STIAEM

Slow Tech Innovation

AI, Ethics and New Teaching Methods

Environmental Footprint of Generative AI: Energy and Water Demands of Data Centres

Co-funded by the European Union

2025-1-IT02-KA220-SCH-000365213

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Education and Culture Executive Agency (EACEA). Neither the European Union nor EACEA can be held responsible for them... Project reference code: 2025-1-IT02-KA220-SCH-000365213

Università di Foggia

Studi Umanistici, Lettere, Beni Culturali, Scienze della Formazione

İSTANBUL - ESENLER İLÇE MİLLİ EĞİTİM MÜDÜRLÜĞÜ

1474

STIAEM

Slow Tech Innovation
AI, Ethics and
New Teaching Methods

Co-funded by
the European Union

2025-1-IT02-KA220-SCH-000365213

eduOpen



CC BY-NC-SA 4.0 International

This license requires that reusers give credit to the creator. It allows reusers to distribute, remix, adapt, and build upon the material in any medium or format, for noncommercial purposes only. If others modify or adapt the material, they must license the modified material under identical terms.

OER 5.3.1: Environmental Footprint of Generative AI: Energy and Water Demands of Data Centres
Reference Module: Module 5 - Safety, Privacy and Sustainability
Related Video-Pill: Video 5.3 - Ecological Footprint & Civic Education
Authoring partner: Konak Science and Art Center (BILSEM) - Türkiye
Topic / Thematic Focus: Environmental footprint, generative AI resource demands and data center energy-water consumption

Introduction and Learning Objectives

This resource, together with OER 5.3.2, accompanies video-pill 5.3 «Ecological Footprint & Civic Education» of Module 5. Its purpose is to make comprehensible, even to those without technical training, the environmental footprint of generative artificial intelligence: how much energy it consumes, how many emissions it produces, how much water and how many material resources it requires. These are data that the school can use for a concrete civic education.

The text connects to competence 4.4 of DigComp 3.0, «Environmental impacts of digital technologies» (Cosgrove & Cachia, 2025), and to the principles of «Clean IT» of Slow Tech (Patrignani & Whitehouse, 2019). The main sources are real scientific studies: Strubell, Ganesh and McCallum (2019), Schwartz, Dodge, Smith and Etzioni (2020), Luccioni, Viguier and Ligozat (2022).

On the institutional level, the topic is framed by the OECD report (2022) on measuring the environmental impacts of computing and of artificial intelligence applications.

By the end of the reading, the teacher will be able to:

- explain with concrete data the energy and emissions cost of training and using AI models;
- distinguish the different layers of the footprint: energy, carbon, water and material resources;
- use the concept of computational efficiency (Green AI) as a criterion of didactic choice.

1. The discovery of a hidden cost

Until a few years ago the environmental cost of artificial intelligence was invisible. The credit for bringing it to the attention of the scientific community goes to a group of NLP researchers, who in 2019 quantified for the first time systematically the energy needed to train language models.

«In this paper we bring this issue to the attention of NLP researchers by quantifying the approximate financial and environmental costs of training a variety of recent successful neural network models for NLP» (Strubell, Ganesh & McCallum, 2019).

The numbers are illuminating because they are compared with familiar magnitudes. Training a large model is not a «light» operation: the fine-tuning and experimentation phase alone can produce emissions comparable to those of a flight; if automated architecture search is added, the average annual emissions of a person are exceeded. In the authors' table, for example, training with fine-tuning reaches 78,468 pounds of CO2 and, with architecture search,

626,155 pounds, while a flight from New York to San Francisco weighs about 1,984 pounds (Strubell et al., 2019).
The point is not to demonise AI, but to make its cost visible: what is not seen cannot be governed. And since training requires resources that few can afford, the environmental cost also becomes a question of equity.
The measurement of the footprint also has a political, as well as technical, value. Since the training of the largest models requires resources that only a few organisations can afford, the environmental cost is intertwined with a question of concentration of power: whoever controls the computing also controls the possibility of developing AI. Bringing this reflection into the classroom means linking sustainability to justice, two themes that civic education naturally holds together.

