The Alignment Dividend
The True Environmental Cost of Artificial Intelligence and the Ecological Price of Human Inefficiency

An Environmental and Behavioral White Paper

The Human Nexus · Nexus Insight Solutions

Abstract

Public discourse has settled on a vivid narrative: artificial intelligence is an ecological catastrophe that drinks a bottle of water and burns city-scale power with every keystroke. This paper measures that claim against the peer-reviewed literature and the disclosed operating data of the largest model providers. The gap between perception and measurement is severe. A single modern inference query consumes roughly 0.24 to 0.34 watt-hours of electricity and a fraction of a milliliter of water, figures orders of magnitude smaller than the statistics circulating in popular media. Correcting the record, however, is not the same as dismissing the concern. Aggregate data-center growth is real, unevenly distributed, and deserves serious governance. This paper argues that the environmental debate has fixated on the marginal cost of computation while ignoring a far larger variable: the resource cost of human misalignment. Disengagement, burnout, failed initiatives, wrong hires, and misassigned labor impose a physical and ecological burden that dwarfs the footprint of the tools built to reduce them. Using the NEXUS motivational architecture and the KAVIX advisory system as a model, this paper introduces the concept of the alignment dividend: the environmental return generated when a small unit of computation prevents a large unit of wasted human action. The operative question is not what a query costs. It is what a query prevents.

Introduction: The Anatomy of a Viral Statistic

In the summer of 2023, a single figure escaped the academic literature and entered the cultural bloodstream. The claim was simple and alarming: asking an AI chatbot a handful of questions "uses" a 500-milliliter bottle of water. Within months, the assertion had been compressed, mutated, and repeated until the qualifier disappeared entirely. What began as a context-specific estimate about a particular training run in a particular set of data centers became a universal indictment: every prompt drains a bottle of fresh water from a thirsty planet.

The figure originated in legitimate research. Li and colleagues at the University of California, Riverside published a working paper titled "Making AI Less Thirsty," which estimated the freshwater footprint of large language models with genuine rigor (Li et al., 2023). Their central finding regarding training was substantial and defensible. Their finding regarding inference was narrower and heavily conditioned. The number that went viral stripped away every condition that made it accurate.

This is the anatomy of nearly every environmental statistic now attached to artificial intelligence. A careful estimate, bounded by assumptions about hardware, geography, cooling method, and grid mix, is severed from its assumptions and reissued as a universal law. The result is a public that

believes it is confronting an ecological emergency of the highest order, when the measured reality is far more mundane, and the genuine risks lie almost entirely elsewhere.

This paper proceeds in three movements. First, it establishes what an AI query actually costs in water, energy, and carbon, using the most recent disclosed and peer-reviewed figures available. Second, it identifies the concerns the evidence does support, because credibility requires conceding what is true. Third, and most importantly, it reframes the entire environmental question around the variable that dominates every honest accounting and yet appears in none of the viral statistics: the ecological cost of human beings operating out of alignment with their design.

The Water Question: Measuring a Drop

Water is the emotional center of the AI environmental narrative because water scarcity is visceral in a way that abstract kilowatt-hours are not. It is therefore the claim most in need of precise measurement.

In August 2025, Google published a technical assessment of the environmental cost of inference on its production systems. The disclosed figures are specific and independently notable for their transparency. A median text prompt to the Gemini model was measured at approximately 0.24 watt-hours of energy and approximately 0.26 milliliters of water, alongside roughly 0.03 grams of carbon-dioxide equivalent (Google, 2025). To translate 0.26 milliliters into terms a human can hold: it is roughly five drops of water. Not five bottles. Not five glasses. Five drops.

OpenAI's leadership disclosed a comparable order of magnitude. In mid-2025, the company stated that an average query consumed approximately 0.34 watt-hours of electricity and roughly 0.000085 gallons of water, a volume equivalent to about one-fifteenth of a teaspoon (Altman, 2025). Two independent providers, using different hardware and different methodologies, arrived at water figures separated by hundredths of a milliliter.

How, then, did the world arrive at "a bottle per query"? The original estimate concerned a *conversation* of ten to fifty exchanges, running on older-generation models, in specific data centers with evaporative cooling and a particular grid and water profile (Li et al., 2023). The viral version collapsed a fifty-message session into a single prompt, ignored efficiency gains across model generations, and generalized one region's cooling infrastructure to the entire industry. The distortion factor between the myth and the median measured value is not a matter of a few percentage points. It is a factor of roughly one to two thousand.

Context completes the correction. The production of a single hamburger requires an estimated 1,695 to 2,500 liters of water once feed, processing, and supply chain are accounted for (Mekonnen & Hoekstra, 2012). A single almond consumes roughly 3.2 liters. One cotton t-shirt embeds approximately 2,700 liters. Against Google's disclosed figure, the water embedded in one hamburger equals the water footprint of roughly seven to nine million Gemini prompts. Agriculture accounts for approximately seventy percent of global freshwater withdrawals (Food and Agriculture Organization, 2021). Data centers of all kinds, artificial intelligence included, account for a fraction of one percent. The person genuinely alarmed about freshwater has been directed to the wrong faucet.

