Climate change and health: the next challenge of ethical AI

Amelia Fiske, Isabella M Radhuber, Theresa Willem, Alena Buyx, Leo Anthony Celi, Stuart McLennan

Artificial intelligence (AI) is one of the world's most resource-intensive digital technologies, but the environmental impact of AI on health remains largely unaddressed in both global health and bioethics. Effects on the environment have, thus far, been understood as a subsidiary consideration in AI ethics and rarely considered as a key ethical concern. AI technologies exacerbate climate change and sociopolitical instability due to their intensive use of natural resources and energy resources linked to the training and deployment of algorithmic systems. In global health, this intensive resource use is particularly concerning, given the explicit emphasis on improving health and advancing equity across the world. To address this, we interrogate how the inclusion of AI's environmental impact necessarily reshapes established ethical commitments in AI ethics frameworks and propose concrete strategies for accountability in the area of global health. This approach includes building a culture of intentional AI, for example through improved reporting, auditing, and intranational cooperation, in order to better align AI development and AI ethics with critical climate goals.

Lancet Glob Health 2025; 13: e1314–20

Institute of History and Ethics in Medicine, Department of Preclinical Medicine, TUM School of Medicine and Health, Technical University of Munich, Munich, Germany (A Fiske PhD, T Willem MA, Prof A Buyx MD, S McLennan PhD); **Equity and Justice Research Group, Population and Just Societies Program, International Institute for Applied Systems Analysis, Laxenburg, Austria** (I M Radhuber PhD); **Department of Political Science, University of Vienna, Vienna, Austria** (I M Radhuber); **Laboratory for Computational Physiology, Massachusetts Institute of Technology, Cambridge, MA, USA** (Prof L A Celi MD); **Division of Pulmonary, Critical Care and Sleep Medicine, Beth Israel Deaconess Medical Center, Boston, MA, USA** (Prof L A Celi); **Department of Biostatistics, Harvard T H Chan School of Public Health, Boston, MA, USA** (Prof L A Celi)

Correspondence to: Dr Amelia Fiske, Institute of History and Ethics in Medicine, Department of Preclinical Medicine, TUM School of Medicine and Health, Technical University of Munich, 81675 Munich, Germany **a.fiske@tum.de**

Introduction

Artificial intelligence (AI) is one of the world's most resource-intensive digital technologies.[1] Nonetheless, the major concerns raised by the environmental impact of AI on health have not been sufficiently addressed in the fields of global health or bioethics. AI ethics frameworks have, to date, primarily focused on the core concerns of transparency, justice and fairness, non-maleficence, responsibility, and privacy, whereas the environmental impact of AI has been relegated to a subsidiary consideration. When included, the environmental impact of AI often falls under the somewhat under-represented ethical concern of sustainability—it is rarely considered as a key ethical concern in its own right.[2] Given the established health burden of climate change, in this Viewpoint, we point to a central, unaddressed tension: the AI solutions proposed to address global health concerns are exacerbating the very health problems they purport to address. AI cannot be a force for good in global health until the adverse effects of AI on climate and health are addressed. Centring the environmental impact of AI on health in the agendas of both global health and the bioethics and technology sector requires refiguring the core principles of AI ethics frameworks and fostering proactive approaches for AI that prioritise reducing environmental impact.

The promise and trade-offs of AI in global health

AI is promised to be transformative for global health, with the potential to reduce some of the world's worst health inequalities, such as unnecessary child mortality from preventable causes. In 2023, the Bill & Melinda Gates Foundation announced the launch of a new US$5 million fund to establish AI large language models in low-income and middle-income countries.[3] Similarly, the World Bank has highlighted the promise of digital health, including AI, to bring innovations to underserved groups.[4] Research in the field of global health includes a wide range of potential applications, including the use of AI to assist medical practitioners in diagnosis, assessments of morbidity and mortality risk, disease surveillance, and health policy and planning.[5] The hope is that the deployment of AI could help to improve health outcomes and address pressing health concerns particular to low-income and middle-income countries.