2. The growing trend and the Green AI proposal

Subsequent studies showed that the phenomenon was not episodic, but structural.
Research on deep learning, measured in computing power, is growing at vertiginous rates, with a surprisingly large carbon footprint.
«The computations required for deep learning research have been doubling every few months, resulting in an estimated 300,000x increase from 2012 to 2018. These computations have a surprisingly large carbon footprint» (Schwartz, Dodge, Smith & Etzioni, 2020).
From this observation comes the proposal of «Green AI», a way of doing research that holds together accuracy and efficiency and also takes on the economic and social cost of the results.
«We advocate increasing research activity in Green AI - AI research that is more environmentally friendly and inclusive» (Schwartz et al., 2020).
The lesson for the school is twofold. On the one hand, computational efficiency becomes a value to be taught: the largest model is not always the best choice. On the other hand, inclusion and sustainability are two sides of the same coin: an AI that consumes enormous resources excludes those who have least.
The criterion of efficiency is not an invitation to stop research, but to measure it more honestly. Assessing a model only on its accuracy, neglecting the resources it required to obtain it, is equivalent to accounting for the revenues and forgetting the costs. The Green AI proposal instead asks to make both terms explicit, so that technical progress is not confused with an indiscriminate increase in computing power.

3. The concrete measurement: the BLOOM case

To move from order of magnitude to precise number, a 2022 study calculated the footprint of BLOOM, a language model with 176 billion parameters, over its entire life cycle.

«We estimate that BLOOM's final training emitted approximately 24.7 tonnes of CO2eq if we consider only the dynamic power consumption, and 50.5 tonnes if we account for all processes ranging from equipment manufacturing to energy-based operational consumption» (Luccioni, Viguier & Ligozat, 2022).

The difference between the two values is instructive: considering only electricity consumption means neglecting the production of the hardware and its supply chain. The study also shows that the cost does not end with training, but continues in the daily use of the model: operation through an interface entails emissions even when it receives no requests, and a significant share of the energy consumed remains «wasted».

«this results in approximately 19 kgs of CO2eq emitted per day of API deployment, or 340 kg over the total period during which we were tracking emissions» (Luccioni et al., 2022).

To this is added the chapter of water and materials. The production of computing devices requires large quantities of ultra-pure water, minerals and chemicals, and transport along carbon-intensive supply chains: a cost that, as the authors observe, often remains little transparent.

«More transparency is needed regarding the environmental impacts of manufacturing computing equipment given the large quantities of chemicals and minerals required (...) and the significant quantities of ultra-pure water and energy needed to manufacture it» (Luccioni et al., 2022).

The distinction between «embodied carbon» and «operational carbon» is a concept useful to make students understand: part of the footprint arises from the production and disposal of devices, another from their operation. The sum of the two explains why even an apparently «light» model, if used by millions of people every day, weighs significantly on the environmental level.

4. The institutional framework and competence 4.4

The topic has also entered the agenda of institutions. An OECD report of 2022 addressed the measurement of the environmental impacts of computing and of AI applications, proposing a methodological framework to quantify energy, water and carbon. The source is available only as an online reference; here we take up its meaning in paraphrase, without textual quotations.

On the level of competences, the European reference is DigComp 3.0, which includes environmental impact in Area 4, dedicated to safety, well-being and responsible use.

«4.4 Environmental impacts of digital technologies», in Area 4 «Safety, well-being and responsible use, including the environmental impact of digital technologies» (DigComp 3.0, It. trans. 2025).

The document explicitly links digital devices to sustainability and the green transition, and asks that people be informed about the environmental impact and energy consumption of the technologies they use. Sustainability, in this framework, is not an appendix: it is a citizenship competence.

The link with the European Declaration on Digital Rights and Principles is explicit: the environmental sustainability of digital technologies figures among the principles that orient

the digital transformation of the Union. It follows that this is not an optional topic, but a structural dimension of digital citizenship, which the school is called to teach together with the other rights and duties.

