Digital Weight: The Invisible Environmental Cost of Our Online Lives

Jekaterina Kaizere
katerina@kaizere.de
Independent artist and designer,
Idar-Oberstein, Germany

Digital Weight is a research-informed media artwork that makes the "immaterial" internet materially legible. A looping video regenerates a near-identical image at a fixed cadence, while on-screen counters accumulate three quantities per iteration: electricity (Wh), carbon (g CO_2), and cooling water (mL). These values derive from published per-unit coefficients: energy per image for text-to-image inference (Luccioni, Jernite, and Strubell 2024), global grid-emission factors of approximately 480 g CO_2/kWh (Ember 2024), and data-centre water-usage estimates of approximately 2 L/kWh (MIT News 2025). The method is simulation-based and auditable: published coefficients are translated through transparent formulas into time-based perception, with impacts reported per fixed quanta, such as per million images, and illustrated through LOW/MID/HIGH scenarios. Conceptually, the work situates the cloud as material infrastructure through Morton's hyperobjects (Morton 2013) and Barad's entanglement (Barad 2007); practically, it turns technical metrics into an embodied, experiential form. The work aims to make the hidden physical cost of digital culture perceptible, so that digital repetition can be understood as an infrastructural and ecological act.

Keywords: Media Installation, Algorithmic Repetition, Infrastructural Aesthetics, Environmental Impact, Cloud Computing, Simulation-Based Modelling, Data Centre Cooling, Generative AI.

DOI 10.34626/2026_xcoax_021

Introduction: Digital Gestures Appear Weightless

Digital gestures appear weightless. Billions of images, streams, and AI prompts circulate each day as if they had no mass, no friction, and no consequence. Interfaces are designed to erase duration: a click becomes an outcome; a prompt becomes an image; a feed refreshes without interruption. The "cloud" is still widely imagined as immaterial and everywhere at once. In practice, each request mobilises material infrastructures – server hardware, electricity from grids, heat dissipation, and water-intensive cooling. The environmental cost is displaced from perception, and that invisibility shapes cultural behaviour: we repeat digital actions because they feel free.

Digital Weight is a research-informed media artwork that makes the supposedly "immaterial" internet materially legible by turning computation into time. The work is presented as a single-channel video loop in which a near-identical image is regenerated at a fixed cadence. With each regeneration, three counters advance: electricity (Wh), carbon emissions (g CO_2), and cooling water (mL). The counters are integrated into the work's generative score: they turn each modelled iteration into a visible record of accumulation. Each iteration is treated as one modelled generative event, translating an inferred resource cost into perceptible form. Duration becomes the medium through which accumulation is perceived.

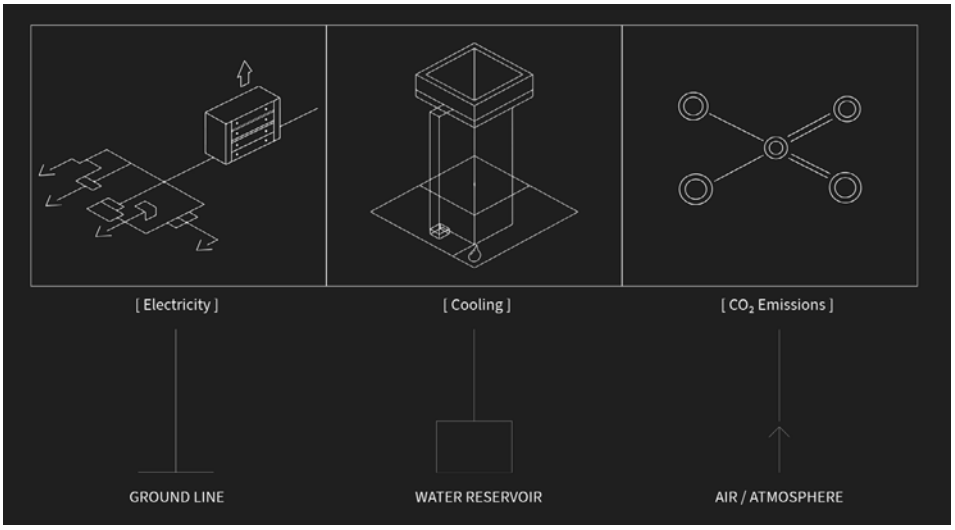

Fig. 1. Infrastructural components mobilised by AI image generation.

