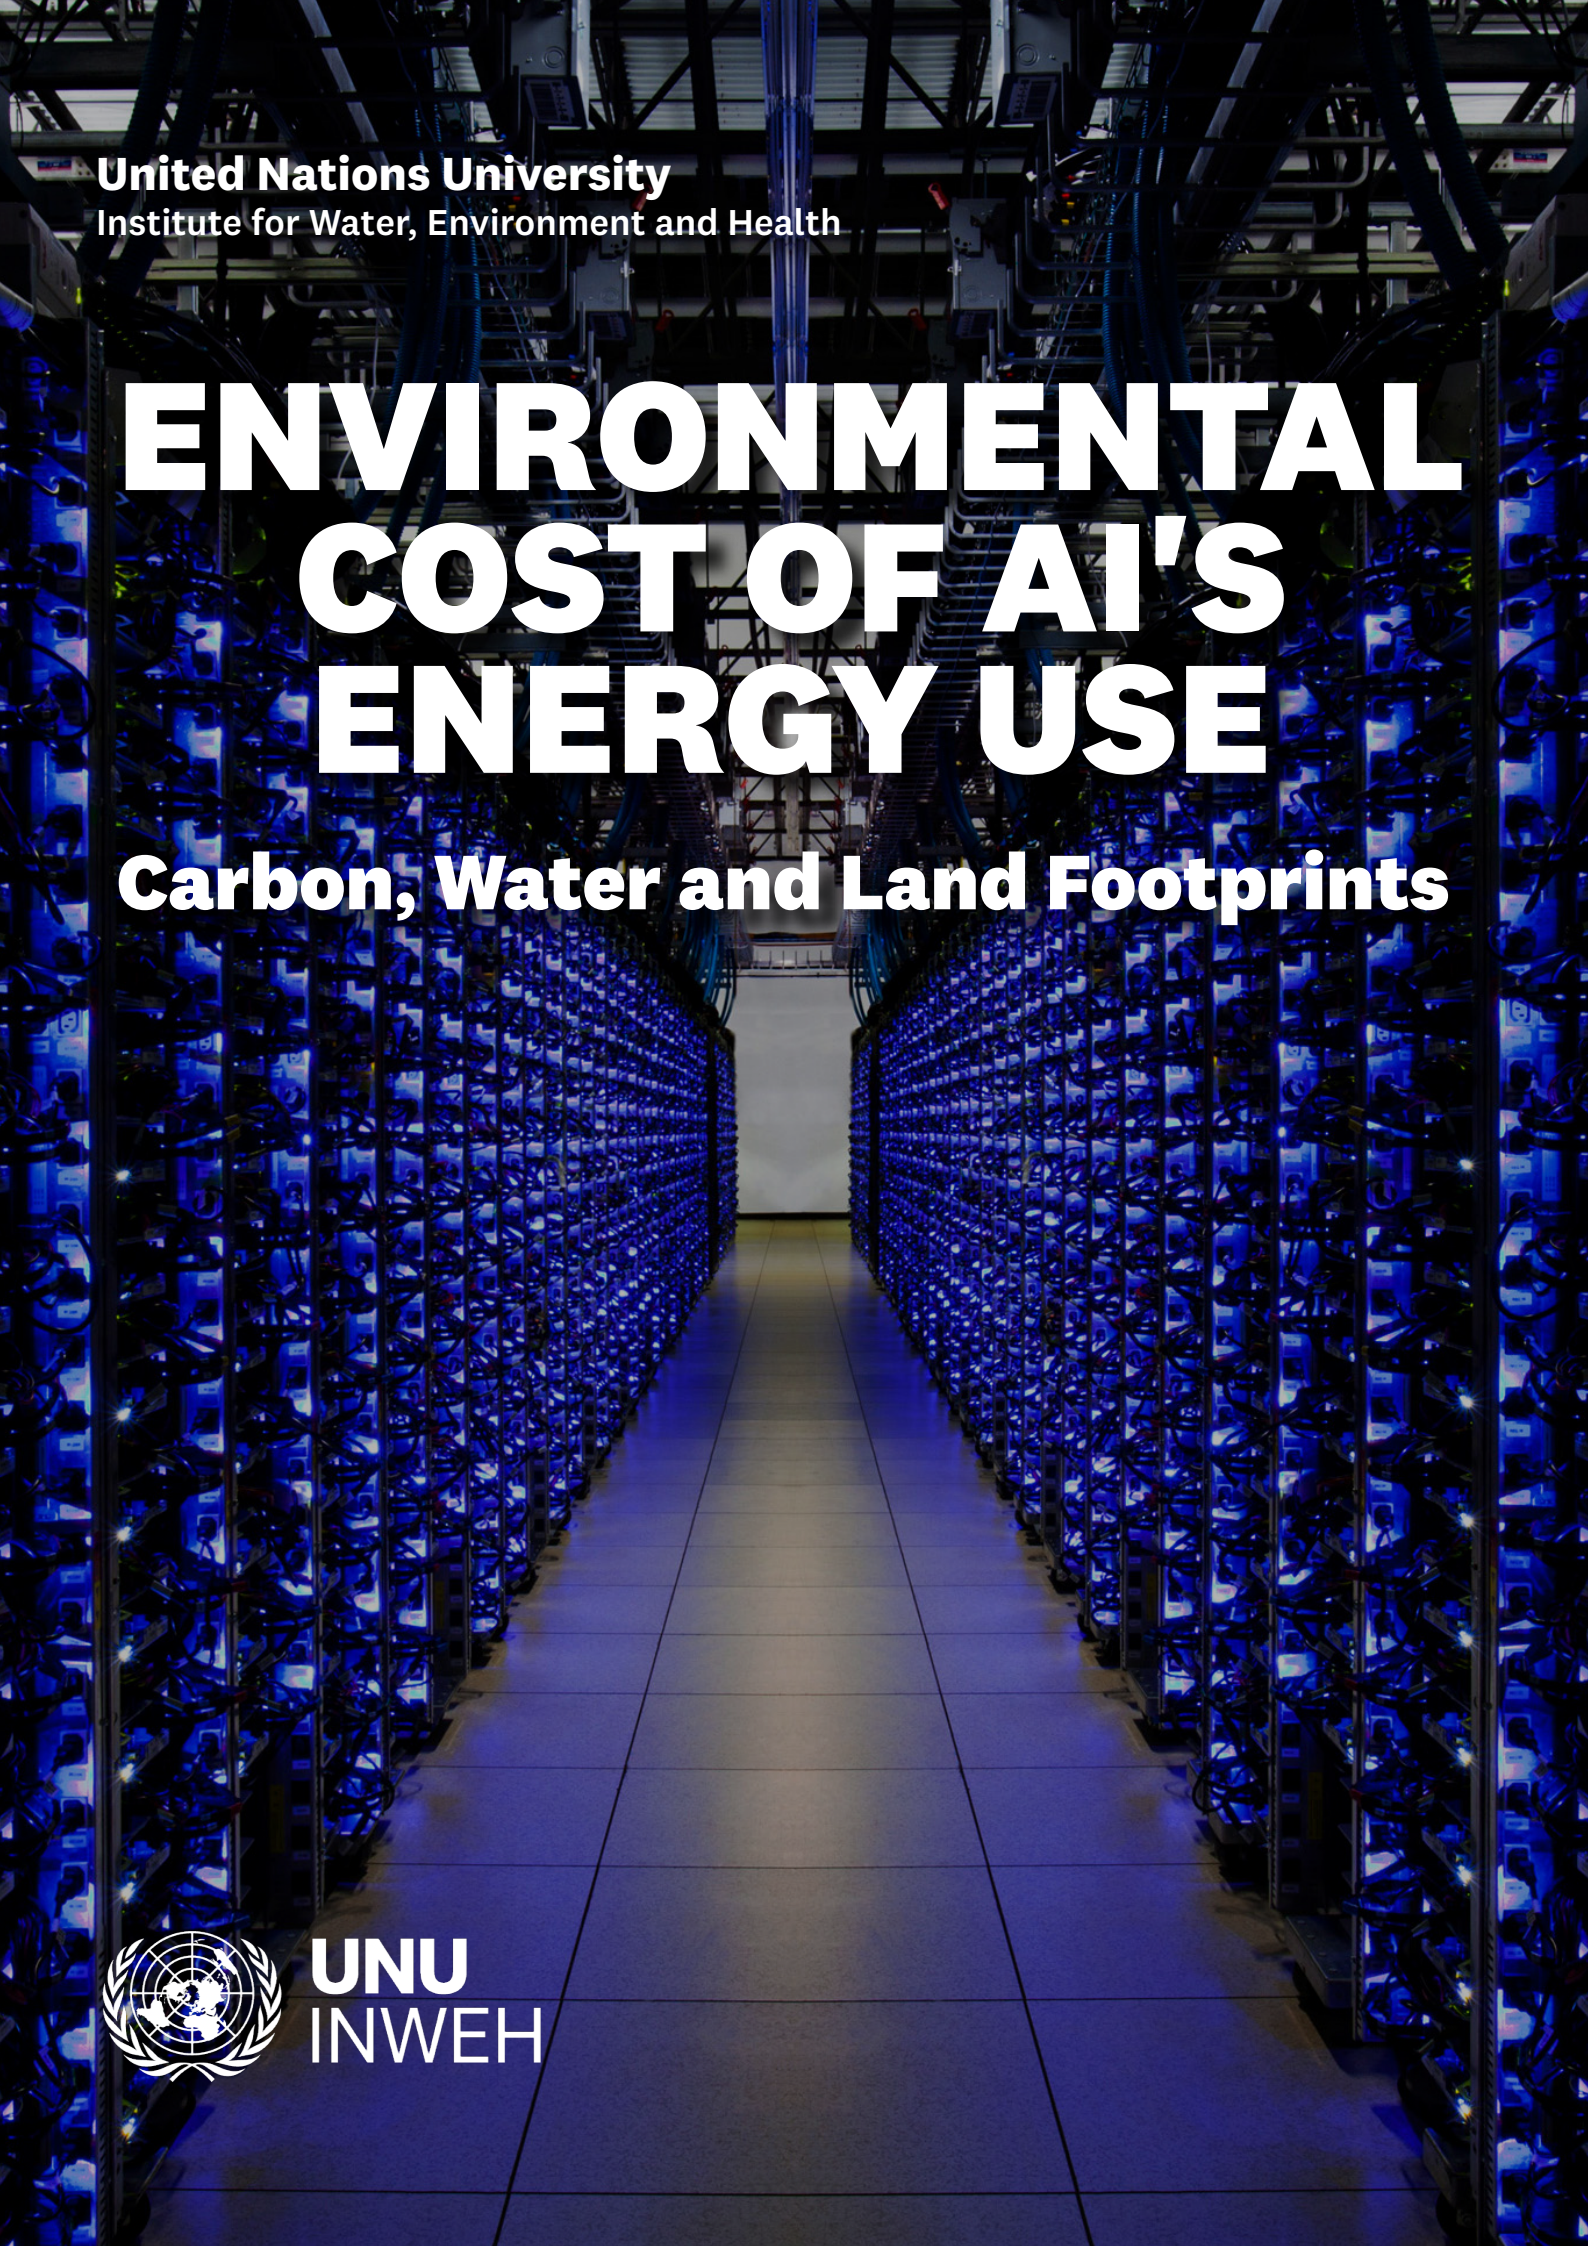

United Nations University
Institute for Water, Environment and Health

ENVIRONMENTAL COST OF AI'S ENERGY USE

Carbon, Water and Land Footprints

UNU INWEH

Acknowledgment
UNU-INWEH gratefully acknowledges its host, the Government of Canada, and ongoing financial support from Global Affairs Canada.

Disclaimer

UNU-INWEH enjoys the generous support of Global Affairs Canada. As a United Nations entity, UNU-INWEH operates independently. The content of UNU-INWEH publications is grounded in scientific evidence and is not intended to be influenced by, nor does it necessarily reflect, the views of the Government of Canada, as host government, or any other government, donor, or funding partner.

About UNU-INWEH