5. What we can do in the classroom

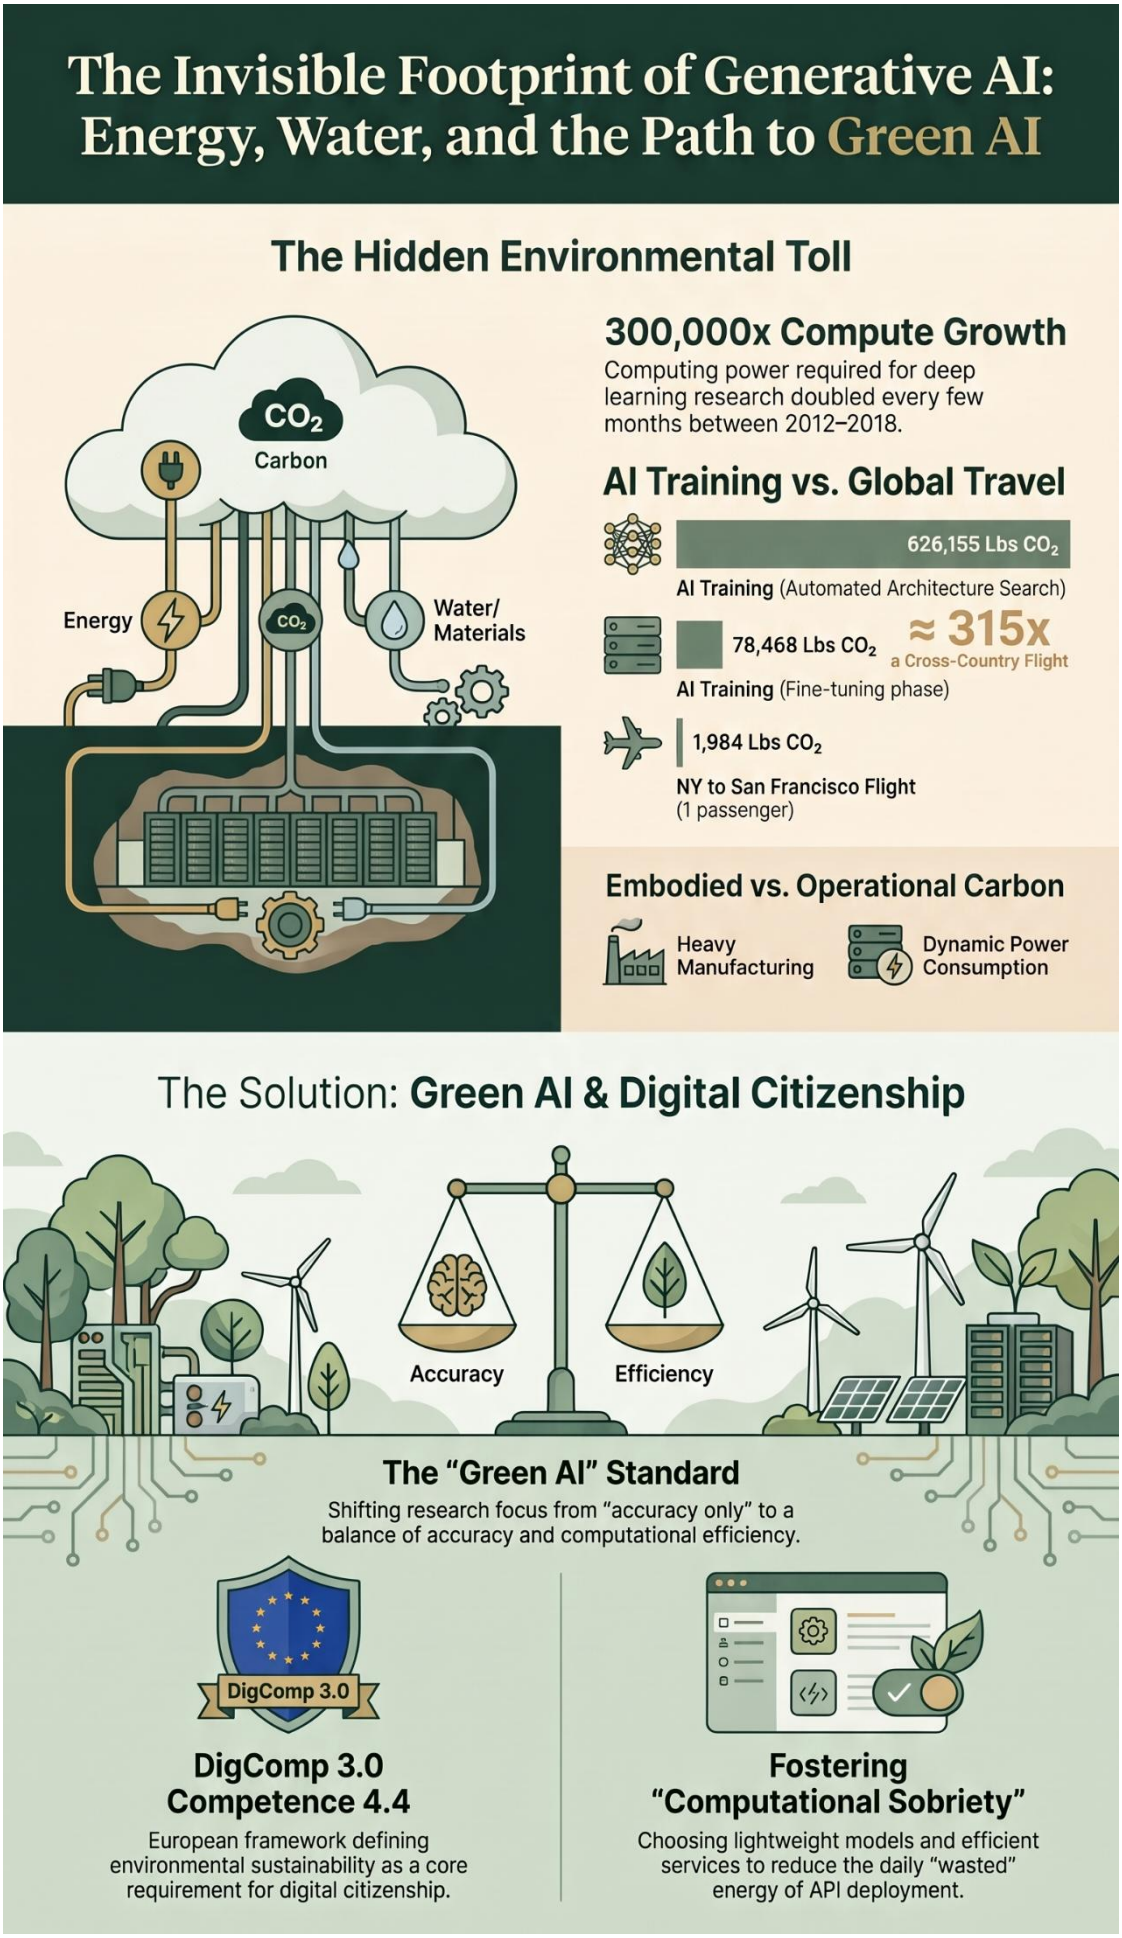
If the footprint is large, the possible action is made up of everyday choices. Using smaller models when they suffice, avoiding useless processing, turning off devices, preferring efficient services, reducing file size: these are modest gestures, but consistent with the logic of Green AI. Even more important is making the class capable of reading a consumption figure and comparing two options.