Regulatory guidance from WHO and other high-level regulatory acts, along with scoping reviews, have identified the ethical challenges and risks of AI in health.[6–10] Yet the major focus of formal guidance, regulation, and ongoing investment continues to neglect the central ethical paradox that, because of the environmental impacts of AI, AI-based solutions are contributing to the global health problems they are intended to address. AI technologies exacerbate climate change and sociopolitical instability due to their intensive use of natural and energy resources linked to the training and deployment of algorithmic systems (panel).[14]

The environmental health burdens of AI are particularly relevant to the field of global health, a field born out of colonial encounters that continue to shape research and investment today.[15,16] For many global majority countries, the first encounters with biomedicine were with colonial medicine, whereby health officials and programmes were an extension of oppressive political and economic powers that also established extractive environmental industries.[16] Dramatic global power imbalances in the political economy of AI are linked to historical patterns of colonialism, including the ongoing extraction of labour, data, and profit from formerly colonised countries, largely for the benefit of technology companies located in the global minority.[17,18] Drawing on work by Adan and colleagues,[17] we use the term global minority to refer to the increasing concentration of access and ability to shape AI development and governance in a few countries with advanced AI capabilities. We use the term global majority to refer to the exclusion of most of the globe—the populations of those living in countries primarily located in Africa, Latin America, south and southeast Asia, and parts of eastern Europe—from the ability to shape AI development or reap its benefits. These different geopolitical and socioeconomic histories shape the ways in which the promises and trade-offs of

Panel: **Water conflicts, health concerns, and data centres in Querétaro, Mexico**

The construction of data centres in Querétaro, Mexico, highlights substantial environmental and health concerns linked to artificial intelligence, particularly in the context of increasing droughts and water scarcity. These facilities require vast amounts of water for cooling—an estimated 15 billion litres annually, about 13% of the metropolitan area's water use. One result has been intensifying competition for water between industry, tourism, and local communities. In Mexico, approximately 77% of the country's electricity comes from fossil fuels. As such, the high energy demand of these centres also exacerbates carbon emissions. Some of the technology companies involved promote data centres as bringing potential economic benefits. Google, for instance, projected that data centres would contribute US$11·2 billion to gross domestic product and create more than 117 000 jobs by 2030. Google, Amazon, and others also claim to pursue water and energy sustainability. However, there has been little transparency regarding resource usage. Experts warn that even if water reuse is feasible, the energy-intensive cooling systems will continue to drive fossil fuel consumption, worsening climate-related health risks.[11–13]

AI are realised in global health, including the distribution of environmental health burdens and benefits.

Mounting environmental impact of AI technologies

The rise of large language models and generative AI in recent years has been linked to alarming environmental trends in relation to water consumption, energy use, and raw material extraction.[19] A preprint article showed that immense quantities of water are required for both on-site cooling and off-site electricity generation for the data centres in which AI models are trained and deployed.[20] In 2022, Microsoft reported consuming 6·4 million cubic meters of water, a 34% increase over the previous year; presumably, most of this additional water was used for cloud data centres.[21] As estimated in a preprint paper,[20] globally, the demand for AI could result in 4·2–6·6 billion m³ of water being used in 2027, which would be more than the total water use of a country such as Denmark. When looked at from a global health perspective, the disparities are ever more striking: the average annual water usage by a cooling system for a small data centre located in the USA is 2000 times greater than that of an individual living in Nigeria.[22]

AI also requires substantial energy resources. According to a a preprint paper on the Carbontracker tool, the energy costs of deep learning increased 300 000-fold between 2012 and 2018.[23] In another study assessing the energy use required by AI, middle range approximations suggested that the electricity demand of new AI servers sold in the year 2027 alone could be as much as the entire annual consumption of countries such as Argentina, the Netherlands, or Sweden.[24,25] Increasing energy consumption points to a growing carbon footprint for AI technologies. In one study, the training of a single big language model was found to generate emissions approximately equivalent to 300 000 kg of CO_2.[26] In some cases, this growing energy consumption has meant the revival of highly polluting ways of producing energy. For example, in Nebraska (USA), one power company reported that it would need to abandon plans to stop burning coal in order to produce electricity to meet the demands of nearby Google and Meta data centres.[27,28]

The push for AI also results in the growing consumption of raw materials, including coal, uranium, cobalt, lithium, and coltan. These materials are necessary for training AI systems and building the computers, smartphones, and infrastructure (eg, undersea cables) that many AI products require to operate, harvest data, or connect to the internet.[18,29] The production of electronics for AI technologies also generates hazardous electronic waste, such as mercury and lead. The manifold environmental costs of developing, training, and deploying AI have substantial and long-term impacts on ecosystems, biodiversity, and human health, many of which are difficult to quantify and will probably have an outsized impact on global majority countries.

