FUTURE THREATS AND RISKS - II: ENVIRONMENTAL IMPACTS OF AI (ENERGY AND WATER)

Dr. Don Liyanage

Tikos Technologies

Prof. Dr. Mehmet Emin Aydın

University of the West of England

Attribution: Liyagene, D, & Aydın, M. E. (2026). Future threats and risks - II: Environmental impacts of AI (energy and water). In H. Arslan, A. Taşdelen, E. Kuzucu Hıdır, & O. A. Çetin (Eds.), *Artificial intelligence and national technology reflections* (pp. 583–598). Turkish Academy of Sciences. https://doi.org/10.53478/TUBA.978-625-6110-86-1.ch24

10.53478/TUBA.978-625-6110-86-1.ch24

FUTURE THREATS AND RISKS - II: ENVIRONMENTAL IMPACTS OF AI (ENERGY AND WATER)

Dr. Don Liyanage
Tikos Technologies

Prof. Dr. Mehmet Emin Aydın
University of the West of England

Abstract

Artificial Intelligence (AI) keeps booming technologically, helping humans in many sectors with succinct skills built in and continually being built. It is apparent and welcomed by humanity, appreciating its assistance in facilitating human life. However, the fact that nothing is gained without any cost in reality, which is applicable to AI, too. As a digital system, AI is a pool of computing tools to process data and produce outcome accordingly, while handling data and producing responsible outcomes remains as one of the paramount concerns. Ethics in AI and data analytics, and responsible AI are one of the key concepts of modern AI studies in order to mitigate the potential threats and risks in the future to any extent. Since Türkiye is one of the fastest-growing countries, the boom in AI and digital systems as well as the advancements in cyber systems, need to be specially cared for. This article overviews and discusses what would be the threats and cause of risks to nature and the environment in the future with respect to energy and water.

Keywords

Future Threats and Risks, Responsible AI, Ethics in AI, Risks upon Energy, Risks upon Water

Introduction

The rapid, almost exponential, growth of Artificial Intelligence (AI) technologies and their sophisticated sub-fields, particularly Generative AI and Large Language Models (LLMs), has heralded a new era of technological capability. Although theoretically the main paradigms of AI remain the same and not much advanced, applications of AI are increasingly hailed as powerful tools for addressing global challenges, from accelerating scientific discovery to enhancing efficiency in resource management and helping to monitor climate change.

The rise of AI in the last two decades of the twenty-first century presents a paradox of materiality. In the public imagination, AI is seen as an ethereal force, a "cloud-based" intelligence composed of weightless code, existing in a frictionless digital common. The industry's discourse reinforces this abstraction with terms such as "virtual assistants", "neural networks" and "serverless computing". Yet the operational reality of the AI revolution is profoundly industrial, grounded in heavy engineering, extractive geology, and thermodynamics. It is a physical colossus that demands the diversion of rivers, the excavation of rare earth minerals, and the consumption of electrical power on a scale that rivals sovereign nations.

As we move from the era of general-purpose computing to one dominated by generative AI and LLMs, the environmental impacts are shifting from minor concerns to major factors influencing geopolitical stability and ecological health. This chapter examines the AI ecosystem's energy and water footprint, arguing that the Twin Transition, aiming for digital progress alongside environmental sustainability, is currently disjointed. Although AI has the potential to improve energy grid efficiency and speed up advancements in material science, its infrastructure expansion poses risks (Dokic et al., 2024) that could hinder global decarbonisation efforts and intensify water shortages in climate-sensitive regions.

However, the physical infrastructure supporting this digital revolution -- primarily massive, high-performance data centres -- requires extraordinary amounts of energy and water, giving rise to an often-overlooked environmental footprint. The expansion of data centre energy use is growing exponentially. The International Energy Agency (IEA) forecasts that worldwide data centre electricity consumption may double from 2024 to 2030, reaching 1,000 terawatt-hours (TWh), approximately equal to Japan's total electricity use. AI significantly drives this growth, with a single query to a generative model consuming ten times more energy than a traditional keyword search. Additionally, supporting infrastructure is exemplified by Microsoft's and OpenAI's proposed $ 500 billion "Stargate" supercomputer project, indicating a future where individual facilities might need up to 5 gigawatts (GW) of power, necessitating a complete overhaul of regional power grids. This chapter delves into the critical and escalating environmental costs associated with the AI boom, specifically focusing on its significant demands on energy and water resources.

