



# Artificial Intelligence and Environmental Sustainability: The Need for a Regulatory Framework

Dhairya Jaiswal<sup>1</sup>, Kavi Shankar<sup>2</sup>, Santhosh V<sup>3</sup>, Sravondev P A<sup>4</sup>

<sup>1</sup> Research Scholar, Faculty of Law, Banaras Hindu University. Email- dhairya.jaiswal31@gmail.com, Orcid id- <https://orcid.org/0009-0002-8967-226X>

<sup>2</sup>Asst. Professor of Law, VIT School of Law, VIT, Chennai. Email- office.kavishankar@gmail.com

<sup>3</sup>Asst. Professor of Law, VIT School of Law, VIT, Chennai. Email- santhoshveeramani@protonmail.com

<sup>4</sup>Advocate, Kozhikode, Kerala. Email- advsravondev@gmail.com

## Abstract

Artificial intelligence (AI) has in no time become both an ally and a quiet threat to the planet. On one side, AI is helping scientists forecast the climate, track pollution, protect wildlife, and run power grids more efficiently. On the other side, the data centres that train and run AI consume large amounts of electricity and freshwater, produce carbon emissions, and add to a growing pile of electronic waste. This paper looks at both sides of that story. It reviews what researchers have learned so far about AI's environmental footprint, summarises the main findings, and asks whether our existing environmental, technology, and data laws are ready to govern AI responsibly. It then turns to the Indian context, where a rapid build-out of AI data centres is meeting a power system that still leans on coal and a regulatory regime designed for factories rather than server farms. The paper argues for a regulatory framework built on transparency, mandatory environmental reporting, climate-justice safeguards, and lifecycle accountability, so that AI can advance sustainability without quietly undermining it.

**Keywords:** artificial intelligence, environmental sustainability, climate justice, AI governance, India

## 1. Introduction

For most of its history, artificial intelligence was treated as a purely digital subject, something that lived inside computers and had little to do with rivers, forests, or the air we breathe. That view no longer holds. As AI systems have grown larger and spread into almost every sector, it has become clear that they have a physical body made of chips, cables, cooling towers and electricity, and that this body leaves a real mark on the environment (Crawford, 2021). At the same time, AI has emerged as one of the most powerful tools available for understanding and fighting environmental problems. The relationship between AI and sustainability is therefore best described not as good or bad, but as deeply two-sided.

On the positive side, AI can support almost every environmental goal we care about. Machine-learning models improve climate forecasts, help manage water and energy more carefully, monitor air and water pollution in real time, and assist conservationists in tracking endangered species (Rolnick et al., 2022). In a widely cited study, Vinuesa et al. (2020) found that AI could help achieve 134 of the 169 targets that sit under the United Nations Sustainable Development Goals (SDGs). That is a remarkable figure, and it explains much of the excitement around "AI for good." But the same study warned that AI could also work against 59 of those targets and later research has shown why (Vinuesa et al., 2020). Training a single large model can consume enormous amounts of energy and release significant carbon dioxide (Strubell et al., 2019). The data centres that host these systems drink freshwater for cooling, and the rapid turnover of specialised hardware creates electronic waste (Li et al., 2023; Wang et al., 2024). These costs fall unevenly, often on communities that gain the least from the technology. This is why scholars now distinguish between "AI for sustainability" and the "sustainability of AI" itself, arguing that we must attend to both (van Wynsberghe, 2021).

This paper examines that tension and asks a governance question at its heart: are our current environmental, technology, and data laws equipped to manage AI's environmental impacts, and if not, what should a better framework look like? The discussion proceeds in six steps. It first sets out the main issues and environmental costs. It then reviews the research done so far and draws out the key findings. After that it turns to the Indian perspective, where these questions are becoming urgent, before closing with a set of practical recommendations for a regulatory framework that is environmentally responsible, fair and consistent with broader climate goals.