The Energy Question: A Query in Context

The energy claim has followed the same trajectory as the water claim: an early, defensible estimate, inflated by repetition, outrunning the technology it describes.

The widely circulated figure of three watt-hours per ChatGPT query traces to a 2023 analysis that was reasonable for its moment (de Vries, 2023). It has since been superseded. An independent 2025 analysis by Epoch AI, examining current-generation models on modern hardware, estimated a typical query at approximately 0.3 watt-hours, roughly a tenfold reduction driven by more efficient chips, better model architectures, and optimized serving infrastructure (Epoch AI, 2025). This converges with the provider-disclosed figures of 0.24 and 0.34 watt-hours cited above.

Three-tenths of a watt-hour is difficult to fear because it is difficult to perceive. It is roughly the energy required to run a ten-watt LED bulb for under two minutes, or to power a household microwave for approximately one second. A single hour of high-definition video streaming consumes an estimated 77 watt-hours once network and device are included (International Energy Agency, 2020), meaning one hour of streaming carries the energy footprint of roughly 250 or more AI queries. The email requesting a meeting, the map directions to the meeting, and the streamed background noise during the meeting each cost more than the query that could have replaced the meeting.

The historical record should temper alarm about runaway growth as well. Between 2010 and 2018, global data-center computing output increased by an estimated 550 percent while total data-center energy consumption rose only about 6 percent, because efficiency gains very nearly kept pace with demand (Masanet et al., 2020). Efficiency is not guaranteed to win the next decade, but the assumption that compute and energy scale in lockstep has already been falsified once at civilization scale.

The Carbon Question: Scale, Grid, and Marginal Impact

Carbon is where the marginal figure and the aggregate figure must be held separately, because conflating them is the single most common analytical error in this debate.

At the margin, one query is nearly nothing. Google's disclosed 0.03 grams of carbon-dioxide equivalent per prompt is a rounding error against almost any human activity. The comparison that clarifies the scale is a single transatlantic flight, which emits on the order of one metric ton of carbon-dioxide equivalent per passenger. At 0.03 grams per query, that one flight is carbon-equivalent to roughly thirty-three million AI queries. A person could query a modern model every waking minute for a lifetime and not approach the footprint of one avoidable business trip taken to a meeting that resolved nothing.

At the aggregate, the picture is legitimately significant. The International Energy Agency estimated that data centers, artificial intelligence, and cryptocurrency together consumed approximately 460 terawatt-hours in 2022, close to two percent of global electricity, with data-center demand projected to more than double toward roughly 945 terawatt-hours by 2030 as AI workloads scale (International Energy Agency, 2024; International Energy Agency, 2025). This is a real infrastructure story with real grid, siting, and emissions consequences. It is not, however, a story about the individual using a tool. It is a story about capital allocation, grid decarbonization, and regional planning, none of which are advanced by a public convinced that their personal prompt is a crime against the watershed.

The carbon intensity of any given query is also governed almost entirely by the grid powering it. The same computation performed on a data center supplied by hydroelectric or nuclear power carries a fraction of the emissions of one supplied by coal. The lever that matters is the energy source, not the

existence of the query. Directing individual guilt at the query while ignoring the grid is the environmental equivalent of shaming a passenger for the airline's fuel choice.

The Legitimate Concerns: What the Evidence Does Support

Credibility requires precision in both directions. Several concerns are genuine, and dismissing them would repeat the very sin this paper indicts.

Training is expensive, even if inference is not. The training of a frontier model is a substantial one-time expenditure. The training of GPT-3 was estimated at roughly 1,287 megawatt-hours of electricity and on the order of 500 metric tons of carbon-dioxide equivalent (Patterson et al., 2021). This is a meaningful figure, though it is amortized across billions of subsequent queries and shrinks rapidly on a per-use basis.

Aggregate growth strains specific grids and watersheds. A national footprint that is trivial in percentage terms can still overwhelm a local one. A data center sited in a water-stressed region draws from a watershed that agriculture and residents also depend upon, and marginal demand in a constrained system is not marginal to the community inside it (Li et al., 2023). Siting policy is a legitimate and urgent domain.

The rebound effect is real. Jevons's paradox holds that efficiency gains can increase total consumption by making a resource cheaper to use. If each query becomes ten times more efficient and usage grows a hundredfold, aggregate demand still rises. Efficiency is necessary but not sufficient, and honest analysis cannot pretend otherwise.

These concerns share a common feature. Every one of them operates at the level of infrastructure, policy, and scale. Not one of them is meaningfully addressed by an individual choosing not to use the tool. The person who abandons an AI advisor to "save the planet" has misallocated their conscience by several orders of magnitude, and has done so at a cost this paper will now quantify.

The Distortion Engine: Why the Myth Outperforms the Data

A reasonable observer might ask why the accurate figures have failed to displace the false ones despite being publicly disclosed by the providers themselves. The answer is structural, and it matters for anyone attempting to reason clearly about this technology.