The following table summarises the dimensions of the footprint and the questions to ask in the classroom.

A simple exercise, suitable for several disciplines, consists in reconstructing the «journey» of a digital activity: where the servers are, with what energy they are powered, how much data is transferred, for how long. Even without precise numbers, the reconstruction develops an awareness that adults too often lack: every digital gesture has a physical infrastructure, made of machines, electricity and materials. It is this awareness that makes a more sober choice possible.

Impact category	Indicator	Order of magnitude (real sources)	Question for the class
Energy	Power and consumption of GPUs/TPUs	Training with tuning: 78,468 lb CO2; with architectures: 626,155 (Strubell et al., 2019)	Why does it take so much energy to train a model?
Carbon	Emissions in CO2eq	BLOOM: 24.7 t (consumption) and 50.5 t (life cycle) (Luccioni et al., 2022)	How much CO2 does a model we use weigh?
Water and materials	Ultra-pure water, minerals	Large quantities for producing devices (Luccioni et al., 2022)	What lies behind the device we hold in our hands?
Trend	Growth of computing	About 300,000 times from 2012 to 2018 (Schwartz et al., 2020)	Is this growth sustainable?
Response	Efficiency (Green AI)	Measuring accuracy and efficiency together (Schwartz et al., 2020)	Can we do the same thing while consuming less?

Visual Scaffolding



Conclusions and Metacognitive Reflection Prompts

This resource has made visible the environmental footprint of generative AI. The training of models requires enormous quantities of energy and produces emissions comparable to familiar consumption; the cost does not stop at training, but continues in use, and also involves water and materials. The growth of computing has pushed the scientific community to propose Green AI, which measures accuracy and efficiency together. DigComp 3.0 has taken up competence 4.4 on the environmental impact of digital technologies: sustainability is thus an integral part of digital citizenship.

