

Information is Physical:

Assessing the Planetary Trade-offs of AI Transformation

Author: Can Şimşek

Abstract:

This paper examines whether the planetary and environmental costs of contemporary artificial intelligence transformation is beginning to outweigh its promised contributions to climate mitigation and environmental sustainability. It first surveys constructive uses of AI, including applications in climate modelling, energy system optimisation, natural capital accounting and precision agriculture. It then analyses the material footprint of AI across the full life cycle of systems and infrastructure, from mineral extraction and semiconductor fabrication to model training, inference, data centre construction and end of life disposal. The analysis highlights the steep growth in electricity demand, greenhouse gas emissions, pollution and freshwater use associated with large scale model training and generative AI services, as well as the concentration of new data centres in regions already experiencing water stress. It argues that current governance, which relies largely on voluntary corporate disclosure, market based instruments and fragmented national regulation, is inadequate in the context of accelerating investment and an emerging international AI arms race. In response, the paper examines various governance approaches, from pre-deployment planetary, environmental and health impact assessments, standardised life cycle reporting for AI infrastructure, and sustainability by design requirements, to more innovative mechanisms aimed at democratising control over computational resources by prioritising public interest use cases, including the establishment of public interest consumption thresholds and the reservation of accelerator capacity for independently certified AI for Good applications. The paper concludes that a benefit sensitive and climate conscious governance framework is needed to ensure that AI development realises its genuine contributions to societies, including climate mitigation and environmental protection, while remaining compatible with planetary boundaries, resource justice and intergenerational equity. Such a framework should also encourage policymakers to ask, for each use case and major application of AI, whether it is in fact worth its planetary costs.

Introduction:

In 1961, physicist Rolf Landauer proposed the thermodynamic limits of computation, demonstrating that information processing is subject to physical laws.¹ He later articulated this principle with the widely cited phrase, “information is physical.”² The act of computation is therefore not abstract in any material sense: it is grounded in physical processes that consume resources and generate waste.

Every irreversible computational operation, such as the erasure of a bit, incurs a thermodynamic cost and results in the dissipation of heat. In the context of artificial intelligence (AI), this means that algorithmic processing is legally and physically inseparable from its material substrate. Although often obscured by terminology such as “the cloud,” digital operations associated with AI, including data storage, model training, and large-scale inference, are underpinned by physical infrastructures and sustained by continuous flows of energy and material resources. The hardware lifecycle depends on the extraction of critical minerals for semiconductor fabrication; operation generates significant waste heat via electrical resistance, often necessitating liquid cooling to manage thermal limits; physical components eventually degrade due to electromigration and thermal stress; and the process concludes with the disposal of obsolete infrastructure as electronic waste.

Against this backdrop, the present article examines whether the constructive utility of machine learning and related AI techniques in advancing climate and broader environmental sustainability objectives is being outweighed by the resource-, energy-, and extraction-intensive realities associated with the development and deployment of contemporary frontier AI systems, in particular large scale foundation models that underpin generative AI services, as well as AI integrated into robotics. It asks whether this emerging imbalance warrants targeted legal and policy intervention and, if so, which forms of regulatory and governance measures are best suited to address the environmental and climate impacts of AI across its lifecycle.

¹ Rolf Landauer, ‘Irreversibility and Heat Generation in the Computing Process’ (1961) 5(3) *IBM Journal of Research and Development* 183 <https://doi.org/10.1147/rd.53.0183>

² Rolf Landauer, ‘Information is Physical’ (1991) 44(5) *Physics Today* 23 [https://w2agz.com/Library/Limits%20of%20Computation/Landauer%20Article.%20Physics%20Today%2044.%205.%2023%20\(1991\).pdf](https://w2agz.com/Library/Limits%20of%20Computation/Landauer%20Article.%20Physics%20Today%2044.%205.%2023%20(1991).pdf)

1. The Utility of AI for Climate and Environmental Objectives

Artificial intelligence (AI) is a dynamic concept: its meaning evolves as techniques for automating cognitive tasks develop and as systems based on such techniques are increasingly deployed in combination with other digital and data-intensive technologies.³ A semantic gap often exists between technical definitions of AI used by engineers or scholars and the interpretations adopted by policymakers and the general public. Until around 2022, the term AI was commonly used to describe systems based on machine learning (ML), particularly those employing discriminative models (differing from generative models).⁴ These models are designed for tasks such as image classification, pattern recognition, and decision support, focusing on the analysis and categorisation of existing data. Within this framework, a substantial body of scholarship highlighted the potential of ML to support climate mitigation and advance broader public interest objectives.

For instance, ML has been deployed to enhance the efficiency of freight consolidation, refine the design and operation of carbon markets, and reduce food waste across supply chains. Moreover, algorithmic systems based on ML can accelerate scientific discovery by identifying novel materials for batteries, construction, and carbon capture, as well as by enabling the optimisation of complex systems.⁵

AI systems offer computational advantages by accelerating resource-intensive physical simulations, including climate models and energy-scheduling models. A prominent illustration of this constructive aspect is the ARIES (Artificial Intelligence for Environment and Sustainability) platform, an international research and innovation initiative that integrates heterogeneous data and models in order to support natural capital accounting and, more broadly, evidence-based decision-making for environmental sustainability. ARIES underpins, for example, the SEEA Explorer, a tool developed with the United Nations to facilitate rapid and standardised natural capital accounting worldwide.⁶

³ See; Can Şimşek and Ayşe Gizem Yaşar, 'From Rejection to Regulation: Mapping the Landscape of AI Resistance' (Sciences Po Digital Governance and Sovereignty Chair, May 2025) Available on: <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5287068>

⁴ Google ML Education, 'Background: What is a Generative Model? | Machine Learning' (Google for Developers, 18 July 2022) <<https://developers.google.com/machine-learning/gan/generative>>

⁵ Rolnick D and others, 'Tackling Climate Change with Machine Learning' (2022) 55(2) ACM Computing Surveys, Article 42, pp.1-96. <https://doi.org/10.1145/3485128> (Accessed November 2025).