This analysis included in this chapter proceeds in three parts. First, we examine the general environmental impact through a cradle to grave Lifecycle Assessment (LCA) of AI hardware, tracing the carbon and toxicological footprint from the mine to the e-waste dump. Second, we analyse the impact upon energy, dissecting the carbon intensity of training versus inference, the crisis of grid capacity, and the “clean energy” procurement challenges facing hyper-scalers. Third, we interrogate the impact upon water, exploring the hydro social conflicts emerging in drought-stricken regions like Chile and Arizona, and the technical trade-offs between water consumption and energy efficiency.

We stand at a critical juncture where the potential benefits of AI risk being overshadowed by its tangible resource consumption. This article explores the current scale of AI's energy consumption, its reliance on often water-intensive cooling systems, and the direct risks these demands pose to climate targets and local ecosystems, especially in water-stressed regions.

1. Environmental Impact of AI

1.1. The material substrate

It is apparent that the growth and spread of AI technologies incurs significant cost and impact upon the environment in many respects. In order to comprehend the true environmental cost of AI, one must look beyond the data centre to the complex, globalised supply chain that constructs the “brain” of the machine.The LCA framework provides the necessary lens, dividing the impact into embodied emissions (manufacturing and transport) and operational emissions (use phase). While operational energy often dominates discussions, the “embodied burden” of AI hardware is growing due to the extreme complexity of modern semiconductor fabrication. This is further elaborated in Table 1.

Table 1

Projected Environmental Impacts of AI (2024 vs. 2030)

Metric	**2024 Baseline (Global)**	**2030 Projection (Global)**	**Primary Driver**
Electricity Consumption	~460 TWh	~1,000 TWh	Generative AI training & inference scaling
AI E-Waste Generation	2.6 thousand tons	0.4 - 2.5 million tons	Rapid hardware obsolescence (GPU turnover)
Water Withdrawal	N/A	4.2 - 6.6 billion m^3 (by 2027)	Cooling needs for high-density compute
Data Centre Emissions share	~0.5% of global CO_2	~1.0 - 1.5% of global CO_2	Grid carbon intensity & demand growth

1.2. The Lithosphere and Extraction: The Mineralogy of Intelligence

The physical foundation of AI is the lithosphere. The specialised hardware required for deep learning, specifically high-performance Graphics Processing Units (GPUs) like NVIDIA's H100 and Tensor Processing Units (TPUs), relies on a specific and scarce mineralogical recipe. Unlike general-purpose CPUs, these accelerators require high-bandwidth memory, I/O, and massive parallel processing cores, which necessitate the use of critical raw materials (CRMs) (Hess, 2024).

There appears an interesting relationship between CRM used in chip production and "Geopolitical Vulnerability". A single AI accelerator card is a compendium of the periodic table. The accelerators contain precious metals such as Gold (Au) and silver (Ag), which are essential for their high conductivity and resistance to corrosion at contact points. The mining of gold is notoriously destructive, often involving cyanide leaching and generating disproportionately high volumes of waste rock per gram of metal extracted. Copper (Cu) is another highly conductive metal commonly used in such industries. It is commonly used wiring and interconnectors of chips and electronic boards of the device. Founding copper is known as highly energy intensive and a significant source of local sulphur dioxide emissions. On the other hand, Rare Earth Elements (REEs) such as neodymium and dysprosium are critical for the permanent magnets in hard drives and cooling fans. At the same time, gallium and indium are used in semiconductor doping (Almasbek et al., 2021).

The supply chain for these materials is fraught with environmental and ethical peril. The extraction of REEs, dominated by processing facilities in China, often involves the release of radioactive tailings and heavy metals into groundwater. Furthermore, the reliance on tantalum and cobalt links the AI supply chain to conflict zones in the Democratic Republic of Congo (DRC), raising questions about the social license of the technology. The "embodied carbon" of these materials is significant; for an NVIDIA H100 GPU, the embodied footprint is estimated at approximately 164 kg CO_2 per card. When deployed in clusters of 100,000 units, the manufacturing footprint alone reaches tens of thousands of tonnes of CO_2 before the system is even powered on.

