



Asian Journal of Distance Education

G for Generative is a Gluttonous Fraud and Shit, G for Green is Genuine Intelligence: Exploring the Unspoken Environmental Cost and the Computational Opulence Paradox in the Artificial Intelligence Landscape

Aras Bozkurt, Ramesh C Sharma

Abstract: Generative Artificial Intelligence (GenAI) is frequently hailed as a revolutionary force, yet its rapid adoption often overlooks the substantial environmental costs hidden behind the metaphor of the cloud. This paper critically examines the physical infrastructure of GenAI, arguing that the current trajectory represents a "gluttonous fraud" of resource consumption compared to the "genuine intelligence" of sustainable computing. Through a synthesis of recent environmental impact studies, this research deconstructs the four pillars of AI's footprint: energy consumption dominated by inference rather than training, the hidden water usage for cooling and electricity generation, the material toll of e-waste and mineral extraction, and the lack of corporate transparency. The analysis reveals that the efficiency gains of current models paradoxically drive higher consumption and that specific tools actively propagate non-green code. The study concludes that without intervention, GenAI will exacerbate the climate crisis through a "computational opulence paradox." Consequently, it proposes actionable pathways toward Green AI, including the adoption of smaller, specialized models, mandatory environmental impact reporting, and regulatory frameworks. These findings imply a necessary paradigm shift for policymakers and academic leaders from unfettered adoption to a model of "slow urgency" and rigorous accountability.

Keywords: generative AI, green AI, environmental impact, carbon footprint, water footprint, sustainable computing, AI ethics, energy consumption, electronic waste, AI policy, educational technology, higher education

Highlights

What is already known about this topic:

- Training large language models requires massive amounts of energy and generates significant carbon emissions.
- Data centers supporting AI infrastructure consume vast quantities of water for cooling systems, often located in water-stressed regions.
- The concept of Green AI promotes energy-efficient practices, but the industry currently prioritizes performance accuracy over sustainability.

What this paper contributes:

- It reframes the energy debate by demonstrating that continuous model inference, not training, constitutes the dominant long-term environmental cost.
- The study exposes the often-overlooked Scope 2 water footprint derived from offsite electricity generation, which exceeds direct onsite cooling usage.
- It synthesizes diverse environmental impacts to define the "computational opulence paradox," contrasting the gluttony of Red AI with the efficiency of Green AI.

Implications for theory, practice and/or policy:

- Policymakers must mandate standardized environmental disclosures for foundation models, treating them as industrial products requiring regulation.
- Organizations should prioritize Green AI strategies like model pruning, quantization, and knowledge distillation to reduce energy consumption.
- Academic and institutional leaders must demand full Environmental Impact Reports (EIRs) from vendors before procuring or deploying generative AI tools.



Introduction: The Mirage of the Digital Cloud

Is the fleeting amusement of generating a surrealist image of an astronaut riding a horse worth the same energy expenditure as fully charging your smartphone? (Heikkilä, 2023). This is not a rhetorical flourish, but a literal accounting of the hidden thermodynamic cost of our current technological obsession. We stand at the precipice of a new era, often heralded as the end of the "AI winter" and the dawn of a new age of intelligence (Bashir et al., 2024). However, this "digital transformation" is increasingly revealing itself to be a physical crisis, masked by the ethereal metaphors of "clouds" and "virtual assistants." The narrative of Generative AI (Gen-AI) is currently dominated by a "Klondike gold rush" mentality (Bashir et al., 2024), characterized by unfettered growth, corporate consolidation, and a dangerous lack of regulatory oversight.

Behind the sleek interfaces of chatbots and image generators lies a sprawling, resource-intensive infrastructure that is fundamentally altering our planetary metabolism. While proponents argue that AI will optimize energy grids and solve climate change, the immediate reality is a "computational opulence paradox": we are burning vast amounts of fossil fuels and evaporating millions of gallons of water to build tools that may eventually, perhaps, help us save the environment. The industry operates in a "black box" of opacity, where tech giants guard energy data as trade secrets, leaving researchers and policymakers to estimate the damage through "grasping at straws" (Chen, 2025). This article argues that the current trajectory of Generative AI represents a gluttonous fraud, a resource-heavy luxury masquerading as essential progress, while "Green AI" offers the only path toward genuine, sustainable intelligence.

This crisis of consumption strikes at the heart of higher education, a sector that champions sustainability goals while simultaneously rushing to integrate these energy-voracious tools into classrooms and research labs. By adopting Gen-AI without scrutinizing its environmental price tag, academic leaders are not merely purchasing software but are subscribing to a massive, ongoing, and cumulative energy liability (Bozkurt, 2025). This uncritical embrace risks normalizing "computational opulence" within the curriculum itself, effectively teaching the next generation to value technological convenience over ecological survival and failing the ethical mandate to foster a livable future. To understand the depth of this betrayal, we must first dissect the mechanics of the "gluttonous fraud" itself.