Material Infrastructures

At the core of the project lies an engagement with contemporary digital infrastructure. Data centres and networks are increasingly discussed as energy actors, and their growth is closely tied to cloud services and AI workloads. The International Energy Agency estimates that data centres accounted for around 1.5% of the world's electricity consumption in 2024 (IEA 2025), with rapid growth over the last decade and strong geographic concentration in specific clusters. Seemingly "virtual" practices are sustained by a distributed, industrial-scale system.

Electricity, however, is only one layer. Computation produces heat, and heat often requires water to remove it – either directly on-site through cooling systems or indirectly through the water used in electricity generation. In an MIT explainer on generative AI's environmental impact, Noman Bashir notes an estimate of roughly two litres of water per kilowatt-hour of data-centre energy use for cooling (Zewe 2025). This figure is site-dependent, since cooling technologies and locations vary. Here it functions as a communicative coefficient for understanding how energy demand couples to hydrological demand.

Selection of AI Image Generation

From the vast landscape of digital activity, *Digital Weight* deliberately isolates one unit of analysis: AI image generation. This choice is methodological. Image generation is computationally intensive, culturally widespread, and measurable in per-unit terms. One generated image functions as a discrete unit that can be repeated at a fixed cadence, allowing a linear model to be mapped onto time without the need for opaque platform telemetry.

Fig. 2. AI image generation isolated within the broader landscape of digital activity.

Digital Weight uses published inference measurements as a baseline for per-image energy. In a comparative analysis of deployment energy across ML tasks, Luccioni, Jernite, and Strubell (2024) show that image generation sits at the high end of inference cost, with values reaching around 2.9 kWh per 1,000 inferences. The work adopts this order-of-magnitude coefficient (≈2.9 Wh per image) as a central value for communicative clarity. Electricity is translated into carbon using a global-average CO_2 intensity factor reported by the Global Electricity Review (≈480 g CO_2/kWh). Electricity is translated into cooling water using the two-litre-per-kWh estimate cited above.

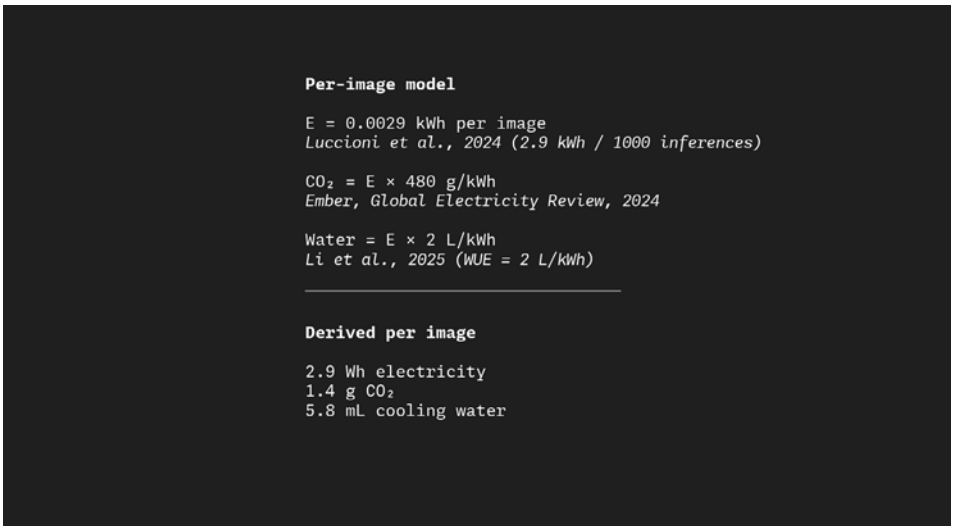

Fig. 3. Per-image energy model and derived environmental metrics (electricity, CO_2, cooling water).[1]

The model is built around two methodological choices. First, the model keeps the three units separate. Electricity, carbon, and water are not collapsed into a single "score," because such aggregation would imply false commensurability. Second, the work is explicit about uncertainty. Per-image energy varies by architecture, accelerator, optimisation, and output size; carbon factors vary by region and time; water intensity varies with cooling technology and siting. *Digital Weight* operates through disclosed assumptions and an auditable structure, allowing viewers to understand how each counter is produced and how scaling works.