The environmental impact of AI worsens climate-related health effects

AI development has immediate consequences for increasing health burdens associated with a warming planet. Climate change contributes to hundreds of thousands of additional deaths relating to malnutrition, malaria, heat stress, and diarrhoea each year.[8] Between 2030 and 2050, climate change is predicted to result in an additional 250 000 deaths annually, with the costs of direct damage to health estimated to be $2–4 billion per year by 2030.[8] Those most affected will be already vulnerable populations living in locations in the global majority with health systems that are least equipped to address this additional health burden. We argue that the growing effects of climate change, compounded by collective inaction on climate goals and continued investments in climate-burdensome technologies, are a form of structural violence.[30]

One area of particular concern with regard to climate-related effects on health is air pollution. Increasing levels of air pollution result from the power plants and diesel generators providing electricity to ever-growing numbers of computer-processing centres.[31,32] Air pollution is associated with major public health concerns (eg, including links to cancers and asthma) and the pollution associated with AI processing centres could result in an additional 1300 premature deaths annually in the USA by 2030.[31,32] Air pollution is the second leading risk factor for deaths globally,[33] and WHO estimates that 7 million premature deaths worldwide are related to air pollution.[34]

The growing demand for resources for training and deploying AI systems is also contributing to sociopolitical instability[35] and exacerbating environmental injustices. The environmental health impacts of the AI supply chain are borne almost exclusively by communities in the global majority and echo long-standing histories of colonial extractivism.[29,36] For example, in addition to environmental effects, the increasing demand for metals and other materials needed to build AI systems has devastating social consequences for local communities, contributing to violent conflict in countries such as the Democratic Republic of the Congo.[29,37] In other cases, the expansion of the AI sector to locations in Latin America where business operations are inexpensive has sparked protests—eg, against Google's proposal to build a data centre in Uruguay[38]—related to the ongoing struggle with droughts. In Mexico, data centres have exacerbated existing competition for water use (panel). In other cases, oil and gas companies, such as Shell, Chevron, BP, and ExxonMobil, have partnered with major cloud-computing providers, including Microsoft, Google, and Amazon, to facilitate the discovery, extraction, refining, and marketing of oil and gas.[39–41]

The worsening of environmental health effects from AI development is, thus, a matter of mounting global injustice. The benefits of AI for health are largely being captured by a minority of wealthy nations equipped to integrate AI systems into health-care systems, whereas the environmental costs of AI are being disproportionately borne by countries in the global majority.[17,18] These countries are also less likely to benefit from the potential gains brought by AI, because of persistent material inequality in health-care systems.[42]

The environmental impact of AI is a central ethical concern for global health and bioethics

In the past decade, ethical frameworks for AI have proliferated in response to growing concerns about the adverse effects of AI technologies.[10] For example, one scoping review identified 84 documents containing ethical principles or guidelines for AI that were issued by a wide range of public and private organisations.[2] Such principle-based approaches to making AI more fair, ethical, or responsible have also been followed by private companies, research institutions, and public sector organisations (eg, the European Commission's High-Level Expert Group on AI, the Organisation for Economic Co-operation and Development, Google, IBM, Microsoft, and the AI Now Institute), all of which have issued principle-based guidelines for ethical AI.

Most existing AI ethics frameworks share the same set of core concerns, including transparency, justice and fairness, non-maleficence, responsibility, and privacy. However, the environmental impact of AI, when considered, often falls under the ethical concern of sustainability, which is relatively under-represented compared with other concerns.[2] We call for the environmental impacts of AI to be understood as a central ethical concern for global health and bioethics. This change requires including environmental consequences as a central ethical concern for AI development and using the inclusion of environmental concerns to prompt a reconfiguration of the existing core ethical principles that are used to guide AI development. Such a reconfiguration is necessary because accelerating natural resource consumption driven by AI poses an ethical conundrum that threatens to undermine the core ethical commitments that form the basis of AI frameworks. Here, we outline how an interrogation of AI's environmental impact would reshape established ethical commitments in AI ethics frameworks.

Transparency

Efforts to improve the technical transparency of AI need to be expanded to include disclosure and accountability for resource use and climate burden for all AI products. Ideally, the entire spectrum of resource use should be included, from the sourcing of metals for AI infrastructure or devices that use AI to the increased water and energy used in training, deploying, and monitoring AI technologies.[43] Transparency should be about both reporting outcomes (ie, source materials, energy used, and CO_2 and e-waste produced) and pushing for changes in how AI is developed. For example, at the moment, it is difficult for many data scientists and engineers to calculate the CO_2 emissions of their training jobs. The responsibility of transparency needs to be borne by the data centre owners, who should share records about their energy sources, energy consumption, and emissions of planned tasks to empower developers to make more intentional choices about model development, training, and resource reporting.