The Gluttony of "Red AI" vs. The Intelligence of Green

The title of this discourse is not merely provocative; it is a taxonomy of two divergent paths in computing. The first, "G for Generative," has become synonymous with "Red AI"; a paradigm that prioritizes state-of-the-art accuracy above all else, buying trivial performance gains with massive increases in computational cost (Bolón-Canedo et al., 2024).

The Gluttonous Fraud: Energy, Water, and Waste

The "gluttony" of Generative AI is physical and quantifiable. While the training phase of Large Language Models (LLMs) is notoriously energy-intensive, training a single large NLP model can emit as much carbon as five cars over their entire lifetimes (Hao, 2019; Tabbakh et al., 2024), the true "fraud" lies in the hidden cost of *inference* (actual usage). It is estimated that 90% of the computing power for AI is eventually used for inference, not training (O'Donnell & Crownhart, 2025). For instance, generating just one image with a powerful model consumes roughly as much energy as driving an average gasoline car for 4.1 miles (Heikkilä, 2023). When multiplied by millions of daily users, this becomes an ecological disaster; a simple query to GenAI consumes roughly 25 times more energy than a standard search (Chen, 2025). One year of inference for a GenAI like service can produce 25 times the carbon emissions of the model's initial training (Chien et al., 2023). Furthermore, AI tools are actively propagating non-green code, defaulting to energy-inefficient functions that scale the environmental problem at the software level (Sikand et al., 2024).

This gluttony extends to water. Data centers are insatiably thirsty, often evaporating potable water to cool servers that run hot with AI calculations. A simple conversation with a chatbot (roughly 20-50 questions) effectively "drinks" a 500ml bottle of water (Ren, 2023). In water-stressed regions, this consumption competes directly with human needs, but companies often obscure these local impacts behind global sustainability reports.

The "Shit" of Embodied Carbon

Beyond electricity, the "gluttonous fraud" includes the "shit" of electronic waste. The relentless demand for newer, faster chips accelerates hardware obsolescence. We are discarding functional hardware to chase marginal gains, creating mountains of toxic e-waste and necessitating the carbon-heavy mining of rare earth metals like cobalt and lithium (Ergönül, 2025; Zewe, 2025). This embodied carbon, the energy used to manufacture the device before it is ever turned on, is often conveniently left out of corporate carbon neutrality claims but can represent a significant portion of the total footprint. A full Life Cycle Assessment (LCA) reveals that end-user terminals and networks account for 85% of the abiotic depletion potential (mineral consumption) and 45% of the carbon footprint of a GenAI service (Berthelot et al., 2025).

G for Green: Genuine Intelligence

In stark contrast, "Green AI" represents genuine intelligence because it treats carbon efficiency not as an afterthought, but as a primary evaluation metric alongside accuracy. Green AI distinguishes between "Green-by AI" (using AI to solve environmental problems) and "Green-in AI" (making the AI itself environmentally friendly) (Bolón-Canedo et al., 2024). Genuine intelligence acknowledges that running a massive, billion-parameter model to perform a simple task is architectural stupidity. It advocates for "sparse" models that prune unnecessary parameters, quantization to lower precision requirements without losing utility, and the use of specialized hardware like Tensor Processing Units (TPUs) or FPGAs that can deliver 15-30% power savings over general-purpose GPUs (Tabbakh et al., 2024). Knowledge distillation, where a smaller "student" model mimics a larger "teacher," can reduce model size by 40% and increase speed by 60% while retaining 97% of performance (Sanh et al., 2019; Tabbakh et al., 2024).

Conclusion and Implications: Toward a "Slow Urgency"

The trajectory of Generative AI, as it stands, is a "gluttonous fraud" because it sells the illusion of frictionless intelligence while hiding a heavy, physical toll on the planet's resources. We are witnessing a "computational opulence paradox" where we deploy nuclear-grade computing power for trivialities, creating a "digital forest" that consumes the resources of the real one. The "shit" of this model is its wastefulness, both in energy and in hardware, and its reliance on opacity to avoid accountability.