## 2. The Core Issues: AI's Environmental Promise and Its Hidden Costs

The first issue is the genuine promise of AI for the environment. Used well, AI helps us see patterns in huge datasets that no human team could process. It sharpens weather and climate models, supports precision agriculture that uses

less water and fewer chemicals, optimises electricity grids so that renewable power is not wasted, and flags illegal logging or poaching from satellite images (Rolnick et al., 2022; Cowls et al., 2023). In each of these cases, AI is not the solution by itself, but it makes existing environmental efforts faster, cheaper, and more accurate.

The second issue is the energy appetite of AI. The infrastructure behind modern AI, the data centres, has become a serious new actor in the global energy system. According to the International Energy Agency (2025), data centres used roughly 415 terawatt-hours of electricity in 2024, about 1.5% of the world's total, and this could roughly double to around 945 terawatt-hours by 2030 as AI workloads expand. The agency noted that electricity demand from AI-focused data centres in particular surged by about 50% in 2025 (International Energy Agency, 2025). When much of that power still comes from fossil fuels, more computation means more carbon emissions (Kaack et al., 2022).

The third issue is water. Cooling thousands of hot processors requires large volumes of clean water, a cost that stayed hidden for years. Li et al. (2023) estimated that training the GPT-3 model in state-of-the-art United States data centres could directly evaporate around 700,000 litres of freshwater, and they projected that global AI demand might account for 4.2 to 6.6 billion cubic metres of water withdrawal by 2027. In regions already facing water stress, this is not a trivial concern.

The fourth issue is electronic waste. AI relies on specialised chips that are replaced quickly as newer, faster hardware arrives. Wang et al. (2024) modelled this stream and found that generative AI alone could add between 1.2 and 5.0 million tonnes of e-waste between 2020 and 2030. Encouragingly, they also showed that circular-economy strategies i.e. reusing, refurbishing, and recycling hardware, could cut that waste by 16% to 86%, which makes management an urgent priority rather than an afterthought.

Closely linked to energy is the question of carbon emissions and a subtler danger known as the rebound effect. The carbon cost of a model depends heavily on where and when it is trained, because a data centre running on coal power emits far more than one running on hydro or solar (Patterson et al., 2021; Dhar, 2020). More worrying is the possibility that efficiency gains simply encourage more use. As models become cheaper to run, they tend to be deployed everywhere, so total consumption can climb even as each individual task becomes greener. Dauvergne (2020) cautions that, without limits, AI can intensify the very patterns of production and consumption that drive environmental damage, meaning that technical efficiency alone will not guarantee a smaller footprint.

It is also worth placing AI's electronic waste within a much larger crisis. The world already generates tens of millions of tonnes of e-waste each year, and only a small fraction is properly recycled (International Telecommunication Union & United Nations Institute for Training and Research, 2024). AI hardware adds a fast-moving new stream to this pile, made up of high-value but hard-to-recycle components. Bender et al. (2021), writing about the risks of ever-larger language models, argued that the field's race toward scale concentrates benefits among well-resourced actors while pushing environmental and social costs outward, a concern that applies directly to the hardware lifecycle as well as to energy.

The fifth issue is fairness, and it ties the others together. The benefits of AI flow mostly to wealthy companies and countries, while many of the costs, emissions, water use, polluting waste sites, and the mining of rare materials, are felt elsewhere, often by poorer communities with little say in the matter (Coeckelbergh, 2021; Brevini, 2021). Robbins and van Wynsberghe (2022) warn that building ever-larger AI infrastructure risks locking society into an unsustainable path that becomes hard to reverse. This raises questions of climate justice that a purely technical view of AI tends to ignore.

Finally, there is the issue of opacity. Companies are generally not required to publish how much energy or water a given model consumes, so the true ecological cost stays buried inside corporate sustainability reports that are hard to verify (Henderson et al., 2020). Without reliable, comparable data, regulators and citizens cannot hold anyone accountable. This information gap is, in many ways, the root problem that any good regulatory framework must address.