Beneficence

Ethical interrogations of how AI can promote better outcomes in health require the frame of consideration to be broadened beyond the individual or clinical cohorts affected by AI, with the global effect of these technologies on the environment and human health being addressed. Potential AI-mediated improvements in the delivery of care for a particular condition—eg, malaria or dengue—need to be weighed against the broader effect of those technologies on a warming climate and associated environmental health effects, including the increasing prevalence of vector-borne diseases. Formalised mechanisms for assessing these trade-offs need to be built into the design and development praxis of the AI field.

Non-maleficence

Preventing or avoiding harm in global health AI applications cannot be limited to averting instances of bias, discrimination, or error. Instead, the explicit environmental costs of AI need to be reflected. Against such assessments of their potential for harm, AI systems should be deployed only in cases where their specific, demonstrated benefit outweighs their harms, including climate burden. Accompanying the publication of model architectures or benchmarking studies with estimations for resource requirements, energy usage, expected emissions, and e-waste during training and in practice[44–46] would enable others using the architectures to create models for specific tasks to make decisions informed by environmental costs.

Equity and fairness

Ethical commitments to equity need to move beyond diversity and representation to interrogate relative questions of social and environmental burden versus health benefit. Such questions concern the current alignment of epistemic expertise and technical know-how for creating and deploying models, currently concentrated in a small number of wealthy countries that are prepared to use AI, for example, in health care. This concentration of expertise excludes the global majority, resulting in AI systems designed to serve the needs of relatively few individuals.[42] The extractivist premises underlying much of AI development—from data to metals—must also be scrutinised, and more just forms of technological co-creation that prioritise environmental protection from the start are needed.[18,47]

Justice

Accounting for the full extent of the socioenvironmental burden across the AI pipeline should raise difficult questions of prioritisation in the development and deployment of AI models. Investments in AI solutions in global health need to be weighed for their relative opportunity costs against established means of advancing health equity, namely investing in robust first-line health-care systems in places that continue to struggle with material limitations to providing care, sourcing medicines, and paying staff a fair wage.[48]

Strategies to foster environmental accountability in AI

Inclusion of environmental impact in the AI ethics agenda is vital because the current trajectory of resource use is unsustainable. As such, it is necessary to develop concrete measures for addressing the climate impact of AI to ensure that the discourse around this topic does not remain limited to ethical reflection but is enacted in technological practices and changed AI workflows.

Fortunately, some efforts are underway within the computer science and technology sectors to address the environmental impact of AI. These include initiatives on energy-efficient AI models, renewable energy for data centres, and carbon tracking and offsetting. For example, tools such as Carbontracker[23] enable AI engineers to track the energy uses of their AI training jobs. Furthermore, efforts to foster so-called green AI seek to improve innovation in AI by encouraging computational efficiency as a criterion for evaluation in addition to established metrics such as accuracy.[49] The inclusion of energy efficiency could facilitate the systematic evaluation of environmental considerations in AI model training. Relatedly, building on Mitchell and colleagues' proposal for model cards,[50] Hershcovich and colleagues have proposed climate model cards—which standardise reporting on the environmental impact of a trained AI model—in the field of natural language processing, one of the most energy-intensive AI fields.[51] Although such efforts are promising, their uptake has been minor within the AI field. Just as there are computational trade-offs in AI research between goals such as performance and efficiency, environmental burdens also need to be justified in order to advance ethical AI.

Work to improve environmental accountability will require an integrative approach that includes strategies to change research culture, bring about new forms of accountability and regulation, and foster international cooperation and incentivisation. Here, we outline several concrete strategies for accountability in the area of global health, to suggest pathways for better aligning AI development and AI ethics with climate goals.[40,52–54]

Reporting

Similar to other sorts of authorship (eg, consolidated criteria for reporting qualitative research) and dataset reporting requirements, journals should require reporting of the carbon footprint of AI use.[52] This reporting would include, for example, the computing resources that were used in the training or deployment of research articles using or referring to AI or machine learning models, as well as the equivalent computed carbon emissions[46] and information about where and how materials for the creation of AI systems would be sourced. As noted previously, carbon footprint recording requires increased transparency on the part of those who own data centres and computational power to make environmental impact information available. Existing examples of systematised reporting mechanisms,[55] including software packages that integrate into coding platforms,[44–46] can be used to publish detailed, accessible information along with the articles. These reporting mechanisms should be considered a routine component of disclosure and academic integrity.

