

Peer Review

Review of: "Adverse Environmental and Public Health Effects of Artificial Intelligence: A Narrative Review"

Evgenii Arsentev¹

1. Independent researcher

1. What the work claims

In this paper, Jose Domingo looks at the harm that large commercial AI, mostly generative models and LLMs, can do to the environment and to public health. As I read it, the paper follows the whole AI lifecycle, from mineral extraction to e-waste, and looks at it through planetary health and environmental justice. First comes the environmental side, which covers energy and carbon, water, e-waste, and air pollution. Then comes the public health side, which covers algorithmic bias, misinformation and deepfakes, privacy and mental health, job losses, and vulnerable groups.

On method, I understand from Section 2 that this is a narrative review. The author searched PubMed, Web of Science, Scopus, and Google Scholar, plus sources like IEA and UNEP, and does not follow PRISMA. Studies that look only at the benefits of AI were left out on purpose. The main claim, in Section 8, is that the current trajectory brings a net increase in harm. So the author asks for standardized sustainability and bias audits, treaties on mineral sourcing and e-waste, and public funding for green AI and a just transition.

2. Whether the conclusion holds on the data

I would like to start with what I found genuinely valuable in version 3. The review puts environmental pathways and health-equity pathways into one frame (Section 1, Figure 1), and Table 1 gives quantitative anchors that a reader can hold on to. I also appreciate how careful the text is about causality in several places. Section 4.2 says it is still unclear how much AI-generated content, as opposed to other sources, contributes to vaccine hesitancy outcomes. Section 4.4 says that evidence on realized health outcomes from AI-driven job displacement remains limited. Section 5 notes that attributing data-center siting

solely to environmental injustice is challenging without more rigorous analysis, and Section 7 lists the limits of the evidence. I value that these limits are stated openly.

My first concern is about scope by design. Section 2 excludes articles focusing solely on beneficial applications, so the evidence base is one-sided by construction. The author says this plainly and calls for a risk-benefit evaluation as future research (Section 2). Section 8 nevertheless concludes "a net increase in harm". A net judgment needs both sides of the balance, and the method left one side out. I would soften this sentence or present it as the author's view rather than a finding of the review.

My second concern is a single estimate that appears to be counted twice. Section 3.3 gives the same range, 1.2 to 5.0 million tons of AI-related e-waste accumulated over 2020-2030, once attributed to Wang et al. [36] and once to Lannelongue [37] ("1.2-5 million tons"). If [37] reports the estimate of [36], the two sentences count one estimate twice. It would help to say which of the two is the primary source, so readers do not take it for two independent confirmations.

Another point I would strengthen is the estimate in Section 3.4. It says AI's direct contribution to air pollution mortality "is likely measured in tens of thousands of deaths globally rather than millions". The reasoning rests on data centers being "approximately 1-2% of global electricity demand" and on the total burden of "over 9 million premature deaths annually" [52]. The 1-2% figure has no reference in that sentence, and the tens-of-thousands figure has neither a calculation nor a citation. Table 1 then repeats it as "estimated in tens of thousands". I would either show the calculation (share of electricity, share from fossil sources, deaths attributed to power generation) or label it in Table 1 as the author's order-of-magnitude reasoning.

A related gap is in Section 3.1. The paragraph says Google reported a 48% increase in total GHG emissions between 2019 and 2023, and Microsoft a 30% increase from 2020 to 2024. Neither sentence has a reference number, unlike the rest of the paragraph. The next paragraph gives projections of 1% by 2030 and up to 1.4% under accelerated adoption, and again there is no reference after the projection. Since Section 7 asks readers to treat grey literature with caution, these are the figures readers will most want to trace, so it would help to add the sources right there.

One more example of a strong claim is in Section 1, which calls this "the first comprehensive review" to connect the full AI lifecycle with environmental and health harms through an environmental justice lens. A claim of priority is hard to support in a narrative review unless a search for earlier reviews is described. I would phrase it more modestly, for example by saying what this review adds and leaving the question of priority to readers.