⁶ See ARIES, 'Artificial Intelligence for Environment and Sustainability (ARIES)' <https://aries.integratedmodelling.org/> (Accessed October 2025)

AI systems have also proven valuable for extreme event prediction and analysis. AI's capacity to handle "big data" enables it to fuse information from diverse sources (satellite data, climate records, field sensors etc.), which enhances early warning systems for environmental hazards. Researchers report that AI methods excel at identifying early signals of extreme weather events (such as hurricanes, floods, and heatwaves) and can forecast these events with improved lead times. These techniques not only enhance prediction but also facilitate event attribution and explanation.⁷

Another critical avenue where AI delivers environmental benefits is in optimizing energy systems and other infrastructure to reduce greenhouse gas emissions. In the power sector, ML techniques enable more accurate forecasting of solar and wind generation, support smart grid management, improve load scheduling, and enhance the operation of battery storage, leading to measurable gains in grid and storage efficiency and associated CO₂ reductions.⁸ AI systems are also used to identify optimal sites for renewable energy installations and to manage diverse resources such as wind, hydropower, and solar in an integrated way.⁹

Beyond electricity systems, AI-driven control technologies improve energy efficiency in buildings and industry by dynamically adjusting heating, cooling, and industrial processes, while in transport and urban planning they optimise routing and traffic flows to cut fuel use. In addition, AI is being applied to the design of low-carbon materials and more efficient supply chains in construction and manufacturing, further lowering emissions across multiple sectors.¹⁰

Last but not least, AI systems can enhance agricultural water management through precision irrigation systems, as ML algorithms predict optimal irrigation schedules and enable proactive responses to water stress.¹¹ These innovations are particularly crucial in water-scarce regions, where such efficiency directly impacts food security and environmental sustainability.

⁷ Gustau Camps-Valls, Miguel Á Fernández-Torres, Kai H Cohrs and others, 'Artificial Intelligence for Modeling and Understanding Extreme Weather and Climate Events' (2025) 16 *Nature Communications*, Article 1919 <https://doi.org/10.1038/s41467-025-56573-8>

⁸ Joanna I. Lewis, Autumn Toney and Xinglan Shi, 'Climate Change and Artificial Intelligence: Assessing the Global Research Landscape' (2024) 4 *Discover Artificial Intelligence* 64 <https://doi.org/10.1007/s44163-024-00170-z> accessed 12 November 2025.

⁹ Ibid.

¹⁰ Ibid.

¹¹ Ahmed Elbeltagi, Nand Lal Kushwaha, Ankur Srivastava and Amira Talaat Zoof, 'Artificial Intelligent-Based Water and Soil Management' in Ramesh Chandra Poonia, and others (eds), *Deep Learning for Sustainable Agriculture* (Academic Press 2022) 129.

These examples represent only a subset of AI applications with constructive potential for climate and environmental goals. As the field rapidly evolves, new use cases continue to emerge across technical domains and policy contexts. On the other hand, artificial intelligence has also been characterised as an extractive industry, even prior to the widespread deployment of generative models. As Kate Crawford observes, the creation of AI systems relies on “exploiting energy and mineral resources from the planet, cheap labour, and data at scale,” situating AI within a broader political economy of resource extraction and infrastructural dependence.¹² With the rapid expansion of generative models capable of producing synthetic content, the environmental and societal implications of AI have become increasingly contested, prompting renewed scrutiny of its overall utility. The following section examines these concerns in greater depth.

2. Generative AI and Accelerated Extractivism

Natural language processing (NLP) has been one of the first subfields of AI to foreground concerns about the environmental impact of large scale models, marking an important shift in debates on responsible AI development. Since around 2018, the development of large language models (LLMs) has accelerated rapidly, with systems such as BERT, GPT-2, GPT-3 and Switch-C increasing dramatically in parameter counts and training data volumes, driven by intense competition to achieve state of the art performance on tasks including machine translation, question answering and other core NLP benchmarks. The trajectory from BERT to GPT-3 and beyond saw model parameters balloon from millions to 175 billion, with Switch-C reaching 1.6 trillion parameters, directly linking improved performance to escalating computational demands.¹³ In 2021, Bender *et al.* warned against the environmental consequences of this exponential growth, and advocated for fundamental shifts toward "energy efficient model architectures and training paradigms," urging researchers to "prioritize energy efficiency and cost to reduce negative environmental impact" and invest in "careful curation of data" rather than simply "ingesting massive amounts of data" from convenient or easily-scraped Internet sources.¹⁴

¹² Kate Crawford, *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (Yale University Press 2021) 14.

¹³ Emily M Bender et al., 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?' in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (ACM 2021) 610, 611.

¹⁴ *Ibid.* pp. 613-618.

