Environmental Science and Ecotechnology 23 (2025) 100520

Contents lists available at ScienceDirect

Environmental Science and Ecotechnology

journal homepage: www.journals.elsevier.com/environmental-science-and-ecotechnology/

Commentary

A social-environmental impact perspective of generative artificial intelligence

1. Introduction

Generative artificial intelligence (GenAI) tools are increasingly developed and deployed with limited consideration for their social and environmental impacts (SEIs). This is partly because not all impacts are ostensible and quantifiable, and it is unclear who should steward the calculation and mitigation of SEIs, as well as which mechanisms or pressure points should be used to ensure they are addressed. Here, we show a preliminary list of focus areas to discuss SEIs of GenAI and stress challenges involved in identifying and quantifying these. Of note, we exclude the consequences of using GenAI (e.g., potential biases, privacy concerns, dual-use, and accountability) because these are context-specific and vary significantly depending on the utilized GenAI, the user, and the use cases. Development and deployment phases are selected for their substantial impacts on resources and social dynamics. We explore SEIs in relation to the required (1) hardware and (2) training and development (Fig. 1).

2. GenAI hardware

2.1. Environmental impacts

The production of graphics processing units (GPUs) required for running GenAI systems has various negative SEIs for the host communities. For example, Taiwan Semiconductor Manufacturing Company (TSMC)'s planned manufacturing site in Taichung (2.85 million residents) will consume 25% of the city's electricity and 6% of its water, raising concerns among residents [1]. Additionally, manufacturing GPUs requires rare metals like tantalum and cobalt. Extracting these minerals has significant environmental impacts, including deforestation and soil and water pollution. Beyond manufacturing GPUs, building and operating data centers essential for GenAI processing represents a growing environmental challenge. Estimates show that data centers and the required energy for cooling them will consume more than 8% of United States electricity and 5% of Europe's by 2030 [2]. Furthermore, other required hardware components such as storage drives and central processing units (CPUs) need plastics for casings and insulation, while microchips and semiconductors rely on silicon and metals like gold, copper, and aluminum. Mining these resources has had disastrous environmental consequences for local communities, including water contamination, air pollution, and soil degradation—all of which exemplify a broader trend of increased environmental footprint, requiring substantial resources from around the world. Another environmental impact of GenAI pertains to e-waste generation and the fact that some of the materials used in GPUs or data centers are non-recyclable, exacerbating environmental pollution and waste management challenges.

2.2. Social impacts

The social implications of GenAI hardware follow systematic disturbance patterns across multiple locations. For example, in Chandler, Arizona (population 280,000, where electricity and space are affordable), city officials have limited the construction of new data centers because they do not create quality jobs for local communities and produce disturbing levels of noise [3]. In Memphis, Tennessee (population 618,000), residents and the city council were taken by surprise when Elon Musk announced that xAI's new center (dedicated to GenAI operations) will be built in a former Electrolux facility. In addition to concerns about xAI's illegal operation of gas turbines, community activists worry about “air quality, water access, and grid stability, especially for nearby neighborhoods that have suffered from industrial pollution for decades.” xAI also undermined local communities by signing a non-disclosure agreement with the Greater Memphis Chamber and Memphis, Gas, Light, and Water Division, thereby effectively concealing details of their negotiations and agreements [4]. Furthermore, the mining of minerals like lithium and cobalt has been reported to include child labor and negatively affect mining communities. Cobalt mining communities like Kolwezi (population 573,000) in the DRC are reported to put children involved in mining operations at major risk.

3. GenAI training and development

3.1. Environmental impacts

GenAI models require vast amounts of data for training and operation, necessitating significant processing capacity and energy. Training a model like GPT3 (with 175 billion parameters) typically consumes about 1287 MWh of electricity and emits approximately 552 tons of CO_2 [5]. Data centers that process and store data consume large amounts of energy to maintain and cool servers. For context, the xAI training facility in Memphis uses as much power as 80,000 households. In July 2022 alone, an OpenAI facility training GPT4 used about 6% of the water used by West Des Moines, Iowa (population 75,000) [1]. These examples show that training GenAI has distinct impacts beyond manufacturing and procuring the required hardware. Nevertheless, emissions also spill over to sustaining the required digital infrastructure for end-to-end

https://doi.org/10.1016/j.ese.2024.100520

	Hardware	Training & Development
Social impact	• Social gentrification • Adverse health impacts • Increased noise pollution • Poor worker safety conditions • Unforeseen psychological impacts	• Increased digital divide • Few languages used in training • Unsuitable labor laws and compensation • Limited perspectives and worldviews
Environmental impact	• Increased e-waste • Extracting rare earth metals • Whole-life carbon footprint • Energy and water consumption • Need for new electrical and data grids • Data centers and GPUs heat generation	• Need for continuous updates • Enormous data storage capacity needed • Increased demand for network traffic • Resource-intensive deployment and scaling