The United Nations University Institute for Water, Environment and Health (UNU-INWEH) is one of 13 institutions comprising the United Nations University (UNU), the academic arm of the United Nations. Established in 1996 through an agreement with the Government of Canada, UNU-INWEH is headquartered in the City of Richmond Hill, Ontario.

UNU-INWEH specializes in addressing critical global security and development challenges at the intersection of water, environment, and health. Through research, capacity development, policy engagement, and knowledge dissemination, the institute bridges the gap between scientific evidence and the practical needs of policymakers and UN member states, with particular attention to low- and middle-income countries. By collaborating with a diverse array of partners—including UN agencies, governments, academia, the private sector, and civil society—UNU-INWEH develops solutions that advance human security, resilience, and sustainability worldwide.

Environmental Cost of AI's Energy Use
Carbon, Water and Land Footprints

United Nations University Institute for Water, Environment and Health (UNU-INWEH)
Richmond Hill, Ontario, Canada, 2026

Authors: Miriam Aczel, Sanaz Chamanara, Mir Matin, Aria Farsi, Tshilidzi Marwala, Kaveh Madani

Contributors: Luying Wang, Jenna Elshurafa

How to cite: Aczel M., Chamanara S., Matin M., Farsi A., Marwala T., Madani K. (2026) Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints, United Nations University Institute for Water, Environment and Health (UNU-INWEH), Richmond Hill, Ontario, Canada, doi: 10.53328/INR26RMA002

ISBN: 978-92-808-6141-9
DOI: 10.53328/INR26RMA002
Download at: inweh.unu.edu/publications

Front cover: A row of servers in Google's Douglas County, Georgia, data center. Photo from Google Gallery