In November 2022, OpenAI released its consumer-facing generative AI chatbot, ChatGPT, built on the GPT-3.5 language model. Within approximately two months, ChatGPT had reportedly reached over 100 million users, making it widely described as the fastest-growing consumer application at that time.¹⁵ Since then, the trajectory of the AI field has largely moved in the opposite direction from that urged by Bender and colleagues, who urged caution in scaling increasingly large language models due to their significant environmental and social costs. Instead, the mainstream paradigm has consolidated around rapidly expanding LLMs and, increasingly, general-purpose “foundation models” trained on web-scale data and deployed across a wide range of applications.¹⁶

In parallel, more efficient architectures and comparatively smaller models have begun to emerge, promising reduced computational and energy requirements for specific use cases. Yet these efficiency gains have so far been outweighed by the “rebound effect”. Soaring demand for generative AI services and an ongoing race to train more capable models have driven unprecedented investment in data centres and high-performance computing infrastructure, reinforcing rather than mitigating AI’s overall resource footprint. This pattern is corroborated by Morand, Ligozat and Névéol, who demonstrate that, despite substantial improvements in GPU efficiency and algorithmic optimisation, the energy use and environmental impacts of model training rose exponentially between 2013 and 2023, even after accounting for shifts to lower-carbon electricity mixes, and interpret this trend as a classic rebound effect in which efficiency gains fuel the development and training of ever larger models.¹⁷

As the field shifts from relatively simple text processing towards high-dimensional image and video generation, the energy cost per query rises sharply. Generative AI is accordingly projected to drive a substantial expansion of data-centre infrastructure: the International Energy Agency estimates that by 2026 the AI industry alone may consume at least ten times more electricity than it did in 2023¹⁸ A study by Desroches *et al.* project corporate AI portfolios to 2030 and, in a high-adoption scenario dominated by generative models, estimate

¹⁵ D Milmo, ‘ChatGPT reaches 100 million users two months after launch’ *The Guardian* (London, 3 February 2023) <https://www.theguardian.com/technology/2023/feb/02/chatgpt-100-million-users-open-ai-fastest-growing-app>

¹⁶ See; Can Şimşek and Ayşe Gizem Yaşar, ‘From Rejection to Regulation: Mapping the Landscape of AI Resistance’ (Sciences Po Digital Governance and Sovereignty Chair, May 2025) pp. 15-16.

¹⁷ C Morand, A-L Ligozat and A Névéol, ‘How Green Can AI Be? A Study of Trends in Machine Learning Environmental Impacts’ (preprint, arXiv, 23 December 2024) <https://arxiv.org/abs/2412.17376> accessed 1 November 2025.

¹⁸ International Energy Agency, *Electricity 2024: Analysis and Forecast to 2026* (IEA 2024).

that AI-related electricity use could increase by a factor of around 24, with large generative models consuming up to 4,600 times more energy than traditional models.¹⁹

Since the environmental impact of AI is not limited to the training phase, the following chapters will break down the full life cycle of AI systems.

3.1. Model Training

Training large-scale models carries significant environmental consequences. These models require large volumes of computation, typically involving thousands of GPUs or TPUs operating in parallel for extended periods. As models have evolved to process not only text but also images, audio, and video, their training demands have intensified. This has increased both the energy consumption and associated carbon emissions of AI development.²⁰

Publicly available data on the carbon footprint of large-scale model training remain scarce. An exception is BLOOM, a 176-billion-parameter open-access language model released in 2022, whose training reportedly consumed approximately 433 megawatt-hours of electricity and generated 24.7 tonnes of CO₂-equivalent in operational emissions, rising to over 50 tonnes when embodied emissions are included.²¹ This represented approximately one-twentieth of GPT-3's 502-tonne footprint, a reduction primarily resulting from lower-carbon energy sources.²² By contrast, most recent commercial frontier models have not disclosed their energy use or even basic architectural information such as parameter counts, a lack of transparency that precludes independent assessment of the environmental impacts of contemporary AI development.

Model training not only consumes electricity but also substantial quantities of freshwater. These concerns are particularly acute in relation to large scale AI models. Li and co authors estimate that training the GPT-3 model in Microsoft's United States data centres directly evaporated around 700,000 litres of clean freshwater as on site cooling water and about 5.4

¹⁹ Clément Desroches *et al*, 'Exploring the sustainable scaling of AI dilemma: A projective study of corporations' AI environmental impacts' (preprint, 24 January 2025). <https://arxiv.org/pdf/2501.14334>

²⁰ Sasha Luccioni, Yacine Jernite and Emma Strubell, 'Power Hungry Processing: Watts Driving the Cost of AI Deployment?' (2024) *FAccT '24: Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* 85, 92–93.

²¹ Sasha Luccioni and others, 'Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model' (2023) *7 Journal of Machine Learning Research* 1, 4–5.

²² *Ibid*.

million litres in total once indirect water use for electricity generation and server manufacturing is included.²³ The training run is only one stage in the model's life cycle.

3.2. Inference and Downstream Use

Inference can be simply defined as the process of using an AI model, trained from data and/or expert rules, to generate predictions, recommendations or other outcomes from new input data.²⁴ Once a model is trained and deployed, inference (such as responding to user queries) can become the dominant source of energy use, since it continues for the model's lifetime and scales with user traffic.