Funding

Proposals for new grants relating to training or deploying AI should be required to include detailed plans for estimating the associated greenhouse gas emissions, information regarding where and how metals and other essential materials will be sourced, and other measures for evaluating the anticipated climate impact of the proposed AI technology. Concurrently, applicants should be required to discuss mitigation measures that they will take to curtail this impact. Similar to the scenario with progress reporting, funders would retain the option of ending funding if a project was found to be incompliant.

Auditing

Algorithmic audits have been proposed as a means of investigating potential bias, discrimination, or other concerns within algorithmic systems and their content. Such audits include repeatedly and systematically querying an algorithm and recording the responding outputs to assess how the algorithm functions in relation to a specific concern—eg, the social determinants of health or health outcomes.[56] We suggest that these auditing techniques could be productively combined with existing models for assessing environmental effects. For example, environmental impact assessments have been applied in the oil and gas industries and are a widely used instrument of regulatory control in the environmental sector. To assess the social and environmental effects of large-scale AI models, the model of environmental impact assessment could be combined with algorithmic auditing to encourage compliance with established, yet non-binding, environmental criteria for sustainable AI development.

Regulation

It is abundantly clear that major multinational technology companies (ie, so-called Big Tech) cannot be left to manage their own environmental or ethical compliance.[57,58] The 2024 EU Artificial Intelligence Act[59] notes the importance of AI contributing to societal and environmental protection and environmentally beneficial outcomes. It recalls earlier guidelines developed by the European Commission's High-Level Expert Group on AI,[60] with Recital 27[59] emphasising the importance of "societal and environmental well-being" such that AI systems "are developed and used in a sustainable and environmentally friendly manner as well as in a way to benefit all human beings, while monitoring and assessing the long term impacts on the individual, society and democracy".[59] However, aside from these broad commitments, the environmental considerations of AI are largely left as a formality that do not translate into specific recommendations and controls and are excluded from the regulatory classifications of risk in the EU Act. We believe that this legal framework was a missed opportunity and that regulatory measures need to be revised to enforce specific social and environmental industry standards for the fair and conscientious sourcing and use of AI materials as well as investments in carbon offsets and renewable energy resources. Environmental consequences should factor into an evaluation of relative risk assessments, particularly in the area of health. This approach could offer a more forceful strategy for addressing the otherwise vague claims of benefit or environmentally friendly AI development proffered by companies leading the AI global health field.

Cooperation

Intranational agreements on energy sources and use for AI are enabling some progress. For example, the Climate Neutral Data Centre Pact[61] is a consortium of more than 100 data centre operators and trade associations that have committed to achieving the greenhouse gas reductions outlined in the European Green Deal, including an agreement to make data centres climate neutral by 2030. However, buy-in and alignment from countries such as China and the USA will be imperative for such consortia to have an appreciable impact. Resource use, energy footprints, and e-waste should also be considered in such agreements.

Incentives

Choices to use environmentally friendly computational power for researching and developing AI models should be recognised and rewarded. In academia, early-career scholars often conduct the programming work and have few incentives for choosing not to do it because of environmental considerations. The same is true at the scale of technology corporations: until the environmental costs of AI are made an explicit and routine part of the social and economic calculus of AI development, companies choosing to invest in less carbon-intensive energy sources or more efficient energy grids will not be rewarded.

Intentionality

Many choices in the way that AI is created affect its environmental impact, from top-level investment decisions to choices made by individual research teams. These include, for example, decisions regarding the type of cloud provider (such as one with documented offsetting of emissions through renewable energy certificates), selection of carbon-friendly regions for data centre locations (the local grid affects the carbon intensity of energy usage by the data centre—eg, some grids are low carbon and are powered by renewable energy sources), prioritisation of algorithm efficiency, choice of computing hardware, and efforts to minimise training runs.[54,55] A culture of intentional AI needs to be incentivised and nurtured so that deliberate decisions are made to prioritise the environment at all levels of AI training and deployment.

Conclusion

The suggestions in this Viewpoint have specific consequences for those researching, implementing, and investing in AI technologies in global health, as well as direct implications for how we conceive of and address ethical issues in AI. More broadly, efforts to holistically understand, mitigate, and reduce the environmental effects of AI across multiple technological, clinical, and applied fields of research and practice should be prioritised. These efforts are as much about the sources and sourcing of materials as they are about changing workplace cultures and norms within Big Tech, data centres, research groups, and clinical practice.[62] The specific costs of environmental resources—relating to water, electricity, and carbon—that result from the training, deployment, and infrastructure required to support AI are rarely made explicit in ethical debates over AI technologies.[52,53] Regarding global health and bioethics specifically, we contend that the environmental effects of AI are an ethical imperative in and of themselves. The promises of sustainability, equity, or good brought by AI technologies will be elusive if we fail to fully reckon with the profound costs of these technologies for the planet.