To transfer these core ideas to one's own practice, the following reflection prompts are proposed:

- Could you estimate, with your students, the energy consumption of one week's digital activities?
- In which classroom tasks would a «lighter» solution achieve the same result with fewer resources?
- How would you connect competence 4.4 of DigComp 3.0 to a civic education project at your school?

Technical Metadata Sheet

Metadata Field	Specification / Parameter Value
Identifier	OER 5.3.1
Title	Environmental Footprint of Generative AI: Energy and Water Demands of Data Centres
Creator	Konak Science and Art Center (BILSEM) - Türkiye
Description	The resource makes comprehensible the environmental footprint of generative AI: how much energy it consumes, how many emissions it produces, and how much water and material resources it requires. It distinguishes the layers of the footprint — energy, carbon, water and materials — and introduces computational efficiency (Green AI) as a didactic criterion, with the concrete case of BLOOM. The practical focus is on what can be done in the classroom, linked to DigComp 3.0 competence 4.4.
Keywords	environmental footprint; energy consumption; carbon emissions; water use; data centres; Green AI; computational efficiency; sustainability
Language	EN (English)
Affiliation	Konak Science and Art Center (BILSEM) - Türkiye
Contributor	STIAEM Transnational Scientific Board (peer review, Rhodes validation protocol)
Media Format	Text; Infographics
Genre	Training Manual; Educational Guide
URL	
Citation	Cosgrove, J., & Cachia, R. (2025). DigComp 3.0: European digital competence framework (Fifth edition) [European framework of digital competences, Italian translation 2025]. Publications Office of the European Union. https://data.europa.eu/doi/10.2760/0001149 Luccioni, A. S., Viguier, S., & Ligozat, A.-L. (2022). Estimating the carbon footprint of BLOOM, a 176B parameter language model. Journal of Machine Learning Research, 24, 1-15. https://arxiv.org/abs/2211.02001 OECD. (2022). Measuring the environmental impacts of artificial intelligence compute and applications. OECD Publishing. https://www.oecd.org/en/publications/measuring-the-environmental-impacts-of-artificial-intelligence-compute-and-applications_7babf571-en.htm Patrignani, N., & Whitehouse, D. (2019). Slow tech: The bridge between

	computer ethics and business ethics. In This changes everything - ICT and climate change (pp. 265-280). Springer. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. Communications of the ACM, 63(12), 54-63. https://arxiv.org/abs/1907.10597 Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 3645-3650). https://aclanthology.org/P19-1355/
Educational Level	Secondary School
Associated Module / MOOC Video Lesson	Module 5 - Safety, Privacy and Sustainability / Video 5.3 - Ecological Footprint & Civic Education
Learning Objectives	explain with concrete data the energy and emissions cost of training and using AI models; distinguish the different layers of the footprint: energy, carbon, water and material resources; use the concept of computational efficiency (Green AI) as a criterion for didactic choices.
Rights	Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)
Date Available	14/09/2026
Date Modified	

References and Webography

Cosgrove, J., & Cachia, R. (2025). DigComp 3.0: European digital competence framework (Fifth edition) [European framework of digital competences, Italian translation 2025]. Publications Office of the European Union. https://data.europa.eu/doi/10.2760/0001149

Luccioni, A. S., Viguier, S., & Ligozat, A.-L. (2022). Estimating the carbon footprint of BLOOM, a 176B parameter language model. Journal of Machine Learning Research, 24, 1-15. https://arxiv.org/abs/2211.02001

OECD. (2022). Measuring the environmental impacts of artificial intelligence compute and applications. OECD Publishing. https://www.oecd.org/en/publications/measuring-the-environmental-impacts-of-artificial-intelligence-compute-and-applications_7babf571-en.htm

Patrignani, N., & Whitehouse, D. (2019). Slow tech: The bridge between computer ethics and business ethics. In This changes everything - ICT and climate change (pp. 265-280). Springer.

Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. Communications of the ACM, 63(12), 54-63. https://arxiv.org/abs/1907.10597

Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 3645-3650). https://aclanthology.org/P19-1355/