@30
30TH ANNIVERSARY
UNU-INWEH
2026

CONTENTS

CONTENTS **05**
FIGURES **06**
ABBREVIATIONS **07**
KEY POINTS **08**
EXECUTIVE SUMMARY **11**
1. AI ACROSS THE GLOBE **14**
1.1 A Technological Explosion 17
1.2 The Expanding Reach of AI Across Sectors 18
1.3 The Global Market for AI 18
1.4 Impacts on the Global Labor Market 19
1.5 Growing Concerns about AI 20
1.6 Environmental Dimensions of AI 22
1.7 A Just and Sustainable AI Future 24
2. AI's Growing Energy Demand **25**
2.1 Energy Intensity of AI Training 26
2.2 AI's Energy Demand and its Environmental Footprints 26
2.3 AI's Training Footprints 26
2.4 Data Centers and AI's Expanding Energy Demand 30
2.5 AI's Contributions to Data Centers' Energy Use 33
2.6 World's Top Data Center Hosts and their Energy Mixes 34
2.7 Local Costs and Distant Benefits 36
3. AI in Use: Tools and Tasks **38**
3.1 Two Forces Shaping AI's Operational Footprint 39
3.2 AI Tasks and their Footprints 40
3.3 AI Use at Scale 41
3.4 Conventional Search Versus Enhanced Search 42
3.5 Efficiency Improvement and Rebound Effects 43
3.6 Behavioral Energy Costs: Model Choice and User Prompts Matter 43
3.7 Toward Sustainable Use and Scale 45
4. The Way Forward **46**
4.1 Guiding Principles for a Responsible AI Ecosystem 47
4.2 From Measurement to Action 48
4.3 Roles and Responsibilities 49
4.4 Putting Principles into Practice 52
4.5 Conclusion 52
REFERENCES **53**

FIGURES

Figure 1. The evolution of artificial intelligence 16
Figure 2. Industrial revolutions from mechanization to intelligent systems 17
Figure 3. Key breakthroughs in deep learning 23
Figure 4. Global electricity supply by source 27
Figure 5. Electricity supply mixes in major data center host locations 28
Figure 6. Global electricity footprints 29
Figure 7. Footprints of electricity in the world's top 20 data center hubs 31
Figure 8. Global distribution of data center locations 32
Figure 9. Global electricity footprint intensities across the world's top 20 data center hubs relative to the global average 37
Figure 10. Energy demand in AI arises from two distinct phases 39
Figure 11. Energy per query for representative text and image generation tasks 41
Figure 12. How video generation settings shape energy demand 44
Figure 13. Six guiding principles for operationalizing responsible AI 48

ABBREVIATIONS

4IR	Fourth Industrial Revolution
AI	Artificial Intelligence
CO_2e	Carbon dioxide equivalent
CPU	Central Processing Unit
GPT	Generative Pre-trained Transformer
GPU	Graphics Processing Unit
GW	Gigawatt, 1,000,000,000 W
GWh	Gigawatt-hour, 1,000 MWh
HPC	High-Performance Computing
IoT	Internet of Things
kW	Kilowatt, 1,000 W
kWh	Kilowatt-hour, 1,000 Wh
LLM	Large Language Model
MoE	Mixture-of-Experts
T2V	Text-to-Video
TPU	Tensor Processing Unit
TWh	Terawatt-hour, 1,000 GWh
USD	United States Dollar
W	Watt, power
Wh	Watt-hour, energy

Key Points in Brief

AI's growth

The global AI market is expanding rapidly, projected to grow from USD 189 billion in 2023 to nearly USD 5 trillion by 2033.

This would represent roughly a 25-fold increase in global AI market size over a decade.

Global AI expenditure is projected to exceed USD 2.5 trillion in 2026.

Generative AI accounted for over 20% of the global AI market in 2026, and is projected to reach 40% by 2030.

Corporate AI investment exceeded USD 580 billion in 2025, while generative AI alone attracted nearly USD 34 billion in private investment.

78% of organizations reported using AI in their work in 2024, and 40% of employers expect to reduce their workforce where AI can automate tasks.

About 60% of jobs in advanced economies incorporate AI, versus 26% in low-income countries.

Nearly half of the world's data centers are in the United States.

Only 16% of countries host AI-specialized cloud compute, and 90% of that capacity is concentrated in just 2 countries (the United States and China).

Training footprints

The training of frontier models demands immense energy. GPT-4 likely consumed 50 to 70 GWh of electricity over 100 days, roughly 40–55 times more than GPT-3 (1.287 GWh over a 34-day period).

GPT-4's training energy is equivalent to the annual residential electricity consumption of over 460,000 people in Sub-Saharan Africa.

GPT-4's training carbon footprint of 25,000 tonnes of CO_2e would require the sequestration capacity of 420,000 tree seedlings grown for 10 years, or about equal to the number of trees in 105 Hyde Parks in London.

The water footprint associated with training GPT-4 was about 600 million liters, enough to meet the minimum annual domestic water needs of 81,000 people in Sub-Saharan Africa, or to fill 237 Olympic-sized pools.