Contributors

AF initiated the conception and design of the manuscript, wrote the first draft, and edited subsequent drafts of the manuscript for intellectual content. IMR, TW, AB, LAC, and SM drafted sections of the manuscript, edited drafts, and contributed to subsequent revision of intellectual content. All authors had full access to all the data in writing the manuscript, approved the final version, and accept responsibility for the submitted publication.

Declaration of interests

LAC declares consulting fees from Philips for evaluation of a haemodynamic monitoring product; payment or honoraria from Stanford University (visiting professor), University of California San Francisco (visiting professor), University of Toronto (visiting professor), and Taipei Medical University (remote speaker); support for attending meetings or travel from the College of Intensive Care Medicine of Australia and New Zealand, University of Bergen, University Medical Center Amsterdam, Académie Nationale de Médecine (France), and the Doris Duke Foundation; leadership or fiduciary roles with *PLOS Digital Health* (Editor-in-Chief) and *The Lancet Digital Health* (International Advisory Board member); and receives compute credits for Oracle. All other authors declare no competing interests.

Acknowledgments

LAC is funded by the US National Institutes of Health through R01 EB017205, DS-I Africa U54 TW012043–01, and Bridge2AI OT2OD032701 and by the US National Science Foundation through ITEST 2148451.

References

1 Hao K. AI is taking water from the desert. March 1, 2024. https://www.theatlantic.com/technology/archive/2024/03/ai-water-climate-microsoft/677602/ (accessed Sept 6, 2024).

2 Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. *Nat Mach Intell* 2019; **1:** 389–99.

3 Gates B. The age of AI has begun. March 21, 2023. https://www.gatesnotes.com/The-Age-of-AI-Has-Begun (accessed Nov 14, 2024).

4 Mrazek MF, O'Neill FM. Artificial intelligence and healthcare in emerging markets. Sept 1, 2020. https://documents.worldbank.org/pt/publication/documents-reports/documentdetail/733971606368563566/Artificial-Intelligence-and-Healthcare-in-Emerging-Markets (accessed Nov 12, 2024).

5 Schwalbe N, Wahl B. Artificial intelligence and the future of global health. *Lancet* 2020; **395:** 1579–86.

6 WHO. Regulatory considerations on artificial intelligence for health. 2023. https://iris.who.int/handle/10665/373421 (accessed Oct 10, 2024).

7 Nuffield Council on Bioethics. Bioethics briefing note: artificial intelligence (AI) in healthcare and research. Nuffield Council on Bioethics, 2018.

8 WHO. Ethics and governance of artificial intelligence for health: WHO guidance. World Health Organization, 2021.

9 UNESCO. Recommendation on the ethics of artificial intelligence. 2021. https://unesdoc.unesco.org/ark:/48223/pf0000380455/PDF/380455eng.pdf.multi (accessed Nov 15, 2024).

10 Morley J, Machado CCV, Burr C, et al. The ethics of AI in health care: a mapping review. *Soc Sci Med* 2021; **260:** 113172.

11 Blanco M. Por consumo de agua, activistas rechazan la llegada de Data Centers. El Universal Querétaro, July 3, 2024. https://www.eluniversalqueretaro.mx/cartera/por-consumo-de-agua-activistas-rechazan-la-llegada-de-data-centers/ (accessed March 10, 2025).

12 Conde C. Querétaro sufre la peor sequía del milenio. Diario de Querétaro, March 11, 2024. https://oem.com.mx/diariodequeretaro/local/queretaro-sufre-la-peor-sequia-del-milenio-13100936 (accessed March 10, 2025).

13 Graham T. Mexico's datacentre industry is booming—but are more drought and blackouts the price communities must pay? The Guardian, Sept 25, 2024. https://www.theguardian.com/global-development/2024/sep/25/mexico-datacentre-amazon-google-queretaro-water-electricity (accessed March 10, 2025).

14 UN Environmental Progamme. AI has an environmental problem. Here's what the world can do about that. Sept 21, 2024. https://www.unep.org/news-and-stories/story/ai-has-environmental-problem-heres-what-world-can-do-about (accessed Nov 15, 2024).