Fig. 1. Social and environmental impacts of developing generative artificial intelligence. These include environmental challenges such as energy and resource consumption and social issues like labor conditions and accessibility gaps, highlighting the systemic resource demands and socio-economic ramifications of generative artificial intelligence.

machine learning pipeline and content generation. For example, in addition to its existing processing capacity, Meta plans to buy 350,000 GPUs ($10.5 billion) to train its large language model. This is reflective of the estimated 12% increase in the annual global GPU demand, increasing the projected energy consumption and strain on data centers, as well as the energy and network grids [6]. Frequent updates and the need for continuous operation further exacerbate environmental impacts. As network traffic increases, so does the energy required to maintain the infrastructure that supports these data transfers.

3.2. Social impacts

Training GenAI requires numerous programs, algorithms, and data structures for preprocessing, data/image/sound categorization and labeling, model architecture design, hyperparameter tuning, and iterative learning cycles. Women and older adults, as well as speakers of most non-English languages, are sometimes underrepresented in these tasks, leading to biased systems that overrepresent the perspectives of the Anglophone world and are not user-friendly for all cohorts [7]. Another concern is the unequal availability and accessibility of GenAI across high- and low-income countries as well as various social groups in the same community, for example, based on education level, race, or digital literacy. GenAI training tasks are also fueled by millions of underpaid workers who perform repetitive tasks under precarious labor conditions.

4. Conclusions, recommendations, and future directions

4.1. Promoting sustainable practices to mitigate environmental impacts

Although vendors promote the ease and speed of creating context-specific GenAI, users and developers need to critically evaluate whether GenAI is truly the most optimal solution for their needs. The advantages of GenAI in industrialized nations often conceal the externalization of SEIs to marginalized communities or those without democratic decision-making processes. One way of mitigating environmental impacts (when GenAI is deemed the most optimal/viable solution) is to adopt sustainable GenAI development measures such as:

(1) Parameter-efficient fine-tuning to help optimize model performance while using fewer resources.
(2) Knowledge distillation enhances the performance of smaller models by transferring knowledge from larger, more complex models.
(3) Sustainable hardware-aware designs.
(4) Edge computing deployment to reduce latency and enhance response times.
(5) Utilize specialized AI accelerators for tasks that involve intensive computations.
(6) Reuse existing models and share them following FAIR principles—*Findable, Accessible, Interoperable, and Reusable*. This means making GenAI models easily discoverable through unique identifiers and standardized metadata, ensuring they are openly accessible without significant barriers, utilizing consistent formats for compatibility across various applications and platforms, and providing clear documentation and licensing to facilitate reuse.

4.2. A call for inclusive efforts and clear accountabilities to mitigate social impacts

The path toward mitigating social impacts requires collaborative and inclusive efforts from various groups and identifying priority areas. Some examples may include:

(1) Standardization of impact metrics.
(2) Investigation of impacts for affected communities.
(3) Developing sustainable facilities that are designed together with community representatives.
(4) Supporting community-driven initiatives and policies.
(5) Improving labor conditions, including job security, fair wages, and workplace safety.
(6) Ethical governance for GenAI.