Researchers observe that while “traditional” ML workloads might balance computation between training and inference, generative AI inference is poised to dominate energy use of AI as these models become ubiquitous.²⁵ The shift from information retrieval to content generation entails a significant energy premium. Every ChatGPT-style query or image-generation request consumes a nontrivial amount of power, and billions of such queries quickly add up. Indeed, the International Electricity Agency (IEA) reports that a single ChatGPT request draws on average about 2.9 Wh of electricity, versus ~0.3 Wh for a typical Google search, which implies roughly a tenfold difference per query.²⁶ In practical terms, if the ~9 billion searches performed daily were instead run through a ChatGPT-like service, it would require on the order of 10 TWh of extra electricity per year.²⁷

Moreover, inference can also consume substantial amounts of water, since most AI services run their computations not on users' devices but in remote data centres, where servers and cooling systems draw on local water resources to manage heat generated by large-scale processing. Li *et al* demonstrate that roughly 10 to 50 medium-length user interactions with a GPT-3-class model can together account for about 500 millilitres of freshwater use, once both cooling water and the water intensity of power generation are included, which implies a significant operational water footprint when such systems are deployed at scale.²⁸

²³ Pengfei Li *et al.*, ‘*Making AI Less “Thirsty”*: Uncovering and Addressing the Secret Water Footprint of AI Models’ (2023) <https://arxiv.org/pdf/2304.03271>

²⁴ ¹ OECD, *OECD Framework for the Classification of AI Systems* (OECD 2022) 48.

²⁵ Adam Zewe, ‘Explained: Generative AI’s Environmental Impact’ (*MIT News*, 17 January 2025) <https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117> accessed 21 January 2025.

²⁶ International Energy Agency, *Electricity 2024: Analysis and Forecast to 2026* (IEA 2024) 34 <https://iea.blob.core.windows.net/assets/6b2fd954-2017-408e-bf08-952fdd62118a/Electricity2024-Analysisandforecastto2026.pdf> accessed 24 March 2025.

²⁷ *Ibid*, p.34.

²⁸ Pengfei Li *et al.*, ‘*Making AI Less “Thirsty”*: Uncovering and Addressing the Secret Water Footprint of AI Models’ (2023) <https://arxiv.org/pdf/2304.03271>

As AI shifts from text to computationally intensive tasks (high-resolution image, video, etc.), per-query environmental costs jump orders of magnitude. Luccioni *et al.* reported that image generation tasks required roughly 60 times more energy than text generation models as of 2023.²⁹ Diffusion-based image generators (like Stable Diffusion or Google’s recent “Nano Banana” model) produce high-resolution images by running dozens of denoising steps for each output. This iterative process means a single image generation can require *billions* of floating-point operations.³⁰ When extended to video (which repeats this work for many frames), the energy use grows even more.

Last but not least, Luers *et al.* note that embedding AI into existing applications in sectors such as health care, transport and entertainment would increase overall electricity demand by expanding digital services. They further observe that AI can improve the efficiency and reduce the cost of oil and gas exploration and extraction, which may in turn incentivise higher fossil fuel production. Thereby, the authors warn that, without appropriate governance, widespread AI deployment could affect political and economic stability and interact with poverty, food security and social inequality in ways that have indirect consequences for greenhouse gas emissions.³¹

After all, the environmental impact of inference depends on query volume (how many users and how often they use the model), the model’s size and architecture, the efficiency of the hardware serving it, and the carbon and resource intensity of the electricity supply. Importantly, inference is continuous and potentially very large-scale, unlike training which is episodic.

3.3. Infrastructure and Supply Chain

The lifecycle emissions of artificial intelligence begin with the extraction of raw materials such as rare earth elements used in semiconductor fabrication, and extend through the manufacture of specialised hardware (e.g. GPUs and TPUs) and the construction of data

²⁹ Alexandra Sasha Luccioni, Yacine Jernite and Emma Strubell, ‘Power Hungry Processing: Watts Driving the Cost of AI Deployment?’ (2023) arXiv preprint arXiv:2311.16863, 5.

³⁰ Aniketh Iyengar et al., ‘Energy Scaling Laws for Diffusion Models: Quantifying Compute and Carbon Emissions in Image Generation’ (preprint, arXiv, 21 November 2025) <<https://arxiv.org/pdf/2511.17031v1>> accessed 22 November 2025.

³¹ Amy Luers et al., ‘Will AI Accelerate or Delay the Race to Net-Zero Emissions?’ (2024) *Nature* <https://www.nature.com/articles/d41586-024-01137-x>

centres, before continuing into the operational phase in which computational services are delivered at scale. The up-front stages are often referred to as “embodied” or “embedded” emissions. It concludes with the end-of-life phase, which includes the decommissioning of data centres, the disposal or recycling of electronic waste, and the environmental costs associated with hardware obsolescence and infrastructure turnover.

While operational impacts and carbon-focused analyses have received growing attention, full lifecycle impacts remain systematically underreported in both academic literature and corporate disclosures, limiting the sector’s ability to quantify and mitigate its total environmental footprint.³²

Notably, decarbonisation technologies and large-scale AI computing rely on many of the same mineral and metal supply chains. For instance, copper is indispensable for the expansion of renewable energy and the reinforcement of electricity grids. At the time of writing, the rapid growth of AI data centres has sharply increased demand for copper, placing additional pressure on the same upstream mining and refining capacity that the green transition depends on.³³

3.4. Water Scarcity as a Standalone Challenge in AI Transformation

The criticality of water resources in the context of AI infrastructure warrants particular scrutiny given water's fundamental role in sustaining human life, ecological systems, and economic development. Approximately half of the world's population are assessed as being subject to severe water scarcity for at least some part of the year due to both climatic and non-climatic factors.³⁴ The United Nations has recognised access to clean water and sanitation as a fundamental human right, enshrined in numerous international instruments.³⁵