15 Keller R. Geographies or power, legacies of mistrust: colonial medicine in the global present. *Hist Geogr* 2006; **34:** 26–48.

16 Affun-Adegbulu C, Adegbulu O. Decolonising global (public) health: from western universalism to global pluriversalities. *BMJ Glob Health* 2020; **5:** e002947.

17 Adan SN, Trager R, Blomquist K, et al. Voice and access in AI: global AI majority participation in artificial intelligence development and governance. Centre for the Governance of AI, 2024.

18 Muldoon J, Wu BA. Artificial intelligence in the colonial matrix of power. *Philos Technol* 2023; **36:** 80.

19 Bender EM, Gebru T, McMillan-Major A, Shmitchell S. On the dangers of stochastic parrots: can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '21; March 3–10, 2021.

20 Li P, Yang J, Islam MA, Ren S. Making AI less "thirsty": uncovering and addressing the secret water footprint of AI models. *arXiv* 2023; published online April 6. http://arxiv.org/abs/2304.03271 (preprint).

21 Moss S. Microsoft's water consumption jumps 34 percent amid AI boom. Sept 12, 2023. https://www.datacenterdynamics.com/en/news/microsofts-water-consumption-jumps-34-percent-amid-ai-boom/ (accessed Nov 2, 2024).

22 Valdivia A. Silicon valley and the environmental costs of AI. Dec 5, 2022. https://www.perc.org.uk/project_posts/silicon-valley-and-the-environmental-costs-of-ai/ (accessed Nov 2, 2024).

23 Anthony LFW, Kanding B, Selvan R. Carbontracker: tracking and predicting the carbon footprint of training deep learning models. *arXiv* 2020; published online July 6. http://arxiv.org/abs/2007.03051 (preprint).

24 de Vries A. The growing energy footprint of artificial intelligence. *Joule* 2023; **7:** 2191–94.

25 Erdenesanaa D. AI could soon need as much electricity as an entire country. The New York Times, Oct 10, 2023. https://www.nytimes.com/2023/10/10/climate/ai-could-soon-need-as-much-electricity-as-an-entire-country.html (accessed Nov 1, 2024).

26 Hao K. Training a single AI model can emit as much carbon as five cars in their lifetimes. June 6, 2019. https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/ (accessed Nov 2, 2024).

27 Robinson D. AI's thirst for power keeps coal fires burning bright. Oct 14, 2024. https://www.theregister.com/2024/10/14/ai_datacenters_coal/ (accessed Nov 1, 2024).

28 Halper E. A utility promised to stop burning coal. Then Google and Meta came to town. Washington Post, Oct 12, 2024. https://www.washingtonpost.com/business/2024/10/08/google-meta-omaha-data-centers/ (accessed Nov 1, 2024).

29 Dauvergne P. Is artificial intelligence greening global supply chains? Exposing the political economy of environmental costs. *Rev Int Polit Econ* 2022; **29:** 696–718.

30 Farmer P. Pathologies of power: health, human rights, and the new war on the poor, 1st edn. University of California Press, 2004.

31 Danelski D. AI's deadly air pollution toll. Dec 9, 2024. https://news.ucr.edu/articles/2024/12/09/ais-deadly-air-pollution-toll (accessed March 10, 2025).

32 Han Y, Wu Z, Li P, Wierman A, Ren S. The unpaid toll: quantifying the public health impact of AI. *arXiv* 2024; published online Dec 9. http://arxiv.org/abs/2412.06288 (preprint).

33 Health Effects Institute. State of global air report 2024: a special report on global exposure to air pollution and its health impacts, with a focus on children's health. Health Effects Institute, 2024.

34 WHO. Air pollution. https://www.who.int/health-topics/air-pollution (accessed March 10, 2025).

35 European Commission. Causes of climate change. https://climate.ec.europa.eu/climate-change/causes-climate-change_en (accessed Nov 14, 2024).

36 Crawford K. Atlas of AI: the real worlds of artificial intelligence. Yale University Press, 2021.

37 Diemel JA, Hilhorst DJM. Unintended consequences or ambivalent policy objectives? Conflict minerals and mining reform in the Democratic Republic of Congo. *Dev Policy Rev* 2019; **37:** 453–69.

38 McGovern G, Branford S. The cloud vs. drought: water hog data centers threaten Latin America, critics say. Nov 2, 2023. https://news.mongabay.com/2023/11/the-cloud-vs-drought-water-hog-data-centers-threaten-latin-america-critics-say/ (accessed Nov 14, 2024).