1.3. Semiconductor Fabrication: The Energy Intensity of the Nanometre

The transformation of raw silicon into a 5-nanometre or 3-nanometre (latest 2-nanometre) logic chip is arguably the most precise and energy-dense manufacturing process in human history. The "Moore's Law" progression towards smaller transistors has necessitated a shift to Extreme Ultraviolet (EUV) lithography.

EUV lithography machines, produced solely by ASML, are the prerequisites for modern AI chips. These machines operate by firing a high-power laser at droplets of molten tin to generate plasma, which emits EUV light. This process is incredibly inefficient in terms of energy conversion; to generate the necessary light output, the laser consumes massive amounts of electricity. Consequently, the electricity consumption of leading-edge fabs (like those of TSMC in Taiwan) has skyrocketed.

Recent studies indicate that manufacturing emissions can account for up to 25% of an AI hardware's total lifecycle emissions (Schneider et al., 2025). This "upfront" carbon debt is exacerbated by the need for ultra clean room environments, which require industrial-scale HVAC systems running 24/7 to filter particulate matter, further driving up the energy intensity per square centimetre of silicon.

1.4. The E-Waste Crisis: The Toxic Afterlife of Compute

The unrelenting pace of AI innovation creates a paradox of obsolescence. To remain competitive in training frontier models, companies must constantly upgrade to the latest hardware architecture (e.g., moving from NVIDIA A100s to H100s, and then to Blackwell B200s). This results in a rapid turnover of server infrastructure, generating a mounting tide of electronic waste (e-waste).

An unavoidable toxicity is projected through the ecosystem of e-waste. The volume of e-waste generated by the AI sector is projected to grow exponentially. From a baseline of 2.6 thousand tons per year in 2023, AI-specific e-waste could reach up to 2.5 million tons annually by 2030. This waste stream is not benign. AI servers contain printed circuit boards (PCBs) laden with lead, mercury, cadmium, and brominated flame retardants (BFRs).

Currently, the global mechanism for handling this waste is failing. Only 22.3% of global e-waste is documented as properly collected and recycled (Baldé et al., 2024). The remaining 78% is often incinerated, landfilled, or exported to the Global South under the guise of used goods, where it is processed in informal, hazardous conditions. This causes a growing deficit in recycling forming materials. The economic loss in this respect emerges through the disposal of previously used materials. This disposal represents a haemorrhage of value. Approximately $ 62 billion worth of recoverable natural resources including the gold and copper mentioned previously are lost annually.

Meanwhile, circular solutions can be managed, where emerging "Urban Mining" initiatives and AI driven sorting technologies attempt to mitigate this, but they are currently outpaced by the volume of waste. The industry lacks a cohesive "take-back" regulation for enterprise AI hardware, meaning much of the retired infrastructure vanishes into the grey market (Rojahn & Grun, 2025).

2. Impact upon Energy: The Carbon Ledger of Intelligence

Energy is one of unavoidable inputs and cost items in the complete ecosystem. If the material impact of AI is the weight of the earth, its energy impact is the heat of the fire. The consumption of electricity by AI systems is the most immediate and scrutinised environmental risk. As models scale from billions to trillions of parameters, the energy required to train and run them is colliding with the physical constraints of power grids and the urgent timeline of climate change mitigation.

2.1. Global Demand and the "Stargate" Trajectory

The scale of energy demand projected for the AI sector is unprecedented in the history of information technology. In 2024, global data centres consumed approximately 415 TWh, or roughly 1.5% of the world's total electricity supply. By 2030, this figure is expected to more than double, reaching nearly 1,000 TWh, roughly equal to the entire annual electricity consumption of Japan (Cossi et al., 2024).

The ambition of the industry is best exemplified by projects similar to Stargate, a joint initiative by Microsoft and OpenAI. This proposed supercomputer campus envisions a power requirement of up to 5 GW. To contextualise, a typical large nuclear reactor generates about 1 GW. Therefore, a single AI facility would require the equivalent of five nuclear reactors dedicated solely to its operation. The US grid is ill-equipped to handle such concentrated loads. The lead time for Large Power Transformers (LPTs), critical for stepping down voltage for data centres, has extended from months to nearly four years (210 weeks) due to supply chain shortages in electrical steel. This "transformer bottleneck" is now a primary constraint on AI deployment, delaying projects and driving up costs.