### 3. The State of Research and Policy

Research on AI and the environment has matured rapidly over the past few years, moving from scattered observations to a recognisable field. An early turning point came when Strubell et al. (2019) measured the carbon footprint of training large language models and showed that the numbers were far higher than most people assumed. Their work pushed the research community to start reporting the energy and emissions behind AI experiments, an idea developed further by Henderson et al. (2020) and by tools for estimating emissions proposed by Lacoste et al. (2019).

From there, two related research streams grew. The first, often called "Green AI," argues that efficiency should be treated as a core goal of AI research, not just accuracy (Schwartz et al., 2020). Studies in this stream have measured the carbon cost of specific models, such as the BLOOM language model (Luccioni et al., 2023), and have catalogued practical ways to reduce AI's footprint, from better hardware to smaller models (Wu et al., 2022; Verdecchia et al., 2023). The second stream, "AI for sustainability," maps where AI can actually help meet environmental goals, with Rolnick et al. (2022) providing an influential roadmap across sectors such as energy, transport, and agriculture.

A separate body of work has tried to weigh these two sides against each other. Vinuesa et al. (2020) produced the most cited attempt, using expert judgement to estimate AI's positive and negative effects on the SDGs. Kaack et al. (2022) refined the picture for climate specifically, distinguishing between AI's direct effects (the emissions from

running it), its immediate application effects (helping cut emissions in other sectors), and its broader system-level effects (changing how economies behave). Their framework is now widely used because it shows that AI's net climate impact depends heavily on how and where it is deployed.

Alongside the technical literature, ethical and conceptual literature has taken shape. Van Wynsberghe (2021) gave the field a useful vocabulary by separating "AI for sustainability" from the "sustainability of AI," and later work has pressed researchers to define these terms more precisely (Falk & van Wynsberghe, 2023; Rohde et al., 2024). Scholars such as Crawford (2021) and Brevini (2021) have placed AI within a larger story about extraction, power, and the planetary costs of computing, while Coeckelbergh (2021) and Sætra (2021) have foregrounded justice and the limits of treating AI as a neutral fix.

On the governance side, progress has been slower and more fragmented. Early management scholarship, such as Nishant et al. (2020), set out a research agenda for using AI to advance sustainability while flagging the risks that come with it, but it stopped short of detailed regulation. Much of the early ethical guidance, such as the AI4People framework, focused on principles like transparency and accountability without saying much about environmental harm (Floridi et al., 2018). More recent contributions argue that AI governance and environmental governance must be brought together. Truby (2020) proposed steering AI development toward the SDGs through regulation, and Galaz et al. (2021) highlighted the systemic environmental risks that fall through the cracks between technology policy and climate policy. International bodies have begun to take notice as well; the United Nations Environment Programme (2024) has urged governments to standardise the measurement of AI's footprint and to weave environmental safeguards into AI policy. There is now broad agreement in the literature that voluntary corporate pledges are not enough and that mandatory measurement and disclosure are the missing foundation (Henderson et al., 2020; Cowls et al., 2023).

#### 4. Findings

Pulling these threads together, several clear findings emerge. First, AI's environmental impact is real and growing, but it is also concentrated. A relatively small number of very large models and the data centres that host them account for a large share of the energy, water, and hardware demand (International Energy Agency, 2025; Luccioni et al., 2023). This is actually good news for governance, because it means that targeted rules aimed at large operators could address most of the impact without burdening every small developer.

Second, the costs are spread across the whole lifecycle, not just one stage. People tend to focus on the electricity used to train a model, but emissions and resource use occur at every step: mining materials for chips, manufacturing hardware, running models for years of everyday use (inference), cooling the facilities, and finally disposing of obsolete equipment (Wang et al., 2024; Crawford, 2021). In fact, recent work shows that the day-to-day use of models, rather than their one-time training, may dominate the total footprint over a system's lifetime, because popular models answer billions of queries (Luccioni et al., 2024). Any honest accounting, and any serious regulation, has to cover the full lifecycle rather than a single eye-catching number.