Data centres consume water through two principal pathways: indirectly, via the water required for electricity generation, especially thermoelectric power, and directly, through cooling systems that dissipate heat from servers. According to an estimation, a medium sized (15 megawatt) data centre uses approximately as much water as three average sized

³² Sophia Falk *et al.*, ‘More than Carbon: Cradle-to-Grave Environmental Impacts of GenAI Training on the Nvidia A100 GPU’ (2025) arXiv preprint arXiv:2507.05401 <https://arxiv.org/pdf/2509.00093>

³³ Camilla Hodgson, ‘The hunt for copper to wire the AI boom’ Financial Times (London, 3 December 2025)

³⁴ IPCC, ‘Climate Change 2022: Impacts, Adaptation and Vulnerability- Chapter 4: Water’ (Contribution of Working Group II to the Sixth Assessment Report, Cambridge University Press 2022) 551.

³⁵ UNGA Res 64/292 (28 July 2010) UN Doc A/RES/64/292

hospitals.³⁶ This amount is highly significant at the local level where data centres draw on the same finite resources as agriculture and households. In regions experiencing water stress, the allocation of potable water to data centre cooling operations creates direct competition with essential human needs, including drinking water, sanitation, food production, and public health services. At the time of writing, recent reports indicate that the expansion of AI data centres is continuing in water-scarce regions, thereby intensifying pressure on already stressed local water resources. For example, three planned Amazon data centres in the water-stressed Aragón region of northern Spain are reportedly licensed to draw an estimated 755,720 cubic metres of water per year.³⁷ This volume is estimated to be sufficient to irrigate roughly 233 hectares of corn, one of the region's main crops. In March 2025 the regional government requested EU assistance to address severe drought conditions, while local groups such as *Tu Nube Seca Mi Río* ("Your cloud is drying my river") have called for a moratorium on new data centres due to water scarcity.³⁸

3. The AI Era in the Anthropocene: A Case for Regulatory Intervention

Although the impact of AI transformation within the context of the Anthropocene and a full life-cycle assessment of AI systems lie beyond the scope of this article and are more appropriately situated within the domain of climate and environmental systems science, existing research increasingly suggests that the net environmental effects of AI transformation may already be unfavourable.

To obtain a holistic understanding of the climate and environmental impacts of AI transformation, its potential efficiency gains must be weighed against its resource demands and associated emissions. For example, in contexts of water scarcity, the benefits of precision irrigation in agriculture should be considered alongside the significant water consumption required for data centre cooling. Such assessments should be conducted at both local levels, to capture ecosystem-specific impacts, and at the global level, to evaluate broader climate and resource implications. Currently, the lack of reliable and standardised data, combined with

³⁶ David Mytton, *Data centre water consumption*. (2021) *npj Clean Water* 4, 11.
<https://doi.org/10.1038/s41545-021-00101-w>

³⁷ Luke Barratt and Costanza Gambarini, 'Revealed: Big tech's new datacentres will take water from the world's driest areas' *The Guardian* (London, 9 April 2025)

<https://www.theguardian.com/environment/2025/apr/09/big-tech-datacentres-water> accessed 4 November 2025.

³⁸ *Ibid.*

the AI sector's opacity and inconsistent voluntary reporting, hinders meaningful evaluation of environmental trade-offs.

Despite these limitations, the scale and pace of investment in AI infrastructure, particularly in support of compute intensive technologies such as generative AI, large foundation models, agentic systems, and embodied AI, indicate a rapidly expanding environmental footprint. A prominent example is the US-based Stargate project, which aims to invest up to US\$500 billion to construct data centres providing approximately 10 gigawatts of capacity for OpenAI.³⁹ This initiative illustrates the magnitude of energy and resource commitments being mobilised to support frontier AI, with significant implications for sustainability and long-term resource governance.

Meanwhile, local communities across jurisdictions now frequently contest proposed data centre developments owing to their immediate environmental and social externalities. Reported concerns include substantial water withdrawals that strain local resources and compete with residential and agricultural needs in water-stressed regions; destabilised local grids or increase in electricity prices; and different types of pollution including noise and light.⁴⁰ At the same time, civil-society organisations, academic researchers and international scientific bodies warn that the rapid expansion of AI infrastructure is contributing to energy use and emissions trajectories that are misaligned with the temperature goals of the Paris Agreement, with repeated emissions gap assessments indicating that existing national pledges remain insufficient for 1.5°C and challenging even for 2°C in the absence of additional mitigation.⁴¹

In this regard, a compelling historical parallel can be drawn between contemporary resistance to data centres and the public opposition that accompanied the rise of industrial factories in the 18th and 19th centuries. The proliferation of chemical plants, tanneries, and other heavy industries during the Industrial Revolution frequently provoked resistance from adjacent communities. These objections were primarily grounded in the severe, localized environmental hazards these facilities produced, including noxious air pollution, the

³⁹ OpenAI, 'OpenAI, Oracle, and SoftBank Expand Stargate with Five New AI Data Center Sites' (23 September 2025) <https://openai.com/index/five-new-stargate-sites/>

⁴⁰ See; Can Şimşek and Ayşe Gizem Yaşar, 'From Rejection to Regulation: Mapping the Landscape of AI Resistance' (Sciences Po Digital Governance and Sovereignty Chair, May 2025), pp 70-74. Available on: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5287068

⁴¹ United Nations Environment Programme, *Emissions Gap Report 2025* (UNEP 2025).

contamination of water supplies, and general fears for public health and safety. While the technology has evolved, the core grievance remains similar: the imposition of environmental externalities and risks upon local populations who often bear the costs without reaping the proportional benefits of technological advancement.