3. Reproducibility

Section 2 gives the databases (PubMed, Web of Science, Scopus, Google Scholar), the date window (January 2015 to December 2025), the language (English), and three Boolean search strings. That is more than many narrative reviews report, and I appreciate it. The same section calls the search "a systematic scientific literature search," but it reports no counts: how many records each database returned, how many were screened by title and abstract, and how many were included. The review does not follow PRISMA (Section 2), and I do not ask for a full PRISMA protocol. A short count of records found, screened, and included, plus the date of the last search, would let a reader repeat or update the search.

Two further points would help reuse. Section 2 says "Other specific sources were also consulted (e.g., IEA, UNEP)." It would help to list which reports were actually used, since several Table 1 numbers depend on them. And Section 7 says some quantitative estimates come from preprints and grey literature (arXiv, SSRN). I would mark in Table 1 which figures are peer-reviewed and which are not. This would make the table easier to cite and to build on.

4. What specifically to fix

1. In Section 3.2, the comparison for GPT-3 training says 700,000 liters of freshwater would fill an Olympic-sized swimming pool "to a depth of about 30 cm". I would recheck the arithmetic. An Olympic pool is 50 m by 25 m, which is 1,250 square meters of surface. 700,000 liters equal 700 cubic meters, and $700 / 1,250$ gives 0.56 m. So the depth is about 56 cm, not 30 cm.
2. Section 3.3 and Table 1 give the e-waste share in two wordings. The text says "less than 22.3%" of the 62 million metric tons of e-waste (2022) was formally collected and recycled, and Table 1 says "only 22.3% formally recycled". It would help to use one wording in both places, so nobody reads them as two different figures.
3. Also in Section 3.3, the same e-waste range is written 1.2-5.0 million tons (Wang et al. [36]) and 1.2-5 million tons (Lannelongue [37]). I would name the primary source once and keep a single form of the range.
4. A few statements in Sections 3.1 and 3.4 need references. In Section 3.1, I would add sources for the Google 48% increase in emissions (2019-2023), the Microsoft 30% increase (2020-2024), and the projections of 1% and 1.4% for 2030. In Section 3.4, "approximately 1-2% of global electricity demand" needs a reference as well. For the "tens of thousands" estimate, it would help to show how it was obtained. If it comes from your own reasoning, please mark it that way in Table 1.

5. In Section 6, the paragraph "Accelerated decarbonization and renewable energy" says planning "alleviates grid pressures that inflate residential electricity costs" and cites [10]. In Section 4.1, [10] is Obermeyer et al. on racial bias in a health-care risk algorithm, so this looks like a misplaced reference. Section 5 cites [54] for the same point about electricity costs, and perhaps that is the one intended here.
6. In Section 4.5, on children and adolescents, the claim that recommendation feeds are "exacerbating prefrontal cortex changes linked to reduced focus and reward sensitivity in developing brains" rests on one citation, [65], placed after the paragraph. A neurodevelopmental claim like this deserves its own source, and I would soften the verb unless [65] shows the effect directly. Two small editing points: "[10], This is not an isolated case" in Section 4.1 has a stray comma, and "are probably to be far more expensive" in Section 6 would read better reworded.

5. Recommendation

Minor revisions. The review is useful as a single map of environmental and public-health harms of large-scale AI, and version 3 already states its limits (Sections 4.2, 4.4, 5, 7). The fixes I would suggest are about numbers and sources, not about the structure. Please check the pool arithmetic (Section 3.2), add citations for the corporate and projection figures (Section 3.1), and give a source for the mortality estimate (Section 3.4, Table 1). The e-waste range appears to be double-counted (Section 3.3), and one reference seems misplaced (Section 6). Given the deliberate exclusion of benefit-focused studies (Section 2), it would help to soften the "net increase in harm" sentence (Section 8). All of this can be done without new research.

Declarations

Potential competing interests: No potential competing interests to declare.