Projections for models like GPT-5 suggest training electricity requirements of 100 GWh, equivalent to the annual residential electricity usage of 770,000

people in Sub-Saharan Africa.

Training models like GPT-5 is estimated to have a carbon footprint of 42,000 tonnes of CO_2e, requiring 700,000 tree seedlings (about equal to the number of trees in 40 Central Parks in New York or 155 times the trees in Toronto's High Park over 10 years to offset).

The water footprint of GPT-5 training is estimated at 1 billion liters, enough to meet the annual domestic water needs of more than 135,000 people in Sub-Saharan Africa.

The land footprints of training GPT-4 and GPT-5 are estimated at roughly 0.9 km^2 (126 football fields) and 1.5 km^2 (210 football fields), respectively.

Inference and "AI in use"

Following its 2022 launch, ChatGPT surpassed 1 million users in 5 days and 100 million users in under 2 months. Currently, ChatGPT processes an estimated 2.5 billion prompts per day.

China's DeepSeek, launched in January 2025, attracted more than 20 million daily active users within three weeks, and had about 125 million monthly active users by mid-2025.

While training is highly resource-intensive, the continuous inference phase used to generate responses for billions of interactions is estimated to account for 80% to 90% of total AI energy use.

A typical ChatGPT-style text query is about 200 times more energy-intensive than text classification (such as spam filtering).

Generating a typical AI image requires 2.9 Wh, making it 60 times more demanding than a short text answer and 1,450 times that of text classification.

Video generation represents the most energy-intensive frontier with high-resolution long clips on large models drawing more than 415 Wh per clip, meaning a single short AI video can draw as much electricity as 200,000 spam classifications.

The energy required to generate a typical AI image is enough to power a 10-watt LED bulb for 17 minutes, and the energy required for a high-complexity AI video is sufficient to run that same bulb for 42 hours. Similarly, the electricity-associated water footprint is about two tablespoons (29 mL) for a single image, but jumps to 4.1 liters for a complex video—almost equivalent to a two-day drinking water need for one person.

Data centers as infrastructure

In 2025, global data centers were estimated to consume 448 TWh of electricity. If data centers' electricity use were considered a country, it would have ranked 11th globally by electricity consumption.

The energy consumed by data centers in 2025 was enough to supply the annual residential electricity needs of the entire population of Sub-Saharan Africa, 1.3 billion people, for 2.6 years.

Data centers' electricity use in 2025 had a carbon footprint of 189 million tonnes of CO_2e, which would require 3.2 billion tree seedlings grown over 10 years to offset, roughly the total number of trees in the entire United Kingdom.

Data centers' 2025 electricity consumption had a water footprint of 4.5 trillion liters of water—enough to fill 1.8 million Olympic-sized pools or meet the annual basic domestic water needs of over 600 million people in Sub-Saharan Africa.

The land footprint of 2025 data centers' electricity demand was 6,900 km^2, nearly 4.5 times the size of Greater London.

AI workloads accounted for roughly 20% of total data centers' electricity use in 2025 and are projected to grow to 40% by 2030.

If AI's share of data center energy rises to 40%, electricity consumption could reach 378 TWh—over 9 times the electricity consumption of Nigeria (world's 6th largest nation by population, with 224 million).

Projected global data centers' electricity consumption could exceed 945 TWh by 2030, accounting for almost 3% of projected global electricity use—enough to supply residential electricity to all 1.3 billion people in Sub-Saharan Africa for about 5.5 years.

If treated as a country, 945 TWh would rank 6th globally by electricity consumption.

The associated water footprint of projected 2030 electricity consumption of data centers is 9.3 trillion liters, or enough to meet the minimum annual domestic water needs of all 1.3 billion residents of Sub-Saharan Africa for a full year.

The associated land footprint of data centers' expected electricity use in 2030 would be over 14,500 km^2, roughly 10 times the size of Mexico City or about twice the Jakarta metropolitan area, home to over 32 million people.

The physical lifecycle of AI hardware presents a growing crisis. AI infrastructure could generate up to 2.5 million metric tons of e-waste annually by 2030, equivalent to discarding nearly 250 Eiffel Towers every year.

EXECUTIVE SUMMARY

AI-generated conceptual image of AI's environmental footprint, generated using OpenAI's ChatGPT/DALL·E, May 2026. Estimated footprints per standard-resolution AI image: 2.9 Wh electricity, 1.22 g CO_2e, 28.6 mL of water, and 0.45 cm^2 of land, based on literature benchmarks and global average electricity footprint factors.

Artificial intelligence (AI) has become a defining technology of the Fourth Industrial Revolution, moving societies beyond digitization toward the optimization and automation of complex systems. AI systems that learn, perceive, reason, and generate content are now embedded across economies and daily lives, supporting services from translation and recommendations to finance, healthcare, energy management, and transportation.