A significant difference from earlier waves of industrial paradigm shifts is that the systemic risks of climate change are now widely recognised, and a relatively dense architecture of international climate governance already exists.⁴² Yet, contemporary legal frameworks address the environmental impacts of AI only partially and without coherence. Under the UNFCCC and the Paris Agreement, AI infrastructure is not recognised as a discrete source of greenhouse gas emissions, water consumption, or material extraction subject to specific mitigation or reporting obligations. Instead, the environmental footprint of AI data centers and hardware is addressed only indirectly through broad sectors, leaving a regulatory gap for the life-cycle impacts of AI. In practice, governance of AI infrastructure’s climate impacts relies predominantly on market-based and voluntary mechanisms rather than binding rules. This creates a patchwork of self-reporting practices and voluntary disclosure schemes with uneven transparency. Furthermore, firms often resort to instruments such as Renewable Energy Certificates (RECs) to claim their electricity use is “green.” Critics note that these accounting tools can obscure the true material impacts of AI infrastructure and enable “*creative accounting*”.⁴³ For example, studies have found that widespread reliance on RECs lets tech companies overstate emissions reductions on paper, without yielding genuine on-site cuts in resource use or carbon output.⁴⁴

At the time of writing, there has been limited progress in addressing this regulatory gap, and the prevailing tendency remains to defer to market-based and self-regulatory approaches within the industry. For instance, the 2025 UN Climate Change Conference (COP30) in Belém did not adopt any specific, binding measures targeting AI infrastructure or data centres. Nonetheless, civil-society organisations used the summit and its parallel events to highlight the rapidly growing energy and resource footprint of AI and cloud infrastructure,

⁴² See, among others: United Nations, *United Nations Framework Convention on Climate Change* (adopted 9 May 1992, entered into force 21 March 1994) 1771 UNTS 107; Paris Agreement (adopted 12 December 2015, entered into force 4 November 2016) 3156 UNTS 3

⁴³ Isabel O’Brien, ‘Data center emissions probably 662% higher than big tech claims. Can it keep up the ruse?’ (The Guardian, London, 15 September 2024)

<<https://www.theguardian.com/technology/2024/sep/15/data-center-gas-emissions-tech>>

⁴⁴ ECOS, Open Future, *From Innovation to Overshoot: How Data Centre Expansion Risks Derailing Climate Goals* (September 2025) <https://ecostandard.org/wp-content/uploads/2025/09/Data-centres-report.pdf>

and to call for measures such as mandatory environmental impact assessments for new data-centre projects and requirements that these facilities be powered by new, additional renewable energy capacity rather than drawing on existing supplies.⁴⁵

Historically, regulations, standards, and risk-containment methods have tended to follow and legitimise technological *faits accomplis* rather than constrain them as *ex ante* norms. Environmental historians have emphasised, for example, that from the nineteenth century onwards processes of administrative centralisation and industrialisation in France facilitated the expansion of highly polluting industries, as the Industrial Revolution reshaped socio-technical reality.⁴⁶ As a result, harms to air, water and soil increasingly ceased to be treated as violations of the environment understood as a strictly protected common good and were instead reframed as individualised injuries remediable through financial compensation in civil courts, or as issues to be managed within administrative regulation.⁴⁷ Arguably, the contemporary transformation driven by artificial intelligence represents a turning point in the Anthropocene, both in its impact on the environment and climate and in the forms of governance emerging around it. The global climate system, traditionally associated with the concept of *res communis omnium*, is increasingly managed through market-based mechanisms, including carbon offset schemes, emissions trading systems, and large-scale infrastructure planning. This approach poses the risk of institutionalising high-carbon and resource-intensive development trajectories over the long term.

On the assumption that scientific communities and policymakers retain some scope for democratic agency, the following discussion investigates which regulatory and policy measures could reduce the adverse climatic and environmental implications of the ongoing AI transition. It invites policymakers, technology leaders and further readers to assess the benefits of AI systems in relation to their societal and ecological costs and to reflect on regulatory options in light of this assessment.

⁴⁵ Anton L. Delgado, ‘Artificial Intelligence Sparks Debate at COP30 Climate Talks in Brazil’ (AP News, 15 November 2025) <https://apnews.com/article/climate-change-artificial-intelligence-brazil-cop30-3f5a9ddc3f2a65cd00d9a0bf06002b91>

⁴⁶ Jean-Baptiste Fressoz, *Happy Apocalypse: A History of Technological Risk* (Verso Books, 2024) pp.109-122.

⁴⁷ *Ibid.*

4. Key Regulatory Policy Options for Sustainable AI

In light of the mounting environmental and climate risks outlined above, there is a compelling case for regulatory intervention that transcends the short-term horizons of market-driven decision-making, which systematically fails to account for climate externalities (greenhouse gas emissions and related atmospheric impacts) and environmental externalities (resource depletion, pollution, and biodiversity loss). Such intervention is essential not only to unlock the positive transformative potential of AI, but also to ensure that this transformation proceeds on a sustainable basis by addressing environmental harms that could feed back into the AI sector through resource constraints, infrastructure vulnerabilities, and reputational risks. More broadly, regulatory frameworks must be designed with intergenerational justice in mind, ensuring that the long-term social and ecological burdens of AI transformation are not shifted onto future populations. To this end, this chapter delineates three key regulatory approaches: enhancing transparency, embedding sustainability by design, and regulating AI use cases through a planetary perspective.