2.2. The Shifting Load: From Training to Inference

Historically, environmental scrutiny focused on the "training" phase of AI models, the months-long process of teaching a model like GPT-4, which consumes gigawatt hours of energy in a single burst. However, as AI is productized and integrated into search engines, software, and consumer devices, the energy profile is shifting towards inference the act of generating a response. Recent analysis indicates that "inference" now accounts for more than half of the total lifecycle carbon emissions of Large Language Models (Özcan et al., 2025). Two major points (as follows) have given rise in this respect:

(i) *Cumulative Impact:* While training is a discrete event, inference is continuous. A single ChatGPT query is estimated to consume between 10 to 50 times the energy of a standard Google search. With billions of daily queries, the cumulative energy draw of inference quickly surpasses the initial training cost.

(ii) *Jevons Paradox:* Efficiency techniques such as quantisation (reducing model precision from 16-bit to 4-bit) and distillation (using small models to mimic large ones) can reduce inference energy by 50-70%. However, studies suggest that as inference becomes cheaper and more energy-efficient, its usage will proliferate into new applications (e.g., AI in every toaster, car, and phone), leading to a net *increase* in total energy consumption (Fernandez et al., 2025).

2.3. The Reality of Emissions: Location vs. Market

The carbon intensity of AI is geographically determined. A model trained in a data centre in Quebec (hydro-powered) has a vastly different footprint than one trained in Virginia (fossil-fuel heavy). Tech companies often obscure this reality through "Market-Based" reporting, which uses Renewable Energy Certificates (RECs) to offset carbon usage. However, "Location-Based" (physical) emissions are rising (Cao et al., 2022).

Despite claims of "net zero", the absolute emissions of major AI players are increasing. In its 2025 environmental report, Google revealed that while it achieved a 12% reduction in data centre energy emissions through efficiency, its *total* scope 1, 2, and 3 emissions had risen by nearly 48% over a five-year period (2019 baseline), driven largely by the expansion of data centre infrastructure and supply chains. Similarly, Microsoft reported a 29% increase in emissions, admitting that the "moonshot" of becoming carbon negative by 2030 is jeopardised by the explosive growth of AI (Nickelsburg, 2024). The core issue is "firm power" Data centres require 24/7 reliability (baseload). Solar and wind are intermittent. Consequently, without massive battery storage or nuclear power, data centres effectively lock in fossil fuel demand to cover the gaps when the sun is not shining or the wind is not blowing (Budinger, 2024).

2.4. Regional Case Studies: The Physical Limits of the Cloud

The energy demands of AI are causing friction in specific geographies, leading to moratoriums and social pushback. An interesting example is that Dublin, often called the "Data Centre Capital of Europe" faced a crisis where data centres began consuming nearly a quarter of the nation's electricity. This threatened the stability of the national grid. In response, the Irish utilities regulator imposed a *de facto* moratorium on new grid connections for data centres in the Greater Dublin Area. This ban was only lifted in late 2025 under strict new conditions: new facilities must generate their own power or install battery storage to operate off-grid during peak demand. This case illustrates that even developed nations have hard physical limits to digital growth (Duggan, 2025).

Another case given rise in North Virginia, US, where Loudoun County, Virginia, hosts 70% of the world's internet traffic. The density of data centres here has led to the "industrialisation" of the landscape, necessitating massive new transmission lines that cut through communities. This has sparked organised resistance from residents and environmental groups, who argue that the region's power grid and land are being sacrificed for global computing demand. As mentioned in Table 2.

Table 2

Comparative Analysis of Data Centre Cooling Technologies

Technology	**Water Consumption**	**Energy Efficiency (PUE[1])**	**Suitability for AI**	**Sustainability Note**
Air Cooling (Chillers)	Zero	Low (High PUE > 1.4)	Limited for high-density	High energy cost; best for water-scarce regions.
Evaporative Cooling	High (High WUE[2])	High (Low PUE < 1.2)	Standard	Trades water for energy; vulnerable to drought.
Liquid/Immersion	Near Zero	Very High (PUE ~1.05)	Excellent	The "Gold Standard" for sustainable AI enables heat recovery.
Direct-to-Chip	Low/Zero	High	Good	Targeted cooling; reduces fan energy by ~80%.