39 Greenpeace. Greenpeace report: oil in the cloud. May 19, 2020. https://www.greenpeace.org/usa/oil-in-the-cloud/ (accessed March 10, 2025).

40 Kaack L, Donti P, Strubell E, Rolnick D. Artificial intelligence and climate change. December, 2020. https://eu.boell.org/en/2020/12/03/artificial-intelligence-and-climate-change (accessed Nov 18, 2024).

41 Victor DG. How artificial intelligence will affect the future of energy and climate. Jan 10, 2019. https://www.brookings.edu/articles/how-artificial-intelligence-will-affect-the-future-of-energy-and-climate/ (accessed Nov 18, 2024).

42 McCoy L, Bihorac A, Celi LA, et al. Building health systems capable of leveraging AI: applying Paul Farmer's 5S framework for equitable global health. *BMC GlobPublic Health* (in press).

43 Bolón-Canedo V, Morán-Fernández L, Cancela B, Alonso-Betanzos A. A review of green artificial intelligence: towards a more sustainable future. *Neurocomputing* 2024; **599:** 128096.

44 Code Carbon. Track and reduce CO2 emissions from your computing. https://codecarbon.io/ (accessed Nov 25, 2024).

45 Green Algorithms. Towards environmentally sustainable computational sciences. http://www.green-algorithms.org/ (accessed Nov 25, 2024).

46 Machine Learning CO2 Impact. Machine learning has a carbon footprint. https://mlco2.github.io/impact (accessed Nov 18, 2024).

47 Couldry N, Mejias U. Data grab: the new colonialism of Big Tech and how to fight back. University Of Chicago Press, 2024.

48 Farmer PE, Nizeye B, Stulac S, Keshavjee S. Structural violence and clinical medicine. *PLoS Med* 2006; **3:** e449.

49 Schwartz R, Dodge J, Smith NA, Etzioni O. Green AI. *Commun ACM* 2020; **63:** 54–63.

50 Mitchell M, Wu S, Zaldivar A, et al. Model cards for model reporting. *Proceedings of the 2019 ACM Conference on Fairness, Accountability, and Transparency*. FAccT '19; Jan 29–31, 2019.

51 Hershcovich D, Webersinke N, Kraus M, Bingler J, Leippold M. Towards climate awareness in NLP Research. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing* 2022: 2480–94.

52 Cowls J, Tsamados A, Taddeo M, Floridi L. The AI gambit: leveraging artificial intelligence to combat climate change-opportunities, challenges, and recommendations. *AI Soc* 2023; **38:** 283–307.

53 Kaack LH, Donti PL, Strubell E, Kamiya G, Creutzig F, Rolnick D. Aligning artificial intelligence with climate change mitigation. *Nat Clim Chang* 2022; **12:** 518–27.

54 Lacoste A, Luccioni A, Schmidt V, Dandres T. Quantifying the carbon emissions of machine learning. *arXiv* 2019; published online Nov 4. http://arxiv.org/abs/1910.09700 (preprint).

55 Henderson P, Hu J, Romoff J, Brunskill E, Jurafsky D, Pineau J. Towards the systematic reporting of the energy and carbon footprints of machine learning. *J Mach Learn Res* 2020; **21:** 1–43.

56 Metaxa D, Park JS, Robertson RE, et al. Auditing algorithms: understanding algorithmic systems from the outside in. *FNT in Human–Computer Interaction* 2021; **14:** 272–344.

57 Hogan M. Data flows and water woes: the Utah data center. *Big Data Soc* 2015; **2:** 1–12.

58 Ochigame R. The invention of "ethical AI": how Big Tech manipulates academia to avoid regulation. Dec 20, 2019. https://theintercept.com/2019/12/20/mit-ethical-ai-artificial-intelligence/ (accessed Dec 7, 2021).

59 Future of Life Institute. EU Artificial Intelligence Act. 2024. https://artificialintelligenceact.eu/ (accessed Nov 18, 2024).

60 European Commission. Ethics guidelines for trustworthy AI. April 8, 2019. https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai (accessed Nov 18, 2024).

61 Climate Neutral Data Centre. Climate Neutral Data Centre Pact: the Green Deal needs green infrastructure. https://www.climateneutraldatacentre.net/ (accessed Nov 20, 2024).

62 Monserrate SG. The cloud is material: on the environmental impacts of computation and data storage. Jan 27, 2022. https://mit-serc.pubpub.org/pub/the-cloud-is-material/release/2 (accessed Nov 14, 2024).