4.1 Sustainability-by-Design Requirements

Sustainability-by-design should be treated as a regulatory principle for AI infrastructure, meaning that environmental and planetary impacts are addressed from the outset and integrated across hardware, software, and operational life cycles, rather than managed only after deployment.

To begin with, site and permitting policies should favor locations with access to low-carbon electricity and where waste heat can be integrated into district heating, industrial processes, or other local heat networks.

At the hardware level, this implies circular-economy strategies. AI servers and accelerators should be designed for durability, repair, and disassembly, with modular and upgradeable architectures that extend product lifetimes. Extended lifespans can reduce the need for premature replacement.

At software level, efficiency should be treated as a first-order design criterion alongside accuracy, so that unnecessarily compute-intensive models are avoided where lower-impact alternatives suffice. For instance, recent scholarship has suggested a “Sustainability Alignment Tax,” a conceptual framework that treats some loss of AI performance,

convenience or profit as a normative “tax” that developers and regulators should be willing to pay to keep AI development aligned with environmental and socio-economic sustainability.⁴⁸

Operational requirements include powering AI data centers with verifiable renewable energy sources and incorporating waste-heat recovery. For example, regulation could mandate additional renewable generation (or direct renewable power contracts). Waste heat should be captured and repurposed (for space heating, industrial use, etc.) wherever practical.

Finally, an internationally coordinated restriction on the use of potable water for data centre cooling in water-stressed regions would offer a focused and effective means of reducing the water footprint of AI infrastructure.

4.2 Planetary, Environmental and Health Impact Assessment

At the time of writing, prevailing approaches to assessing the sustainability of data centres and AI systems remain incomplete, as they predominantly focus on greenhouse gas emissions and, more recently, water consumption, while failing to incorporate systematic assessment of criteria air pollutants and their associated public health impacts across the full life cycle. This omission is consequential: air pollutant emissions arising from electricity generation, backup power systems, and upstream hardware production can generate substantial, spatially uneven, and potentially inequitable health burdens that are not well captured by carbon centric reporting alone.⁴⁹

To address such gaps, a robust planetary computing regulation could introduce a Planetary, Environmental and Health Impact Assessment (PEHIA) conducted prior to the construction of AI infrastructure such as data centres. This assessment should incorporate site selection criteria that prioritise access to low-carbon energy and distance to residents, evaluate availability for cooling, analyse public health risks and estimate projected life-cycle impacts including material extraction, construction emissions, and hardware production.

As for orbital data centres, a PEHIA could extend beyond terrestrial considerations to account for the full life cycle and transboundary impacts of space based infrastructure. Although

⁴⁸ Giulio Amore and Andrea Gentili, ‘AI and Climate: An Ethical Sustainability Framework for Balancing Risks and Responsibilities’ (2025) *AI & Society* (advance access) <https://doi.org/10.1007/s00146-025-02615-0>

⁴⁹ Han Yu et al, 'The Unpaid Toll: Quantifying and Addressing the Public Health Impact of Data Centers' (2024) arXiv <https://arxiv.org/abs/2412.06288>

orbital data centres remain largely conceptual or at very early demonstration stages as of 2026, anticipatory governance is warranted given the pace of commercial space development. For instance, rocket launches introduce pollutants directly into the upper atmosphere, where they persist far longer than at ground level and contribute to stratospheric chemical disruption.⁵⁰ The accumulation of orbital debris poses well established collision risks that have been intensified by the growing number of satellite constellations. The occupation of finite orbital resources and the interference of satellites with astronomical observation represent opportunity costs for scientific research and other critical uses of the orbital commons. Upstream impacts from material extraction, hardware manufacturing and energy system production remain significant across the full life cycle, as do the environmental and safety risks posed by re-entry and decommissioning. Incorporating these factors into a comprehensive assessment framework ensures that environmental, public health and shared orbital resource burdens are neither displaced nor overlooked, providing a necessary foundation for transparency and precautionary governance of an emerging domain.

4.3 Transparency Requirements

The environmental governance of AI infrastructure is currently characterised by limited transparency, fragmented disclosure practices, and a heavy reliance on market-based instruments and voluntary corporate reporting. This situation hampers the ability of regulators, policymakers, and civil society to assess whether current trajectories are aligned with e.g. the Paris Agreement's temperature goals or to verify corporate climate commitments in a consistent and verifiable manner. A first policy necessity, therefore, is the introduction of mandatory, standardised life-cycle reporting for AI systems and associated infrastructure.