Furthermore, the lifecycle carbon usage is detailed in Table 3.

Table 3

Lifecycle Carbon Footprint of AI Hardware (Illustrative H100 GPU)

Lifecycle Stage	Share of Emissions	Key Impact Drivers	Mitigation Strategies
Manufacturing (Embodied)	~15-25%	EUV Lithography, ultrapure water, fluorinated gases.	Renewable energy in fabs; extended hardware lifespans.
Operational (Use Phase)	~70-85%	Electricity consumption during 3–5-year lifespan.	Decarbonised grids; liquid cooling; high utilisation rates.
End of Life (E-Waste)	<5%	Hazardous material leaching; recycling energy.	Circular economy; precious metal recovery (Urban Mining).

1 PUE: Power Usage Effectiveness.

2 WUE: Water Usage Effectiveness.

3. Impact upon Water: The Hidden Crisis

While energy consumption dominates the headlines, the impact of AI on water resources is arguably more acute and localised. Data centres are voracious consumers of freshwater, primarily for cooling. As climate change exacerbates drought conditions globally, the "thirst" of AI is placing the industry on a collision course with local communities and agricultural sectors (Li et al., 2025).

3.1. The Metrics of Thirst: Withdrawal and Consumption

In order to clarify and understand the water footprint, two key metrics need to be distinguished: (i) *water withdrawal* is known to be the total volume of water taken from a source (e.g., a river or aquifer). This water is used and may be returned, albeit often at a higher temperature or with chemical additives. (ii) Water consumption is the second metric in this regard, which is known to be the volume of water that is lost permanently from the local watershed, typically through evaporation in cooling towers. Projections indicate that by 2027, global AI demand could account for 4.2 to 6.6 billion cubic meters of water withdrawal, a volume roughly four to six times the total annual water withdrawal of Denmark (Kenny, et al., 2025).

The consumption of water involves operational aspects as a footprint. The consumption is driven by the thermodynamics of high-density computing. As chips like the H100 run hotter, removing that heat requires massive thermal transfer (Privette, 2024). More example cases can give more insights in this regard, where it is estimated that a simple conversation with ChatGPT (comprising roughly 20-50 queries) consumes approximately 500 millilitres of water. While negligible for a single user, when multiplied by hundreds of millions of daily active users, the aggregate consumption is torrential.

The second point appears in the use of water in energy generation, which indicates that there is also a "hidden" water cost in the electricity used. Thermoelectric power plants (coal, nuclear, gas) consume vast amounts of water for cooling. Therefore, the energy inefficiency of a model directly translates into a water footprint at the power plant.

3.2. The Physics of Cooling: Efficiency Trade-offs

The use of compute power and data infrastructure require significant cooling mechanisms/ instruments, either conducted with using water or air cooling. This involves data centre operators facing a fundamental trade-off between energy efficiency and water conservation. (i) Evaporative Cooling is the traditional method involving cooling towers that evaporate water to lower the temperature of the air entering the server hall. This is highly energy-efficient, lowering the Power Usage Effectiveness (PUE) metric. However, it spikes the Water Usage Effectiveness (WUE). Essentially, operators "burn" water to save electricity. (ii) Air Cooling is another alternative method that uses chillers and fans (similar to residential AC) to cool the air without water evaporation. It saves water (zero consumption) but significantly increases electricity usage (worsening PUE).

The use of either of these two methods drives a dilemma, so-called The Drought Dilemma. In water-stressed regions, the reliance on evaporative cooling is becoming politically untenable. This has driven interest in Liquid Cooling (Direct-to-Chip or Immersion), where fluid circulates directly against the hot components. This technology can reduce cooling energy by up to 80% and reduce water consumption to near zero, representing the “gold standard” for sustainable AI (Yañez-Barnuevo, 2025).

3.3. Hydro-Social Conflicts: AI vs. The Community

The use of water for cooling purposes has caused a number of social and communal issues to emerge. The abstraction of the “cloud” breaks down when it competes with local farmers and residents for water rights (Yañez-Barnuevo, 2025).