The pace of adoption has been extraordinary. After OpenAI released ChatGPT in 2022, it surpassed one million users in five days and reached 100 million in under two months; by mid-2025, around 700 million people were using it, sending about 18 billion messages each week. Generative AI tools are proliferating across text, images, code, and multimodal applications, with market projections anticipating that AI will become a trillion-dollar-scale industry within the next decade. Global AI market projections span USD 2.4 to 4.8 trillion by the early 2030s.

Yet the infrastructure and benefits of this boom are not evenly distributed. Frontier AI relies on high-end compute, specialized chips, and hyperscale data center capacity that is concentrated in a small number of locations. Many countries lack the domestic infrastructure to train or run frontier models at scale and instead depend on external providers for access to advanced compute. This concentration shapes who captures strategic and economic advantage and who sets the terms of access, pricing, and data governance. **The result is a widening of the existing digital divide between countries that build and control AI systems and those that consume them.**

When a technology scales this quickly, unintended social, economic, geopolitical, and environmental consequences can accumulate quietly and then become complex and inequitable to correct once systems, investments, and dependencies lock in. Public debate has rightly focused on AI risks such as bias, privacy, disinformation, labor disruption, and inequity. Yet, one of the most consequential dimensions of AI that remains comparatively underexamined is its environmental footprint and the justice implications that follow from where and how AI infrastructure expands. AI is not "just code"; it also involves physical infrastructure and supply chains, including data centers, chips, electricity generation, cooling systems, water withdrawals, land occupation, critical minerals, and eventual e-waste.

This report by the United Nations University Institute for Water, Environment and Health (UNU-INWEH) on its 30th anniversary is a step forward in addressing the current gap in AI's environmental governance by assessing its environmental footprints. The investigation goes beyond the carbon-only lens that typically dominates the conversation. It examines AI's indirect environmental footprints through energy use, quantifying the carbon, water, and land footprints associated with generating the electricity required to operate AI at scale, and highlighting how outcomes vary substantially by location depending on electricity supply mixes. This matters because "low-carbon" is not automatically "low-water" or "low-land," and evaluating sustainability through a single metric can hide trade-offs and shift burdens onto places already facing water stress or land pressure. These asymmetries can reinforce the environmental problems of local communities while strategic advantages of AI flow elsewhere.

The study reveals some striking numbers. In 2025, data centers—the physical backbone of AI—consumed an estimated 448 TWh of electricity. If data centers were a country, that level of electricity use would rank it 11th globally. On current trajectories, data center electricity demand could roughly double to 945 TWh by 2030, nearly triple the combined annual electricity use of Pakistan, Bangladesh, and Nigeria, together home to more than 650 million people. Producing that much electricity would have a carbon footprint of 399 million tonnes CO_2e, requiring 6.7 billion trees grown over 10 years to offset—roughly twice the number of trees in the United Kingdom. The associated water footprint of 9.3 trillion liters would be equivalent to the annual domestic water needs of all 1.3 billion residents of Sub-Saharan Africa. The land footprint associated with this electricity would exceed 14,500 km², nearly 10 times the size of Mexico City.

AI is now one of the most significant drivers of that data center growth. In 2025, AI workloads alone accounted for around 20% of total data center electricity use, and if that share rises to 40% by 2030 as projected, its electricity demand could reach roughly 378 TWh—**enough to meet the residential electricity needs of the entire population of Sub-Saharan Africa for over 2 years.**

AI's environmental impacts are shaped not only by data center growth and electricity supply mixes, but also by the escalating cost of building ever-larger models. For example, GPT-3 training consumed an estimated 1.3 GWh of electricity over 34 days, while GPT-4 is estimated to have consumed 50 to 70 GWh over 100 days, roughly 40 to 55 times GPT-3. Yet training is only part of the picture as AI's operational footprint is increasingly driven by inference. Once models are deployed, billions of everyday interactions account for the bulk of energy use, with inference estimated at roughly 80–90% of total energy consumption.

ChatGPT alone is estimated to process around 2.5 billion prompts per day. At scale, small per-task costs become infrastructure-level loads: at a conservative 0.42 Wh per text prompt, ChatGPT-scale use translates into roughly 383 GWh of electricity per year. Offsetting associated carbon emissions would require 2.6 million tree seedlings grown for 10 years, enough trees to cover a land area the size of Manhattan. The water footprint is equivalent to the minimum annual domestic water needs of roughly 500,000 people in Sub-Saharan Africa, and the land footprint is equal to over 800 football (soccer) fields.