Complementing pre-deployment PEHIA, post-deployment transparency obligations should include continuous, standardised reporting of energy use during training and inference, associated emissions, water consumption, and end-of-life hardware disposal. Furthermore, this reporting should extend across the full life cycle of AI systems, encompassing raw material extraction, semiconductor fabrication, and hardware manufacturing. At a minimum, providers of large foundation models and operators of AI-intensive data centres should be required to publish, on an annual basis: (i) total energy consumption, disaggregated by energy

⁵⁰ Robert G. Ryan et al., 'Impact of Rocket Launch and Space Debris Air Pollutant Emissions on Stratospheric Ozone and Global Climate' (2022) 10(6) *Earth's Future* e2021EF002612 <https://doi.org/10.1029/2021EF002612>

source (including the share of renewable and low-carbon energy and any applicable minimum green energy thresholds); (ii) location-based and, where relevant, market-based carbon dioxide and other greenhouse gas emissions, reported in accordance with recognised standards; (iii) water withdrawals and consumption; (iv) key information on hardware turnover, reuse and recycling; and (v) the proportion of data centre and comparable AI-computing capacity allocated to different categories of use cases, including, insofar as reasonably assessable, whether such uses are predominantly environmentally harmful or beneficial. Ideally, these disclosures should be subject to independent audit and use harmonised methodologies, so that figures are comparable across firms and over time. Without such mandatory transparency, further policy instruments cannot be designed in a manner that is both effective and fair.

While sustainable design requirements, impact assessments and transparency obligations are approaches widely discussed and partially adopted in various jurisdictions including the EU, they have proven insufficient for aligning AI transformation with planetary goals, as they remain largely procedural and vulnerable to rebound effects.⁵¹ The final section therefore discusses a different approach that aims to democratise compute governance and contribute to building international computational rules and structures capable of constraining the sector's environmental footprint in a manner consistent with planetary boundaries.

5. Regulating the Use Cases

Going beyond current approaches, regulating the use cases of AI from an ecological standpoint would constitute a fundamental shift for mitigating environmental, planetary and social harms while better targeting potential benefits.

An interesting cost-benefit approach has been discussed by Philipp Hacker, in which regulators would classify AI applications into broad “social usefulness classes” (high, medium and low) according to their expected societal utility, and then link those classes to differentiated “consumption caps” on the energy use and associated greenhouse gas emissions of model training and deployment.⁵² Although Hacker’s proposals are developed with the EU

⁵¹ Jessica Commins and Kristina Irion, *Towards Planet Proof Computing: Law and Policy of Data Centre Sustainability in the European Union*, Technology and Regulation, 2025, 1-36 • <https://doi.org/10.71265/c1nnwh92> • ISSN: 2666-139X

⁵² Philipp Hacker, ‘Sustainable AI Regulation’ (2023) arXiv preprint 2306.00292, 3–4, 21–22, 25–28 <https://arxiv.org/abs/2306.00292>

legal framework in mind, the same regulatory approach could in principle be applied at a global level. In this scheme, sectors such as medicine, employment, public administration and transport would typically fall into the high-benefit category, while applications in entertainment, marketing and advertising would generally be treated as low-benefit, with most remaining use cases occupying a medium-benefit middle ground. These value judgements are, he concedes, inherently contestable, yet he argues that societies facing a climate emergency cannot avoid deciding how much energy and carbon they are prepared to “spend” on different AI uses, and must therefore confront such prioritisation explicitly in AI governance. This tripartite benefit classification could also be normatively refined by designating projects that verifiably reduce environmental or climate impacts, such as the ones mentioned above in chapter 1, as paradigmatic “high-benefit” uses, which would receive more generous resource caps than entertainment, advertising and other comparatively low-benefit applications.

A further policy option is to impose a public service obligation on large data centre operators or cloud providers, requiring that a defined, auditable share of accelerator capacity (for example GPU hours) be reserved for independently certified “AI for good” applications in domains such as climate modelling, public health, and other public interest research. To be administrable and non symbolic, the obligation would need precise eligibility criteria, transparent allocation and audit mechanisms, and safeguards ensuring additionality and consistency with binding energy and emissions constraints.

A significant challenge to adopting such policy frameworks is that industry leaders are actively monetizing generative AI through low-benefit AI use cases. For example, OpenAI announced that, from December 2025 onwards, ChatGPT will allow verified adult users to generate erotic content as part of a broader relaxation of its mature-content rules.⁵³ Such policies are likely to increase engagement and usage time, with attendant risks of reinforcing problematic or compulsive patterns of use for some users. By driving up usage in this way, they could entail substantial inference costs if widely used, while their broader societal value remains highly contested among policymakers, experts and civil-society organisations. Within benefit-sensitive regulatory frameworks, such use cases would typically qualify as

⁵³ BBC News, ‘ChatGPT will soon allow erotica for verified adults, says OpenAI boss’ (15 October 2025)

low-benefit AI and, in certain instances, as net-harmful, yet their profitability generates a strong resistance to more stringent environmental governance.

Another major challenge is the lack of effective multilateral coordination and cooperation in the governance of AI, which contributes to an international arms race dynamic. As states and corporations compete to lead in frontier AI development without a level playing field, there is a growing risk of a regulatory race to the bottom, where environmental standards are weakened or overlooked to accelerate innovation and deployment. This competitive pressure undermines efforts to align AI infrastructure with climate goals and reinforces short-term incentives that prioritise scale and performance over sustainability and long-term responsibility.

Regardless of the challenges involved in implementing the proposed regulatory measures, the question of which AI applications are normatively defensible in light of their environmental and climatic impacts must be addressed through global deliberation. In an era of accelerating AI transformation, the central issue is not only what can be developed or deployed but what should be developed, once planetary costs are taken into account. Information processing is a materially grounded activity that consumes energy, raw materials, as well as ecological and scientific capacity, and therefore cannot be considered environmentally neutral. When certain AI use cases impose significant environmental burdens while offering limited or uncertain social value, a benefit-sensitive and climate-conscious governance framework requires direct engagement with a simple but critical question: is it worth it?