Third, AI's benefits for the environment are real but conditional. AI does not automatically reduce emissions; it does so only when it is deliberately pointed at the right problems and powered by clean energy (Kaack et al., 2022; Rolnick et al., 2022). Used carelessly, the same technology can speed up resource extraction or encourage more consumption, cancelling out its gains. The net effect depends on human choices and policy, not on technology alone.

Fourth, the central obstacle to good governance is missing information. Because companies are rarely required to disclose model-level energy and water use, neither regulators nor the public can verify claims or compare options (Henderson et al., 2020). This data gap weakens every other policy tool, since one cannot regulate what one cannot measure.

Fifth, the burdens and benefits of AI are distributed unequally, raising genuine climate-justice concerns. Vulnerable communities and water-stressed or coal-dependent regions can bear real costs while having little influence over the decisions that create them (Coeckelbergh, 2021; Robbins & van Wynsberghe, 2022). A framework that ignores this risks being efficient on paper but unjust in practice.

Sixth, the existing legal toolkit is poorly matched to the problem. Environmental laws were largely written for visible, point-source polluters like factories and mines, technology and data laws focus on privacy and competition rather than ecology, and climate policy rarely names AI explicitly. The result is a set of frameworks that, taken individually, leave AI's environmental footprint only lightly governed (Galaz et al., 2021; Truby, 2020).

#### 5. The Indian Perspective

India sits at a difficult crossroads in this debate. The country has embraced AI as an engine of growth and development. Its national AI strategy, branded #AIForAll, framed AI as a tool to improve healthcare, agriculture, education, and inclusive development (NITI Aayog, 2018), and the more recent IndiaAI Mission has committed substantial public funding to build computing capacity, deploy thousands of graphics processing units, and support home-grown models (Ministry of Electronics and Information Technology, 2024). India also hosted a major global AI summit and has positioned itself as a voice for the Global South in shaping how AI develops.

Crucially, India is not only a consumer of AI's costs but also a beneficiary of its environmental promise. Indian agencies and researchers are already using AI to sharpen monsoon and disaster forecasting, to monitor air pollution and greenhouse-gas emissions, to map deforestation and biodiversity loss, and to support precision farming that cuts water use and methane from rice cultivation (Rolnick et al., 2022; Vinuesa et al., 2020). For a country that is highly exposed to climate impacts and that must feed a vast population on stressed land and water, these applications are not luxuries; they are tools for resilience. The challenge is that the same digital build-out that enables them also generates the emissions, water demand, and waste described above, so India experiences both halves of the AI–environment relationship at once.

This ambition comes with a heavy infrastructure build-out. India's data centre capacity, around 1.5 gigawatts in 2025, is projected to expand several-fold by 2030 as AI demand grows, and one industry analysis estimated that AI-linked data centres alone could require an additional 40 to 45 terawatt-hours of electricity by 2030 (Deloitte, 2025). The environmental catch is that India's electricity still comes mostly from coal, so a surge in AI computing can translate fairly directly into higher carbon emissions unless clean energy keeps pace. Because AI data centres need round-the-clock power, they may even encourage additional fossil-fuel capacity to guarantee reliability.

Water is an equally pressing concern in the Indian setting. Many of the cities competing to host data centres are already water-stressed, yet cooling these facilities can consume large volumes of freshwater. This creates a direct tension between digital ambition and the everyday water needs of local communities, and it has prompted calls for facilities to rely on treated wastewater, closed-loop cooling, and zero-liquid-discharge systems rather than potable supplies.

