

Title	Environmental contradictions of large language models in higher education
Authors	O'Sullivan, James; Kelly, Mark
Publication date	2025-12-16
Original Citation	O'Sullivan, J & Kelly, M 2025, 'Environmental contradictions of large language models in higher education', UNESCO Ideas LAB. < https://www.unesco.org/en/articles/environmental-contradictions-large-language-models-higher-education >
Type of publication	Article
Link to publisher's version	https://www.unesco.org/en/articles/environmental-contradictions-large-language-models-higher-education
Rights	© 2026, the Authors and the UNESCOs. - https://creativecommons.org/licenses/by-sa/4.0/
Download date	2026-09-25 01:53:44
Item downloaded from	https://hdl.handle.net/10468/18706

**unesco**

Press and news

Calendar

Partnerships

Governance

Quick Links 

English ▼

Idea >

Environmental contradictions of large language models in higher education

Universities are racing to adopt AI, even as its environmental costs remain largely obscured. Can higher education pursue responsible innovation while upholding sustainability and its commitment to the public good?

16 December 2025 - Last update: 6 January 2026



©Vecteezy/Asih Wahyuni

Summary

- Material infrastructures and hidden environmental costs
- Complicity, greenwashing, and the university

- References
- Biography

By James O'Sullivan and Mark Kelly

The emergence of generative artificial intelligence (AI) systems, particularly, large language models, holds the potential to transform higher education to an extent not witnessed since the emergence of the web some 30 years ago.

Large language models (LLMs) are a type of generative AI system trained on vast quantities of text data to generate, interpret, and manipulate human language with remarkable fluency. Built using deep learning architectures—most commonly the transformer model introduced in 2017—LLMs are designed to predict the probability of a given word or sequence of words based on the context in which they appear. This predictive capacity allows them to perform a wide range of language-related tasks, including writing essays, summarising documents, answering questions, and translating between languages.

What distinguishes LLMs like OpenAI's GPT series, Google's Gemini, and Anthropic's Claude from earlier forms of natural language processing is their scale: they are trained on billions—or even trillions—of words scraped from (often plagiarised) books, websites, academic papers, and other digital sources. While their 'understanding' is statistical rather than semantic, they exhibit a form of general linguistic competence which has prompted both enthusiasm and concern.

"To sustain the university as a site of ethical, critical, and democratic inquiry, further explorations of what LLMs can do should be deprivileged, even paused, in favour of analyses of how they are reshaping the values and viability of education itself. "

Framed by proponents as engines of efficiency, access, and pedagogical innovation, LLMs are increasingly integrated into university platforms, tools, and curricula. But the narratives of progress that accompany this rapid integration often obscure the environmental consequences of such technologies. The widespread deployment of LLMs in higher education raises urgent questions about sustainability—not only in terms of technological capacity or institutional readiness, but in relation to the material viability of the infrastructures that underwrite AI.

To sustain the university as a site of ethical, critical, and democratic inquiry, further explorations of *what* LLMs can do should be deprivileged, even paused, in favour of analyses of *how* they are reshaping the values and viability of education itself. At a moment when universities are committing to sustainability targets and climate leadership, the adoption of computationally intensive generative AI tools presents a paradox. Can the university continue to position itself as a site of sustainability responsibility and leadership while embracing technologies that are fundamentally extractive?

Material infrastructures and hidden environmental costs

Despite their disembodied appearance on our screens, LLMs are grounded in highly material infrastructures. Training a large-scale model requires vast computational resources. Researchers calculate that training and deploying LLMs can, depending on model architecture and scale, emit tens to hundreds of metric tonnes of CO₂ (Faiz et al. 2024). It is estimated that training a single large model could emit over 626,000 pounds of carbon dioxide, the equivalent of 300 round-trip flights between New York and San Francisco (Strubell et al. 2020), or five times the lifetime emissions of an average car, including its manufacture, in the US (Hao 2019).

Commercial LLM providers treat precise energy consumption and carbon emission data as proprietary, making accurate environmental impact assessments difficult or impossible for external stakeholders, including universities (Dodge et al. 2022). This lack of transparency effectively shields these providers from scrutiny and complicates efforts by institutions to align their technology procurement with explicit sustainability goals. Furthermore, such models are being developed at a time when the realities of climate inequality have never been so stark (Bender et al. 2021).