(i) *The Chilean Resistance:* A landmark case occurred in Santiago, Chile, involving a proposed Google data centre. The region of Cerrillos had been suffering from a decade-long “mega-drought”. Local activists and municipal authorities opposed the project, arguing that the facility's cooling towers would deplete the local aquifer, prioritising server cooling over drinking water. In 2024, an environmental court ruling forced Google to pause the project and revise its design. Google subsequently announced it would abandon water-cooled designs in favour of air-cooled technology at that location. This victory for local activism set a global precedent: water scarcity is a hard constraint on AI expansion.

(ii) *The Taiwan Semiconductor Drought:* The water risk extends upstream to chip manufacturing. TSMC, the foundry for the world’s advanced AI chips, requires “Ultrapure Water” (UPW) for cleaning wafers. In 2021, Taiwan faced its worst drought in a century. The government was forced to halt irrigation to tens of thousands of hectares of farmland to preserve water for the semiconductor fabs. This event highlighted the fragility of the AI supply chain; without water, there are no chips, and without chips, there is no AI.

(iii) *The US Southwest:* In the United States, the migration of data centres to “cheap power” states like Arizona and Texas has placed them in some of the most water-stressed basins in North America. New facilities in these regions face increasing scrutiny, with some municipalities debating whether to grant water permits to new industrial users.

Conclusion

The environmental impact of AI is the defining material challenge of the digital age. Unmanaged AI growth risks undermining broader sustainability efforts. The technology that promises to solve global problems including complex climate problems is currently one of the fastest-growing contributors to the problem itself. To align the transformative potential of AI with global environmental protection, the industry requires urgent policy intervention, including mandatory transparency in reporting energy and water usage, incentives for building resource-efficient systems, and prioritising AI applications whose environmental benefits demonstrably outweigh their footprint. The risks are not merely atmospheric but operational; without securing sustainable energy and water, the AI industry faces physical limits to its own growth. The future of AI will not be determined solely by the elegance of its code, but by the resilience of the physical world that sustains it. The sustainable AI development in the future hinges on recognizing and aggressively mitigating these growing resource demands.

Türkiye is one of the fast-growing countries in technology and use of AI, while energy and water sources are limited. Given that the source of the risks is not atmospheric, but operational, a smart and efficient resource management policy needs to be set up and continuously monitored and updated subject to emerging needs and environmental circumstances.

References

Almasbek, A., Kozlovskiy, A., & Zdorovets, M. V. (2021). The effect of doping with gallium and indium oxides on the optical and shielding characteristics of 0.5TeO2-(0.5-2x)MoO3-xGa2O3-xIn2O3 glasses. *Optical Materials, 118,* 111271. doi:10.1016/j.optmat.2021.111271

Baldé, C. P., Kuehr, R., Yamamoto, T., McDonald, R., D'Angelo, E., Althaf, S., Bel, G., Deubzer, O., Fernandez-Cubillo, E., Forti, V., Gray, V., Herat, S., Honda, S., Iattoni, G., Khetriwal, D. S., Di Cortemiglia, V. L., Lobuntsova, Y., Nnorom, I., Pralat, N., & Wagner, M. (2024). *THE GLOBAL E-WASTE MONITOR 2024*. SCYCLE, UNITAR. https://ewastemonitor.info/the-global-e-waste-monitor-2024/

Budinger, W. D. (2024, February 16). Why renewables cannot replace fossil fuels. Democracy: A Journal of Ideas. https://democracyjournal.org/arguments/why-renewables-cannot-replace-fossil-fuels/

Cao, Z., Zhou, X., Hu, H., Wang, Z., & Wen, Y. (2022). Toward a Systematic Survey for Carbon Neutral Data Centers. *IEEE Communications Surveys & Tutorials, 24*(2), 895–936. doi:10.1109/COMST.2022.3161275

Cossi, L., Spencer, T., & Singh, S. (2024). *Energy and AI*. https://www.iea.org/reports/energy-and-ai/energy-supply-for-ai

Dokic, D., Groen In'T Woud, F., & Maass, W. (2024). *Towards Sustainability of AI: A Systematic Review of Exisiting Life Cycle Assessment Approaches and Key Environmental Impact Parameters of Artificial Intelligence*. https://doi.org/10.24251/HICSS.2024.518