India's regulatory framework, however, was not designed with any of this in mind. The Environment (Protection) Act, 1986 and the country's environmental impact assessment process were built around traditional industries such as mining and manufacturing, and they do not clearly capture large, cloud-based computing facilities. As a result, GPU clusters with a substantial energy and water footprint can be set up without the kind of environmental scrutiny applied to a comparable factory. On the data side, the Digital Personal Data Protection Act, 2023 governs how personal data is handled but says nothing about the ecological cost of the systems that process it (Ministry of Electronics and Information Technology, 2023). The most directly relevant rule is arguably the E-Waste (Management) Rules, 2022, which require electronic waste to be channelled to registered recyclers (Ministry of Environment, Forest and Climate Change, 2022); applied seriously to AI hardware, these rules could help, but they were not written with rapid server turnover in mind.

There is, then, a regulatory gap in India that mirrors the global one but with sharper stakes. The country has strong reasons to pursue AI and equally strong reasons to protect its energy security, scarce water, and vulnerable populations. India's own tradition of environmental thought, captured in ideas like treating the Earth as one family, and its constitutional commitments to environmental protection, give it a values-based foundation for insisting that digital progress not come at the planet's expense. The opportunity for India is to design “green AI” rules early, before the infrastructure locks in, and in doing so to offer a model that other developing countries can follow.

## **6. The Way Forward: Toward a Responsible Regulatory Framework**

The findings point toward a regulatory framework with a few clear pillars. The first and most important is transparency. Large AI operators should be required to measure and publicly report the energy, carbon, and water footprint of training and running their major models, using standardised methods so the numbers can be compared (Henderson et al., 2020; Cowsls et al., 2023). Standard metrics such as power usage effectiveness and water usage effectiveness for data centres can anchor this reporting. Mandatory disclosure is the foundation on which every other rule depends, because it closes the information gap identified earlier.

The second pillar is lifecycle accountability. Rules should cover the full journey of AI hardware and computation, from the sourcing of materials and manufacturing of chips, through years of operation, to end-of-life recycling. Extended-producer-responsibility schemes and circular-economy requirements can push operators to reuse and refurbish hardware, which research suggests can dramatically cut e-waste (Wang et al., 2024; Wu et al., 2022). Tying data-centre approvals to credible clean-energy commitments would address the emissions that arise during operation. The third pillar is to bring AI inside existing environmental governance rather than inventing a wholly separate regime. Environmental impact assessment processes can be updated so that large data centres are evaluated for their energy and water demand, much as a factory would be. Where dedicated AI laws or climate laws exist, they should explicitly name and account for the environmental footprint of computing, closing the gap between technology policy and climate policy (Truby, 2020; Galaz et al., 2021).

The fourth pillar is climate justice. A responsible framework must check who bears the costs and who reaps the benefits. This means protecting water-stressed communities from having scarce supplies diverted to cooling, ensuring that polluting facilities and waste sites are not simply pushed onto the least powerful, and giving affected communities a real voice in siting decisions (Coeckelbergh, 2021; Robbins & van Wynsberghe, 2022). Fairness should be treated as a design requirement, not an afterthought.

The fifth pillar is to steer AI deliberately toward sustainability. Because AI's benefits are conditional, governments can use funding, procurement, and incentives to direct AI talent and computing toward genuine environmental problems, grid optimisation, climate modelling, pollution monitoring, and conservation, while discouraging wasteful

or purely extractive uses (Rolnick et al., 2022; Kaack et al., 2022). Public money should reward efficiency and real-world impact rather than sheer model size.

For India specifically, the way forward is to act early and turn constraint into leadership. India can require new data centres to move toward renewable power and treated-wastewater cooling, update its environmental assessment rules to cover computing infrastructure, enforce the E-Waste Rules rigorously for AI hardware, and embed environmental audits within the IndiaAI Mission itself. By doing so while its AI build-out is still young, India can avoid locking in an unsustainable path and can credibly lead the Global South toward a model of AI development that is both ambitious and ecologically responsible (Deloitte, 2025; NITI Aayog, 2018).