Linear Model and Scale

At the methodological level, the model is intentionally linear: multiplying the number of images multiplies energy, emissions, and water. This linearity is conceptually central because it contradicts the everyday intuition that individual digital actions are negligible. The work therefore reports impacts both per unit (per generated image) and through LOW/MID/HIGH scenario bands designed to frame plausible platform-scale activity (e.g., one, ten, and one hundred million images per day). The scenarios function as illustrative frames for scale, making the difference between "one image" and "mass adoption" perceptible.

Within the loop, accumulation unfolds as a micro-simulation. In the current cut, the loop runs for eight minutes at a cadence of two regenerations per second, producing 960 generative iterations. Under the coefficients above, the loop's totals rise into the range of a few kilowatt-hours of electricity, around a kilogram of CO_2, and several litres of cooling water. The precise totals matter less than the perceptual structure: the counters climb steadily even when the image appears almost unchanged. The work insists that "nothing happening" is not the same as "nothing being consumed."

1. From Pengfei Li, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren. 2025. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. Communications of the ACM 68, 7 (July 2025). https://doi.org/10.1145/3724499.

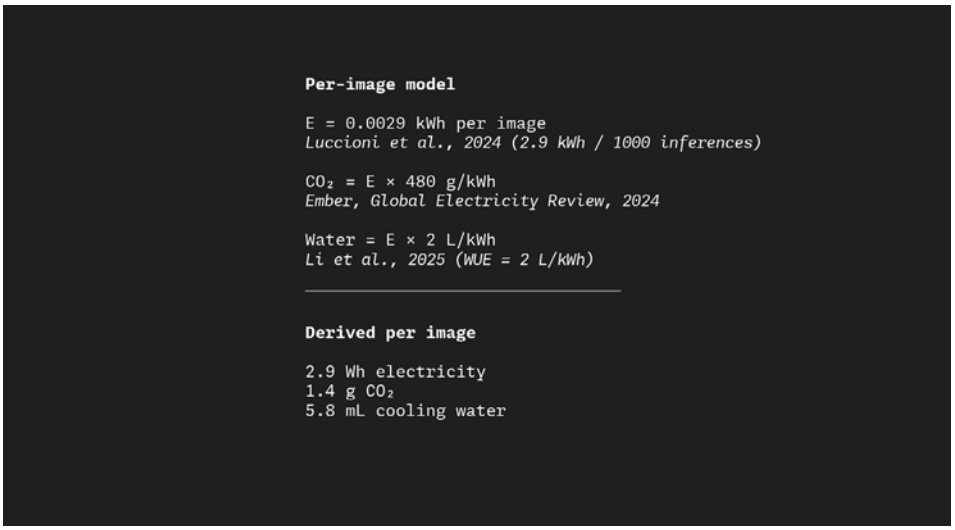

Fig. 4. Linear scaling model with LOW/MID/HIGH scenario bands.

Algorithmic Repetition as Form

Digital Weight visualises environmental data as a temporal aesthetic situation. Computation is experienced as a temporal machine, with accumulation unfolding through repetition, cadence, and near-stasis. The regenerated frame remains almost stable, refusing the usual promise of generative novelty. This near-stasis is deliberate: it stages the paradox of digital culture, where consumption continues even when perception registers minimal visible change.

Repetition is therefore both subject and medium. The work treats regeneration as an algorithmic ritual: perform the same operation, register an increment, and accumulate externalities. This formal strategy mirrors platform dynamics – auto-refresh and infinite scroll – where repetition is normalised and its costs are displaced from view. By bringing cost inside the image, *Digital Weight* collapses the distance between interface time and infrastructure time. The counters materialise what is usually abstracted away.

Fig. 5. Stills from the video loop showing incremental accumulation over repeated generations.

Closed-Loop Dramaturgy and Sound

Its temporal structure follows a closed-loop dramaturgy, without climax or resolution. The same image returns at the beginning, midpoint, and end, so that the loop ends where it begins. The structure foregrounds an energetic cycle that consumes resources without producing mean-

ingful change. In this way, the artwork's temporal organisation becomes its argument.