"Despite their disembodied appearance on our screens, LLMs are grounded in highly material infrastructures. "

Whatever the reality of the resource requirements of LLMs, it is well established that data centre energy consumption has significant repercussions for social equity and fuel poverty. Data centres dedicated to training and deploying large-scale AI models extract considerable energy resources from national grids,

exacerbating the inequitable distribution of energy in societies already facing acute fuel poverty. Data centres in the EU used an estimated 45-65 TWh of electricity in 2022, representing 1.8-2.6% of total EU electricity use (Kamiya & Bertoldi 2024). In Ireland, data centres consumed 21% of all metered electricity from the grid in 2023, an increase of 3% from 2022 (Walsh 2024). This is at a time when 377,000 people in Ireland were unable to keep their house adequately warm in 2022 (Joint Committee 2024).

This tension between technological advancement and basic energy access directly challenges the [United Nations Sustainable Development Goals \(SDGs\)](#), particularly SDGs 3, 7, 10, 12, and 13, which aim to ensure healthy lives and promote well-being for all; ensure access to affordable, reliable, sustainable, and modern energy for all; reduce inequality and promote responsible consumption and production patterns; and take urgent action to combat climate change and its impacts. As universities increasingly adopt computationally intensive AI technologies, they risk contributing to energy inequities, effectively prioritising technological innovation over the fundamental needs and rights of disadvantaged communities. Such choices necessitate deeper institutional reflection on the broader societal implications of technological adoption, and they underline the importance of embedding equity considerations within higher education's sustainability strategies.

"Unlike the corporate actors who develop and profit from these tools, universities bear a different kind of responsibility: to model values and practices oriented towards the public good. As such, academic engagement with LLMs must include environmental reckoning. "

It is not only training that carries environmental costs. The inference phase—when models are actually used in real-world applications—also requires significant computational power. As universities increasingly integrate LLMs into their platforms and services, the cumulative energy demands of routine student and staff use begin to scale dramatically. A 2023 study estimated that a single user prompting GPT-3 for just 50 queries a day over a year could consume as much electricity as a modern refrigerator. When scaled to millions of students i.e. 235 million student enrolled in higher education in 2020 and growing (UNESCO, 2022) and staff across multiple institutions, the resource implications become staggering.

Furthermore, these systems rely on global data centre networks, many of which are powered by fossil fuels and located in regions already experiencing environmental stress. Water usage is another hidden cost: cooling data centres consumes millions of litres, often drawn from vulnerable local sources. In regions affected by drought or climate instability, the expansion of AI infrastructure poses direct risks to local ecosystems and communities.

Universities, often committed to the language of sustainability in their strategic plans, find themselves in a precarious position. LLMs offer—or certainly promise—potential remedies to the pressures of massification, but their adoption deepens institutional entanglement with carbon-intensive digital economies.

The discourse around LLMs often celebrates their frictionless utility while obscuring the ecological externalities of their existence. In adopting such systems at scale, higher education institutions risk becoming complicit in extractive digital economies that are environmentally unsustainable. Yet, unlike the corporate actors who develop and profit from these tools, universities bear a different kind of responsibility: to model values and practices oriented towards the public good.

As such, academic engagement with LLMs must include environmental reckoning. This might involve demanding transparency from vendors regarding energy use and training data; supporting low-carbon, open-source alternatives; or embedding climate literacy into digital policy. Sustainability in the era of LLMs begins with acknowledging that even the most intangible forms of intelligence require material substrates—and that these substrates are finite.

Complicity, greenwashing, and the university

The university is not a neutral actor in the story of AI. As early adopters, legitimators, and sites of technological reproduction, academic institutions play a key role in shaping the trajectory of LLM integration into society. When universities endorse or license AI systems developed by commercial actors, they do more than adopt a tool—they participate in the construction of a social imaginary in which such systems are necessary, desirable, and inevitable.

This complicity is often masked by sustainability rhetoric. Institutions may highlight the energy efficiencies of automating administrative processes or reducing paper usage while overlooking the broader environmental toll of the systems they deploy. Such gestures amount to a form of academic

greenwashing: the selective framing of digital adoption as ecologically virtuous without full transparency about its material impacts.

Universities are uniquely positioned to resist such reductive logics. They are among the few remaining institutions with both the expertise and the ethical mandate to challenge unsustainable technological trajectories. In the context of LLMs, this means moving beyond reactive accommodation and taking a proactive role in shaping responsible innovation.

To engage seriously with environmental sustainability, higher education must adopt a framework of environmental accountability in its dealings with AI. This begins with transparency. Universities should demand clear environmental impact assessments from AI vendors, including disclosures of training energy, water usage, and ongoing inference costs. Just as ethical data usage and privacy considerations have become standard in procurement, so too should environmental criteria.

"Institutions may highlight the energy efficiencies of automating administrative processes or reducing paper usage while overlooking the broader environmental toll of the systems they deploy. Such gestures amount to a form of academic greenwashing: the selective framing of digital adoption as ecologically virtuous without full transparency about its material impacts."