## 7. Conclusion

Artificial intelligence is neither a saviour nor a villain in the environmental story; it is a powerful tool whose effect depends entirely on how we build, power, and govern it. The evidence shows that AI can meaningfully advance sustainability across many sectors, but that it also carries real and rising costs in energy, water, carbon, and electronic waste, costs that are often hidden and unequally shared. Our current environmental, technology, and data laws were not designed for this challenge and leave important gaps.

Closing those gaps does not require us to choose between innovation and the environment. With transparency at its core, supported by lifecycle accountability, integration into existing environmental rules, climate-justice safeguards, and the deliberate steering of AI toward green ends, a regulatory framework can let societies enjoy AI's benefits while keeping its footprint in check. For India and other rapidly developing countries, acting now, before the infrastructure hardens, offers a rare chance to get the balance right from the start, ensuring that the age of artificial intelligence is also an age of environmental responsibility.

## References

- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 610–623). Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>
- Brevini, B. (2021). *Is AI good for the planet?* Polity Press.
- Coeckelbergh, M. (2021). AI for climate: Freedom, justice, and other ethical and political challenges. *AI and Ethics*, 1(1), 67–72. <https://doi.org/10.1007/s43681-020-00007-2>
- Cowls, J., Tsamados, A., Taddeo, M., & Floridi, L. (2023). The AI gambit: Leveraging artificial intelligence to combat climate change-opportunities, challenges, and recommendations. *AI & Society*, 38(1), 283–307. <https://doi.org/10.1007/s00146-021-01294-x>
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Dauvergne, P. (2020). *AI in the wild: Sustainability in the age of artificial intelligence*. MIT Press.
- Deloitte. (2025). *Attracting AI data centre infrastructure investment in India*. Deloitte India.
- Dhar, P. (2020). The carbon impact of artificial intelligence. *Nature Machine Intelligence*, 2(8), 423–425. <https://doi.org/10.1038/s42256-020-0219-9>
- Falk, S., & van Wynsberghe, A. (2023). Challenging AI for sustainability: What ought it mean? *AI and Ethics*. Advance online publication. <https://doi.org/10.1007/s43681-023-00323-3>
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- Galaz, V., Centeno, M. A., Callahan, P. W., Causevic, A., Patterson, T., Brass, I., Baum, S., Farber, D., Fischer, J., Garcia, D., McPhearson, T., Jimenez, D., King, B., Söderberg, P., & Levy, K. (2021). Artificial intelligence, systemic risks, and sustainability. *Technology in Society*, 67, 101741. <https://doi.org/10.1016/j.techsoc.2021.101741>
- Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D., & Pineau, J. (2020). Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research*, 21(248), 1–43.
- International Energy Agency. (2025). *Energy and AI*. IEA. <https://www.iea.org/reports/energy-and-ai>
- International Telecommunication Union & United Nations Institute for Training and Research. (2024). *The global e-waste monitor 2024*. ITU & UNITAR.
- Kaack, L. H., Donti, P. L., Strubell, E., Kamiya, G., Creutzig, F., & Rolnick, D. (2022). Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*, 12(6), 518–527. <https://doi.org/10.1038/s41558-022-01377-7>
- Lacoste, A., Luccioni, A., Schmidt, V., & Dandres, T. (2019). Quantifying the carbon emissions of machine learning. *arXiv*. <https://arxiv.org/abs/1910.09700>
- Li, P., Yang, J., Islam, M. A., & Ren, S. (2023). Making AI less “thirsty”: Uncovering and addressing the secret water footprint of AI models. *arXiv*. <https://arxiv.org/abs/2304.03271>