Sound operates as measurement rather than narration. A low infrastructural hum, a per-iteration tick, and intermittent compute-like bursts mark the cadence and make the work's temporal structure legible through listening. The audience encounters the data as a rhythm of extraction.

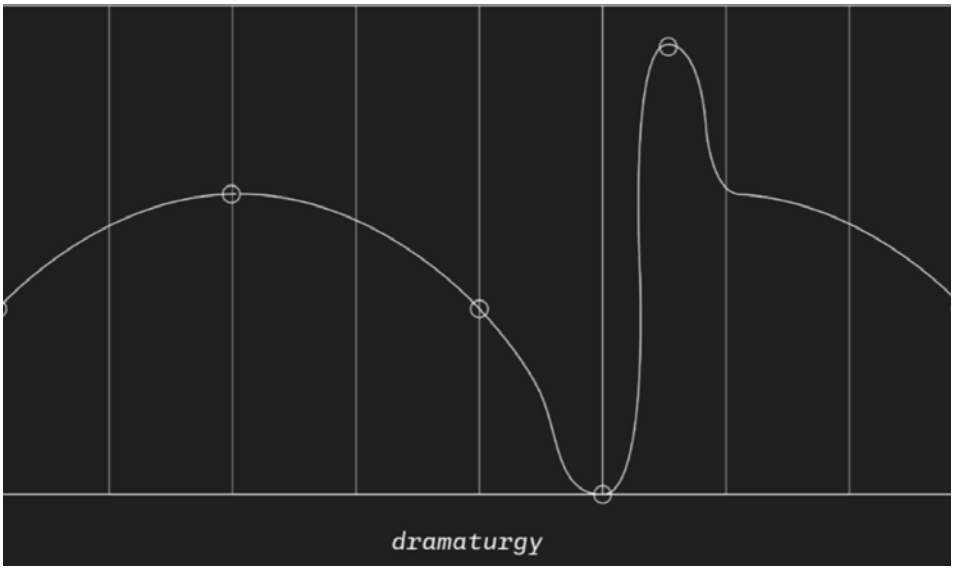

Fig. 6. Closed-Loop Dramaturgy.

Why This Matters Culturally

Generative AI and always-on digital services have normalised a culture of frictionless repetition: endless prompts, infinite feeds, perpetual regeneration. In this condition, environmental costs are structurally absent from perception. *Digital Weight* intervenes at that perceptual level. By binding algorithmic repetition to material counters, it treats the cloud as a concrete extension of planetary metabolism, drawing from the same grids, rivers, and atmosphere as other industrial systems.

The artwork proposes perceptual recalibration: to make the duration of digital activity felt and to make the fiction of immateriality harder to sustain. The "first frame" and the "last frame" of the loop are nearly indistinguishable; what changes is the record of expenditure. This apparent stillness is the point.

Everything digital is material.

Project documentation and preview:
https://kaizere.de/digitalweight_xcoax2026/

References

Barad, Karen.
2007. *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*. Duke University Press.

Ember.
2024. *Global Electricity Review 2024*. May 8, 2024. https://ember-energy.org/latest-insights/global-electricity-review-2024/

IEA.
2025. *Energy and AI*. IEA (International Energy Agency), Paris. April10, 2025 https://www.iea.org/reports/energy-and-ai

Li, Pengfei, Jianyi Yang, Mohammad A. Islam, and Shaolei Ren.
2025. "Making AI Less 'Thirsty': Uncovering and Addressing the Secret Water Footprint of AI Models." *Communications of the ACM* 68, no. 7. https://doi.org/10.1145/3724499

Luccioni, Alexandra Sasha, Yacine Jernite, and Emma Strubell.
2024. "Power Hungry Processing: Watts Driving the Cost of AI Deployment?" In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, 85–99. ACM. https://doi.org/10.1145/3630106.3658542

Zewe, Adam.
2025. "Explained: Generative AI's Environmental Impact." *MIT News*. Massachusetts Institute of Technology. January 17, 2025. https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117

Morton, Timothy.
2013. *Hyperobjects: Philosophy and Ecology after the End of the World*. University of Minnesota Press.