Duggan, J. (2025). Ireland set to end moratorium on new power links to data centers. Bloomberg. https://www.bloomberg.com/news/articles/2025-12-12/ireland-set-to-end-moratorium-on-new-power-links-to-data-centers

Fernandez, J., Na, C., Tiwari, V., Bisk, Y., Luccioni, S., & Strubell, E. (2025). Energy Considerations of Large Language Model Inference and Efficiency Optimizations. *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 32556–32569. https://doi.org/10.18653/v1/2025.acl-long.1563

Hess, J. C. (2024, June 20). *Chip Production's Ecological Footprint: Mapping Climate and Environmental Impact*. Interface. https://www.interface-eu.org/publications/chip-productions-ecological-footprint

Kenny, R., Lunj, A., & Kis, A. (2025). AI's thirst for water. UK Government Sustainable ICT. https://sustainableict.blog.gov.uk/2025/09/17/ais-thirst-for-water/

Li, P., Yang, J., Islam, M. A., & Ren, S. (2025). Making AI Less "Thirsty." *Communications of the ACM*, *68*(7), 54–61. https://doi.org/10.1145/3724499

Nickelsburg, M. (2024). Microsoft's carbon emissions surge despite goal of becoming 'carbon negative' by 2030. KUOW. https://www.kuow.org/stories/microsoft-s-carbon-emissions-surge-despite-goal-of-becoming-carbon-negative-by-2030

Özcan, M., Wiesner, P., Weiß, P., & Kao, O. (2025). Quantifying the energy consumption and carbon emissions of LLM inference via simulations. *ArXiv Preprint, arXiv:2507.11417.*

Privette, A. P. (2024, October 11). AI's Challenging Waters. *Civil and Environmental Engineering - University of Illinois*. https://cee.illinois.edu/news/AIs-Challenging-Waters

Rojahn, M., & Grum, M. (2025). Green AI: A systematic review and meta-analysis of its definitions, lifecycle models, hardware and measurement attempts. arXiv [Cs.AI]. Retrieved from http://arxiv.org/abs/2511.07090

Schneider, I., Xu, H., Benecke, S., Patterson, D., Huang, K., Ranganathan, P., & Elsworth, C. (2025). An Introduction to Life-Cycle Emissions of Artificial Intelligence Hardware. *IEEE Micro*, *45*(5), 9–19. https://doi.org/10.1109/MM.2025.3592568

Yañez-Barnuevo, M. (2025). Data centers and water consumption. Environmental and Energy Study Institute. https://www.eesi.org/articles/view/data-centers-and-water-consumption

About the Authors

Dr. Don LIYANAGE
Chief Scientist & CTO at Tikos | don.liyanage[at]tikos.tech
ORCID: 0000-0002-1358-1970

Dr Don Liyanage is the Co-Founder, Chief Scientist, and CTO of Tikos (2024–Present), a UK-based company specialising in AI assurance and trust. Dr Liyanage trained as a Computer Scientist and holds a PhD focused on AI trust, reasoning, and explainability from the University of Gloucestershire (in collaboration with the University of Manchester). He leads the development of technical frameworks for AI compliance, specifically addressing the "black box" problem in artificial intelligence through 3rd Generation AI technologies. He has authored technical memoranda and papers on the EU AI Act, ISO/IEC AI standards, and Neural Network Reasoning.

Prof. Dr Mehmet Emin AYDIN
University of the West of England | mehmet.aydin[at]uwe.ac.uk
ORCID: 0000-0002-4890-5648

Mehmet Emin Aydin is Senior Lecturer in Computer Science at the University of the West of England. He has received BSc, MA and PhD degrees from Istanbul Technical University, Istanbul University and Sakarya University, respectively. His research interests include machine learning, intelligent agents and multi agent systems, metaheuristics, swarm intelligence, resource planning, scheduling and optimization, combinatorial optimization. Alongside his research and academic portfolio, he has conducted several guest editorial roles for special issues of peer-reviewed international journals. In addition to being the member of advisory committees of many international conferences, he is editorial board member of various peer-reviewed international journals. He is currently fellow of Higher Education Academy, member of EPSRC College in the UK, and senior member of ACM and IEEE.