Institutions must consider investing in or co-developing lower-carbon alternatives to proprietary models. Open-source LLMs, such as BLOOM or Meta's LLaMa, can be fine-tuned locally, offering opportunities for energy-efficient, targeted models trained on smaller and more relevant datasets. These approaches align with the principles of 'small data' and 'minimal computing', which privilege context-sensitive, low-impact digital practices.

Digital literacy initiatives in universities must expand to include climate literacy. If students are to engage critically with AI, they must understand not only its functional capacities but also its ecological footprints. Curriculum design should address the environmental dimensions of digital tools, fostering a generation of scholars and professionals who see sustainability as integral to technological fluency.

Sustainability considerations must be integrated into institutional policy frameworks. This includes procurement policies, platform governance, and academic planning. Where LLMs are deployed, their use should be justified not only in terms of pedagogical or administrative benefit, but also in relation to long-term environmental strategy.

Finally, universities should work collaboratively—across regions and sectors—to establish shared norms and benchmarks for sustainable AI. Given the global nature of both climate crisis and digital infrastructure, isolated action is insufficient. Networks of higher education institutions can leverage collective bargaining power to demand change from technology providers and to model alternative futures grounded in planetary ethics.

"Curriculum design should address the environmental dimensions of digital tools, fostering a generation of scholars and professionals who see sustainability as integral to technological fluency."

Large language models mark a profound shift in how we produce, distribute, and interact with knowledge. Their integration into higher education is both an opportunity and a warning. If adopted uncritically, universities risk further entrenching extractive and unsustainable digital infrastructures under the guise of innovation. If approached with care, however, they offer an occasion to reassert their public mission to shape technologies in ways that are just, inclusive, and ecologically attuned. The future of higher education depends not on how fast we adopt new tools, but on how wisely we decide which tools—and which futures—are worth sustaining.

References

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 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*, 610–23. FAccT '21. New York, NY, USA: Association for Computing Machinery.

<https://doi.org/10.1145/3442188.3445922> .

Dodge, Jesse, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A. Smith, Nicole DeCario, and Will Buchanan. 2022. 'Measuring the Carbon Intensity of AI in Cloud Instances.' In *Proceedings of the 2022 ACM*

Conference on Fairness, Accountability, and Transparency, 1877–94. FAccT '22. New York, NY,

USA: Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533234> .

Faiz, Ahmad, Sotaro Kaneda, Ruhan Wang, Rita Osi, Prateek Sharma, Fan Chen, and Lei Jiang.

2024. 'LLMCarbon: Modeling the End-to-End Carbon Footprint of Large Language Models'. *arXiv*.

<https://doi.org/10.48550/arXiv.2309.14393> .

Hao, Karen. 2019. 'Training a Single AI Model Can Emit as Much Carbon as Five Cars in Their Lifetimes'. *MIT Technology Review*, 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/> .

Joint Committee on Social Protection, Community and Rural Development, and the Islands. 2024.

Report on Energy Poverty. 33/JCSPCRI/026. Houses of the Oireachtas. <https://urlr.me/6GbWjc> .

Kamiya, George, and Paolo Bertoldi. 2024. *Energy Consumption in Data Centres and Broadband Communication Networks in the EU*. JRC135926. Luxembourg: Publications Office of the European

Union. <https://data.europa.eu/doi/10.2760/706491> .

Strubell, Emma, Ananya Ganesh, and Andrew McCallum. 2020. 'Energy and Policy Considerations for Modern Deep Learning Research'. In *Proceedings of the AAAI Conference on Artificial*

Intelligence, 34 (9): 13693–96. <https://doi.org/10.1609/aaai.v34i09.7123> .

UNESCO. 2022. *Higher Education Global Data Report: A Contribution to the World Higher*

Education Conference 18–20 May 2022. Paris: United Nations Educational, Scientific and Cultural

Organization. <https://www.right-to-education.org/resource/higher-education-global-data-report>

.

Walsh, Katie. 2024. *L&RS Spotlight: Data Centres and Energy*. No. 2 of 2024. Dublin: Oireachtas Library & Research Service.

https://data.oireachtas.ie/ie/oireachtas/libraryResearch/2024/2024-07-23_spotlight-data-centres-and-energy_en.pdf .

The ideas expressed here are those of the authors; they are not necessarily the official position of UNESCO and do not commit the Organization. Unless otherwise indicated, content published on the UNESCO IdeasLAB platform is licensed under the Creative Commons Attribution-ShareAlike 3.0 IGO licence ([CC BY-SA 3.0 IGO](https://creativecommons.org/licenses/by-sa/3.0/igo/) .