4.3. Future directions

Here, we used a range of examples to articulate some SEIs for developing and deploying GenAI. We did not attempt to quantify the impacts, recognizing that assessing the cumulative SEIs of GenAI is highly complicated. Factors contributing to this complexity include various energy-intensive processes with a complicated value chain and unforeseen socio-economic ramifications for marginalized communities. Let's also not forget that it is nearly impossible to assign a definitive value to natural resources, such as a forest, prairie, or lake, or the cultural and social well-being of affected communities. Therefore, we recommend developing a comprehensive model that evaluates the lifecycle impacts of GenAIs at every stage—from design and development to deployment, maintenance, decommissioning, recycling, and disposal. Existing theoretical frameworks for evaluating technology impacts, such as the Social Life Cycle Assessment and Environmental Impact Assessment, could provide structured approaches for understanding these impacts [8]. That said, any impact assessment model should be developed through a transdisciplinary and inclusive approach that actively engages researchers in metrics development, industry partners in data collection and implementing sustainable practices, policymakers in regulations, and local community members in providing feedback. A transdisciplinary and inclusive approach can help capture a wider range of perspectives and ensure that the resulting assessment model addresses relevant SEIs. Furthermore, while currently, some commercial and public parties voluntarily declare hardware usage and other human/financial resources used for deploying and developing GenAI, this is not a mandatory requirement. We recommend those developing and deploying GenAI disclose these details in their annual reports. This requires pressure from policymakers and regulatory bodies overseeing and auditing technology use. Future research could look into these avenues and offer incentives for disclosure.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work originated through the 2023–2024 Biomedical Data Science Innovation Lab (BDSIL; NIGMS grant No. R25GM139080) on “Building Partnerships for Generative AI Training in Biomedical and Clinical Research.” Mohammad Hosseini was supported by the NIH National Center for Advancing Translational Sciences (NCATS, UM1TR005121). The funders have not played a role in the design, analysis, decision to publish, or manuscript preparation. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NCATS, NIH, or the US government.

References

[1] M. Tobin, V. Chiang, P.-L. Wu, Economy vs. environment: some Taiwanese consider cashing in their chips, Wash. Post (2023).
[2] Goldman Sachs. AI poised to drive 160% increase in data center power demand. Goldman Sachs, https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand.
[3] Terri, The staggering ecological impacts of computation and the cloud, MIT Schwarzman College of Computing (2022). https://computing.mit.edu/news/the-staggering-ecological-impacts-of-computation-and-the-cloud/.
[4] Time. Elon Musk's xAI launches Grok, a chatbot with a personality. Time, https://time.com/7021709/elon-musk-xai-grok-memphis/.
[5] Y.I. Alzoubi, A. Mishra, Green artificial intelligence initiatives: potentials and challenges, J. Clean. Prod. 468 (2024) 143090, https://doi.org/10.1016/j.jclepro.2024.143090.
[6] M. Kan, Zuckerberg's Meta Is Spending Billions to Buy 350,000 Nvidia H100 GPUs, PCMag, 2024. https://www.pcmag.com/news/zuckerbergs-meta-is-spending-billions-to-buy-350000-nvidia-h100-gpus.
[7] Z. Chen, Ethics and discrimination in artificial intelligence-enabled recruitment practices, Humanit. Soc. Sci. Commun. 10 (2023) 567, https://doi.org/10.1057/s41599-023-02079-x.
[8] S. Yang, K. Ma, Z. Liu, J. Ren, Y. Man, Development and applicability of life cycle impact assessment methodologies, in: J. Ren, S. Toniolo (Eds.), Life Cycle Sustainability Assessment for Decision-Making, Elsevier, 2020, pp. 95–124, https://doi.org/10.1016/B978-0-12-818355-7.00005-1.

Mohammad Hosseini**
Department of Preventive Medicine, Northwestern University Feinberg School of Medicine, Chicago, IL, USA

Galter Health Sciences Library and Learning Center, Northwestern University Feinberg School of Medicine, Chicago, IL, USA

Peng Gao*
Department of Environmental and Occupational Health, and Department of Civil and Environmental Engineering, University of Pittsburgh, Pittsburgh, PA, USA

Department of Environmental Health, Harvard T.H. Chan School of Public Health, Boston, MA, USA

Carolina Vivas-Valencia
Department of Biomedical Engineering and Chemical Engineering, The University of Texas at San Antonio, San Antonio, TX, USA

** Corresponding author. Department of Preventive Medicine, Northwestern University Feinberg School of Medicine, Chicago, IL, USA.

* Corresponding author. Department of Environmental Health, Harvard T.H. Chan School of Public Health, Boston, MA, USA.
E-mail address: mohammad.hosseini@northwestern.edu (M. Hosseini).
E-mail address: pgao@hsph.harvard.edu (P. Gao).

24 September 2024