However, a shift to "Genuine Intelligence" is possible. This requires a transition from "Red AI" to "Green AI," where efficiency is prized over raw size. Implications for the future are stark:

- **Regulatory Intervention:** We must move beyond voluntary reporting. Policies like the EU AI Act are essential to mandate transparency regarding the energy efficiency of high-risk models (Bolón-Canedo et al., 2024). Governments must regulate foundation models as a new class of industrial product, mandating standardized environmental disclosures (Bozkurt, 2025).
- **Relational Footprinting:** We need assessment frameworks that account for local impacts (water stress in specific counties) rather than just global averages, ensuring that data centers do not cannibalize local resources (Bashir et al., 2024; Pasek et al., 2023). This includes geographical load balancing to shift workloads to regions with cleaner grids (Ding et al., 2025; Ren & Wierman, 2024).
- **Adopting "Slow Urgency":** We should embrace the concept of "slow urgency" (Akomolafe, 2023; Bashir et al., 2024); a measured approach that refuses the frenetic "gold rush" pace in favor of

thoughtful, sustainable development that considers long-term societal and environmental health. This means prioritizing "Green Lean" operational strategies and demanding full Environmental Impact Reports before procurement (Bozkurt, 2025; Setyadi et al., 2025).

As a final remark, it should be noted that true intelligence is not defined by *the ability to generate the most pixels per second*, but by *the wisdom to do so without consuming the world that sustains us*.

References

- Akomolafe, B. (2023). *A Slower Urgency*. <https://www.bayoakomolafe.net/post/a-slower-urgency>
- Bashir, N., Donti, P., Cuff, J., Sroka, S., Ilic, M., Sze, V., Delimitrou, C., & Olivetti, E. (2024, March 27). *The climate and sustainability implications of generative AI*. An MIT Exploration of Generative AI From Novel Chemicals to Opera. <https://mit-genai.pubpub.org/pub/8ulgrckc/release/2>
- Berthelot, A., Caron, E., Jay, M., & Lefèvre, L. (2025). Understanding the environmental impact of generative AI services. *Communications of the ACM*, 68(7), 46–53. <https://doi.org/10.1145/3725984>
- Bolón-Canedo, V., Morán-Fernández, L., Cancela, B., & Alonso-Betanzos, A. (2024). A review of green artificial intelligence: Towards a more sustainable future. *Neurocomputing*, 599, 128096. <https://doi.org/10.1016/j.neucom.2024.128096>
- Bozkurt, A. (2024). GenAI et al.: Cocreation, Authorship, Ownership, Academic Ethics and Integrity in a Time of Generative AI. *Open Praxis*, 16(1), 1–10. <https://doi.org/10.55982/openpraxis.16.1.654>
- Bozkurt, A. (2025). AI's Thirst, AI's Heat, AI's Waste: Exposing the Hidden Environmental Impact of Every Artificial Intelligence Interaction. *Open Praxis*, 17(4), 638–647. <https://doi.org/10.55982/openpraxis.17.4.1058>
- Chen, S. (2025). How much energy will AI really consume? The good, the bad and the unknown. *Nature*, 639(8053), 22-24. <https://doi.org/10.1038/d41586-025-00616-z>
- Chien, A. A., Lin, L., Nguyen, H., Rao, V., Sharma, T., & Wijayawardana, R. (2023). Reducing the Carbon Impact of Generative AI Inference (today and in 2035). Proceedings of the 2nd Workshop on Sustainable Computer Systems, 1-7. <https://doi.org/10.1145/3604930.3605705>
- Ding, Z., Wang, J., Song, Y., Zheng, X., He, G., Chen, X., Zhang, T., Lee, W.-J., & Song, J. (2025). Tracking the carbon footprint of global generative artificial intelligence. *The Innovation*, 6(5), 100866. <https://doi.org/10.1016/j.xinn.2025.100866>
- Ergönül, D. (2025). *The Digital Forest: Exploring AI's Ecological Footprint and Sustainability Challenges*. Bachelor Thesis, University of Oulu. <https://doi.org/10.5281/zenodo.17189940>
- Hao, K. (2019). *Training a single AI model can emit as much carbon as five cars in their lifetimes*. MIT technology Review. <https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/>
- Heikkilä, M. (2023, December 1). *Making an image with generative AI uses as much energy as charging your phone*. MIT Technology Review. <https://www.technologyreview.com/2023/12/01/1084189/making-an-image-with-generative-ai-uses-as-much-energy-as-charging-your-phone/>