18. Luccioni, A. S., Jernite, Y., & Strubell, E. (2024). Power hungry processing: Watts driving the cost of AI deployment? In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (pp. 85–99). Association for Computing Machinery. <https://doi.org/10.1145/3630106.3658542>
19. Luccioni, A. S., Viguier, S., & Ligozat, A.-L. (2023). Estimating the carbon footprint of BLOOM, a 176B parameter language model. *Journal of Machine Learning Research*, 24(253), 1–15.
20. Ministry of Electronics and Information Technology. (2023). *The Digital Personal Data Protection Act, 2023*. Government of India.
21. Ministry of Electronics and Information Technology. (2024). *IndiaAI Mission*. Government of India. <https://indiaai.gov.in/>
22. Ministry of Environment, Forest and Climate Change. (2022). *E-Waste (Management) Rules, 2022*. Government of India.
23. NITI Aayog. (2018). *National strategy for artificial intelligence: #AIForAll*. Government of India.
24. Nishant, R., Kennedy, M., & Corbett, J. (2020). Artificial intelligence for sustainability: Challenges, opportunities, and a research agenda. *International Journal of Information Management*, 53, 102104. <https://doi.org/10.1016/j.ijinfomgt.2020.102104>
25. Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon emissions and large neural network training. *arXiv*. <https://arxiv.org/abs/2104.10350>
26. Robbins, S., & van Wynsberghe, A. (2022). Our new artificial intelligence infrastructure: Becoming locked into an unsustainable future. *Sustainability*, 14(8), 4829. <https://doi.org/10.3390/su14084829>
27. Rohde, F., Wagner, J., Meyer, A., Reinhard, P., Voss, M., Petschow, U., & Mollen, A. (2024). Broadening the perspective for sustainable artificial intelligence: Sustainability criteria and indicators for artificial intelligence systems. *Current Opinion in Environmental Sustainability*, 66, 101411. <https://doi.org/10.1016/j.cosust.2023.101411>
28. Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A. S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., Luccioni, A. S., Maharaj, T., Sherwin, E. D., Mukkavilli, S. K., Kording, K. P., Gomes, C. P., Ng, A. Y., Hassabis, D., Platt, J. C., ... Bengio, Y. (2022). Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2), Article 42. <https://doi.org/10.1145/3485128>
29. Sætra, H. S. (2021). AI in context and the sustainable development goals: Factoring in the unsustainability of the sociotechnical system. *Sustainability*, 13(4), 1738. <https://doi.org/10.3390/su13041738>
30. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>
31. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 3645–3650). Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1355>
32. Truby, J. (2020). Governing artificial intelligence to benefit the UN Sustainable Development Goals. *Sustainable Development*, 28(4), 946–959. <https://doi.org/10.1002/sd.2048>
33. United Nations Environment Programme. (2024). *AI has an environmental problem. Here's what the world can do about that*. UNEP.
34. van Wynsberghe, A. (2021). Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics*, 1(3), 213–218. <https://doi.org/10.1007/s43681-021-00043-6>
35. Verdecchia, R., Sallou, J., & Cruz, L. (2023). A systematic review of Green AI. *WIREs Data Mining and Knowledge Discovery*, 13(4), e1507. <https://doi.org/10.1002/widm.1507>
36. Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M., & Fuso Nerini, F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11, Article 233. <https://doi.org/10.1038/s41467-019-14108-y>
37. Wang, P., Zhang, L.-Y., Tzachor, A., & Chen, W.-Q. (2024). E-waste challenges of generative artificial intelligence. *Nature Computational Science*, 4(11), 818–823. <https://doi.org/10.1038/s43588-024-00712-6>
38. Wu, C.-J., Raghavendra, R., Gupta, U., Acun, B., Ardalani, N., Maeng, K., Chang, G., Aga, F., Huang, J., Bai, C., Gschwind, M., Gupta, A., Ott, M., Melnikov, A., Candido, S., Brooks, D., Chauhan, G., Lee, B., Lee, H.-H. S., ... Hazelwood, K. (2022). Sustainable AI: Environmental implications, challenges and opportunities. *Proceedings of Machine Learning and Systems*, 4, 795–813.