The numbers grow drastically once the AI embedded in mass platforms (such as Google Search) is counted. Crucially, per-use energy varies by orders of magnitude across modalities and output lengths, so product defaults and user choices are footprint determinants. A typical ChatGPT-style query is about 200 times more energy-intensive than basic text classification; long GPT-style responses can cost around 1,000 times more than text classification; and AI-generated images can require roughly 1,450–2,000 times the energy of text classification. Video is the new energy frontier: a single short AI-generated video can draw as much electricity as 200,000 spam classifications or hundreds of AI-generated images. Longer responses and richer media can be far more energy-intensive than lightweight text tasks, meaning that image- or video-first routing can materially increase electricity demand and associated carbon, water, and land impacts. Efficiency improvements are important, but they are not sufficient on their own: if lower per-use impacts drive higher volumes of use, total impacts may still rise, a rebound effect often described as the Jevons Paradox.

While this UNU-INWEH report is focused on the environmental footprints of AI's energy use, AI has other major environmental impacts. AI hardware relies on critical minerals whose extraction and processing can cause environmental and social harms, often concentrated in the Global South and in places with weak oversight. At end of life, poorly managed e-waste can expose frontline communities to hazardous substances. By 2030, AI infrastructure could generate up to 2.5 million metric tons of e-waste each year, **roughly equivalent to discarding 250 Eiffel Towers annually.** These impacts show that responsible AI requires full value-chain governance, from mineral sourcing to recycling and safe disposal.

To address these challenges, the report calls for a responsible AI ecosystem grounded in six operational principles: **transparency; efficiency by design; equity and environmental justice; lifecycle responsibility; global cooperation; and sustainable use**. Measuring carbon, water, and land footprints makes impacts visible, comparable, and actionable, enabling governance and supporting accountability across jurisdictions. By recognizing AI's physical impacts and implementing strict lifecycle management, the international community can ensure that technological advancement remains environmentally manageable.

The report identifies priority implications for key stakeholders. Governments must integrate AI infrastructure into energy system planning, carbon accounting, water governance and land use permitting. This requires standardized environmental footprint reporting so impacts can be verified and compared across providers and jurisdictions. Industry and AI developers should treat model selection, default outputs, and routing decisions as footprint determinants, while also improving efficiency-by-design.

Users and deploying organizations also shape impacts through volume, frequency, and modality choices. They should adopt "fit-for-purpose" use: selecting the lightest model and lowest-energy format that meets the task and limiting high-cost features when not needed. Data center operators and utilities should recognize siting and procurement as environmental footprint decisions, apply environmental impact screening and cumulative impact assessment. They should implement transparent mitigation and community safeguards where expansion occurs. Investors and financiers should treat electricity, carbon, water, and land footprints as material risks for AI infrastructure portfolios and use comparable footprint metrics in due diligence.

Communities and civil society should be involved early in siting decisions, with enforceable transparency, consultation, and grievance mechanisms, especially in environmentally-stressed regions. International institutions should support harmonized measurement standards and disclosure practices, reduce incentives for cross-border burden shifting, and support participation and capacity in regions excluded from AI compute. Taken together, these findings provide a practical basis for integrating AI into electricity, carbon, water, and land-use planning, so that innovation advances within environmental limits and without shifting burdens onto vulnerable communities across the globe.