- O'Donnell, J., & Crownhart, C. (2025, May 20). We did the math on AI's energy footprint. Here's the story you haven't heard. *MIT Technology Review*. <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>
- Pasek, A., Vaughan, H., & Starosielski, N. (2023). The world wide web of carbon: toward a relational footprinting of information and communications technology's climate impacts. *Big Data & Society*, 10(1). <https://doi.org/10.1177/20539517231158994>
- Ren, S., & Wierman, A. (2024, July 15). The uneven distribution of AI's environmental impacts. *Harvard Business Review*. <https://hbr.org/2024/07/the-uneven-distribution-of-ais-environmental-impacts>
- Ren, S. (2023, November 30). *How much water does AI consume? The public deserves to know*. OECD.AI. <https://oecd.ai/en/wonk/how-much-water-does-ai-consume>
- Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv*. <https://doi.org/10.48550/arXiv.1910.01108>
- Setyadi, A., Pawirosumarto, S., & Damaris, A. (2025). Sustainable operations strategy in the age of climate change: Integrating green lean practices into operational excellence. *Sustainability*, 17(13), 5954. <https://doi.org/10.3390/su17135954>
- Sikand, S., Mehra, R., Sharma, V. S., Kaulgud, V., Podder, S., & Burden, A. P. (2024). Do Generative AI Tools Ensure Green Code? An Investigative Study. Proceedings of the 2nd International Workshop on Responsible AI Engineering, (pp. 52-55). Lisbon, Portugal. <https://doi.org/10.1145/3643691.3648588>
- Tabbakh, A., Al Amin, L., Islam, M., Mahmud, G. M. I., Chowdhury, I. K., & Mukta, M. S. H. (2024). Towards sustainable AI: A comprehensive framework for Green AI. *Discover Sustainability*, 5(408). <https://doi.org/10.1007/s43621-024-00641-4>
- Zewe, A. (2025, January 17). *Explained: Generative AI's environmental impact*. MIT News. <https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117>

About the Author(s)

- Aras Bozkurt; arasbozkurt@gmail.com; Anadolu University, Türkiye. <https://orcid.org/0000-0002-4520-642X>
- Ramesh C. Sharma; rc_sharma@yahoo.com; Dr B R Ambedkar University Delhi, India. <https://orcid.org/0000-0002-1371-1157>

Author's Contributions (CRediT)

Aras Bozkurt: Conceptualization, Methodology, Visualization, Writing – original draft, Writing – review & editing; Ramesh C Sharma: Writing – original draft, Writing – review & editing.

Sustainable Development Goals (SDGs)

This study is linked to the following SDG(s): Clean water and sanitation (SDG 6), Affordable and clean energy (SDG 7), Industry, innovation and infrastructure (SDG 9), Responsible consumption and production (SDG 12), Climate action (SDG 13), Life below water (SDG 14), Life on land (SDG 15).

Acknowledgements

I extend my deepest gratitude to the *Greenpeace Movement* for their relentless courage in exposing the environmental reality of the digital age. Their groundbreaking work has been instrumental in piercing the "black box" of corporate opacity. By tirelessly documenting the hidden costs of our "digital romance", from the fossil fuels powering chip manufacturing in East Asia to the water stress in local communities, Greenpeace has forced a necessary global reckoning. Their commitment to "digital-sustainable co-transformation" and their refusal to let the "cloud" remain an abstract concept have provided the ethical and empirical foundation for critical inquiries like this one.

Funding

This paper is funded by Anadolu University with grant number YTS-2024-2559 and YTS-2025-2776.

Ethics Statement

Because this study doesn't involve any living entities, an ethics review is not applicable.

Conflict of Interest

The authors do not declare any conflicts of interest.

Data Availability Statement

Data sharing is not applicable to this article as no datasets were generated or analyzed during the current study.

Author(s) Notes

Based on *Academic Integrity and Transparency in AI-assisted Research and Specification Framework* (Bozkurt, 2024), the authors of this paper acknowledge that the paper was reviewed, edited, and refined with the assistance of Google's Gemini (Versions as of December 2025), complementing the human editorial process. The human authors critically assessed and validated the content to maintain academic rigor. The authors also assessed and addressed potential biases inherent in the AI-generated content. The final version of the paper is the sole responsibility of the human authors.

Article History

Submitted: November 20, 2025 – Accepted: December 13, 2025.

Suggested citation:

Bozkurt, A., & Sharma, R. C. (2025). G for Generative is Gluttonous Fraud and Shit, G for Green is Genuine Intelligence: Exploring the Unspoken Environmental Cost and the Computational Opulence Paradox in the Artificial Intelligence Landscape. *Asian Journal of Distance Education*, 20(2), i-vi. <https://doi.org/10.5281/zenodo.17665240>



Authors retain copyright. Articles published under a Creative Commons Attribution 4.0 (CC-BY) International License. This licence allows this work to be copied, distributed, remixed, transformed, and built upon for any purpose provided that appropriate attribution is given, a link is provided to the license, and changes made were indicated.