1. AI Across the Globe

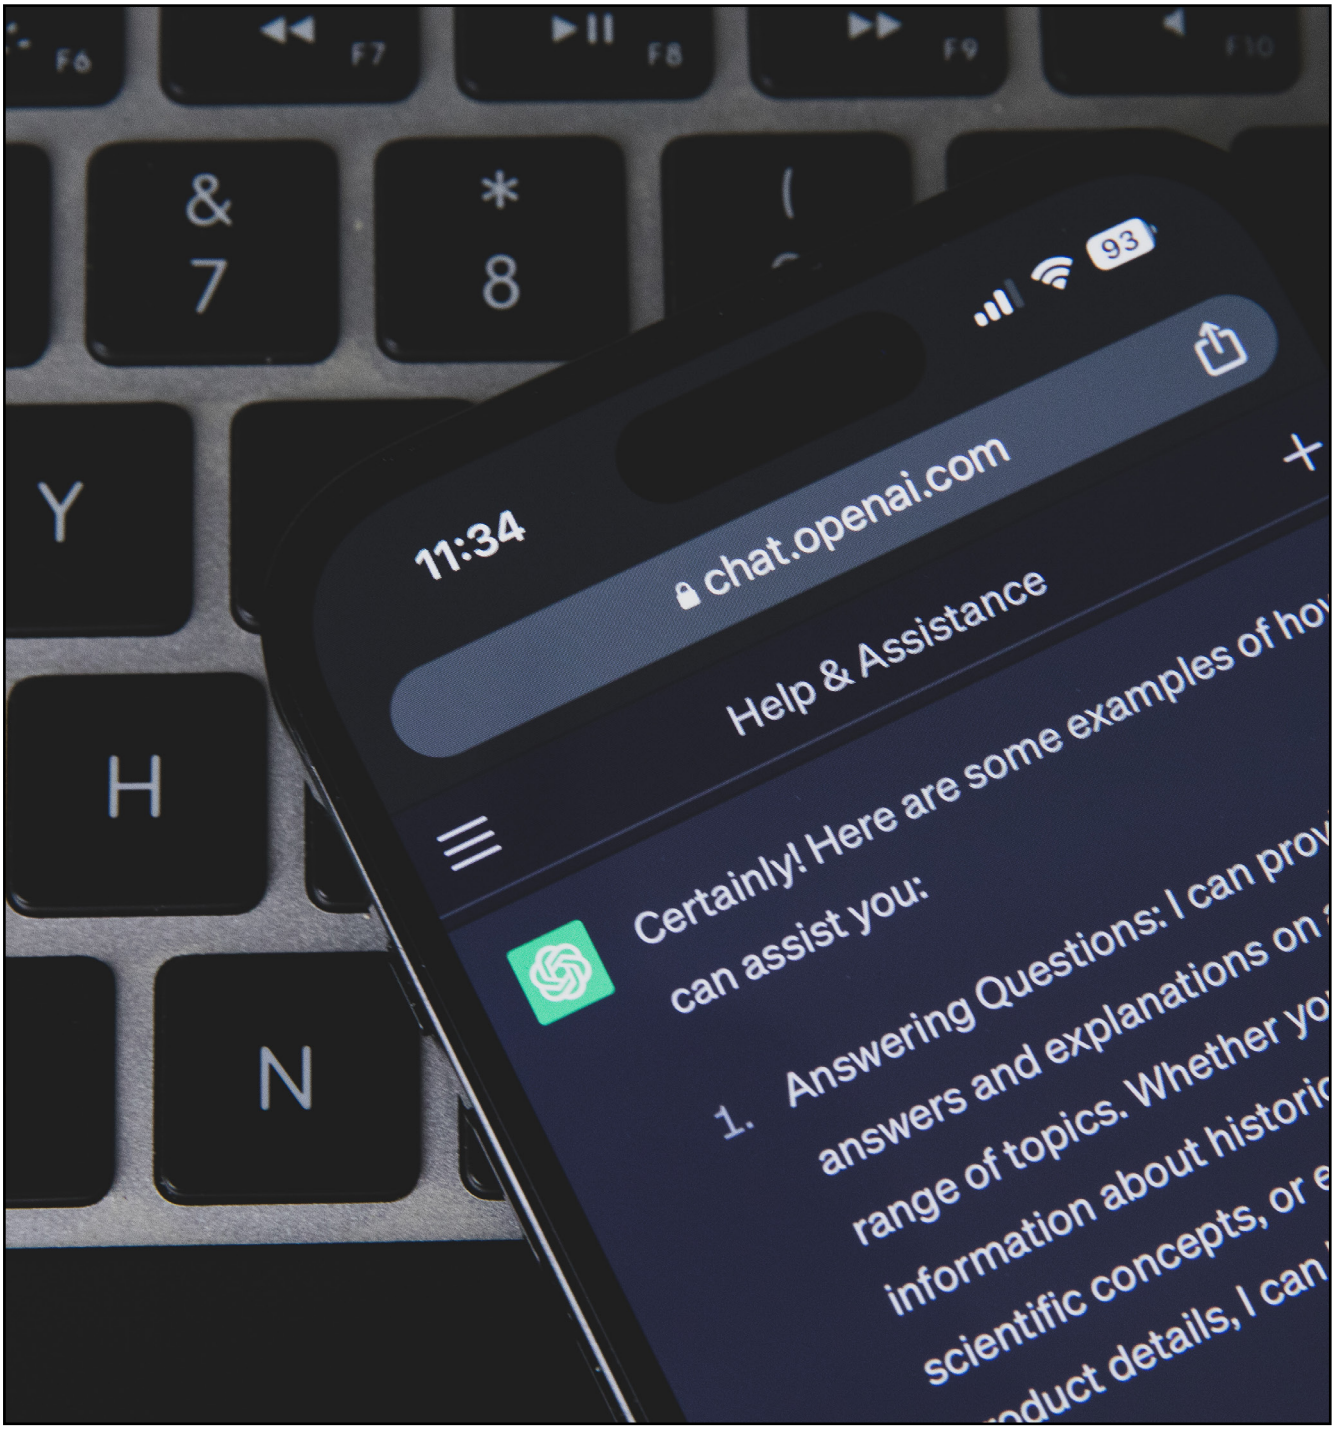

Generative AI is becoming part of everyday life, shaping how millions of people communicate, work, learn, and solve problems. However, this rapid expansion is also increasing the energy, water, and material resources required to sustain the digital world. Photo: Jernej Furman (Wikimedia Commons).