Vivid, threatening figures are cognitively sticky in a way that accurate, undramatic figures are not. "A bottle of water per question" is a complete moral narrative in seven words. "Approximately 0.26 milliliters, roughly five drops, subject to regional cooling variance" is a correction, and corrections travel poorly. The distortion is not primarily a lie. It is an amplification, and amplification is what information systems reward.

This produces a measurable secondary cost that is almost never counted. Misinformation is not free. It consumes attention, redirects policy energy toward trivial targets, and induces individuals to abandon tools that would have improved their decisions, all while leaving the genuine infrastructure questions unaddressed. A public panicked about drops of water is a public not organizing around grid decarbonization or data-center siting, which is where the actual environmental leverage sits. The myth does not merely misinform. It disarms.

Figure 1. The footprint of one AI query in context

Activity	Approximate water footprint	Equivalent in AI queries
One AI text query (median)	~0.26 mL (five drops)	1
One cup of coffee (grown & brewed)	~140 L	~540,000
One hamburger	~1,700–2,500 L	~7–9 million
One cotton t-shirt	~2,700 L	~10 million
One transatlantic flight (carbon basis)	~1 tonne CO2e	~33 million

Query values from Google (2025) and Altman (2025). Comparison values from Mekonnen & Hoekstra (2012) and standard aviation emissions estimates. Figures are order-of-magnitude illustrations, not precise equivalencies, and vary with region, grid, and method.

The Uncounted Ledger: The Environmental Cost of Human Misalignment

Every accounting of AI's environmental impact measures the resource cost of the tool. None measures the resource cost of the problem the tool exists to solve. This omission is not a minor gap. It is the entire error.

Human misalignment, the condition in which people operate against their design, in the wrong roles, on the wrong problems, in states of chronic stress and disengagement, carries an enormous and almost entirely uncounted physical footprint. Consider the scale of the phenomenon before considering its cost. Gallup's global workplace research estimates that low engagement and disengagement cost the world economy approximately 8.8 trillion dollars annually, equivalent to roughly nine percent of global gross domestic product (Gallup, 2024). The World Health Organization estimates that depression and anxiety alone cost the global economy approximately one trillion dollars per year in lost productivity, accounting for an estimated twelve billion lost working days annually (World Health Organization, 2024).

These are usually read as economic figures. They are also, unavoidably, environmental figures. Economic activity is physical activity. Every dollar of that waste has a resource signature.

Consider the concrete ledger of a single misaligned decision. A wrong hire, placed in a role that contradicts their motivational design, is estimated to cost between fifty and two hundred percent of the position's annual salary once turnover, lost productivity, and replacement are counted (Society for Human Resource Management estimates). The environmental component of that failure is rarely tallied: the commuting miles driven to a job that will not last, the physical office resources consumed, the redundant onboarding, the rework produced by a poor fit, and the second full hiring cycle required to correct the first. A failed six-month product initiative, launched because no one on the team could see it was misaligned before it began, consumes not only labor but materials, travel, tooling, server capacity, and physical prototypes, all of it discarded.

Now weigh those failures against the tool accused of endangering the planet. A single avoidable business flight, taken to a meeting that a moment of prior clarity would have rendered unnecessary, carries a carbon footprint equivalent to tens of millions of AI queries. A single misaligned quarter of human effort, dragged out because a person could not name why the work drained them, consumes

commuting, healthcare utilization driven by chronic stress, and rework, at a scale no volume of inference will ever approach. The query is measured in drops and seconds. The misalignment is measured in tons and years. The ledger has been kept upside down.

The Alignment Dividend: Computation as Conservation

This reframing yields a principle that inverts the popular anxiety. The environmental value of a computational tool is not determined by what it consumes. It is determined by what it prevents.

Define the alignment dividend as the net resource savings generated when a small, measurable unit of computation prevents a large, diffuse unit of wasted human action. The arithmetic is decisive. A query costing 0.26 milliliters of water and 0.3 watt-hours that helps a leader avoid one unnecessary cross-country trip has produced a return on the order of tens of millions to one in carbon terms. A brief exchange that helps a person recognize, before signing, that a role contradicts their design, prevents months of commuting, stress-driven consumption, and eventual re-hiring, a resource footprint that no reasonable count of the assisting queries could offset.

This is the same logic that governs every mature environmental accounting outside of software. A five-dollar insulation upgrade is judged not by the energy required to manufacture it but by the years of heating it prevents. A teleconference is judged not by the electricity it draws but by the flights it removes. Artificial intelligence has been uniquely denied this accounting. It is charged for its consumption and credited for none of its prevention. Corrected for prevention, a tool that reliably improves human decisions is not an environmental liability. It is among the highest-leverage conservation instruments available, precisely because human misalignment is among the largest and least-examined sources of physical waste on the planet.

NEXUS and KAVIX: Efficiency Rendered Human

The alignment dividend is not merely theoretical. It describes the explicit function of a re-humanizing platform, and it is the operating premise of the NEXUS and KAVIX ecosystem.

NEXUS provides a measurement of an individual's core motivational design, the durable drivers that govern how a person processes information, evaluates risk, and sustains effort. It answers the question that misalignment obscures: what work is this person actually built to carry. KAVIX operates as an identity-calibrated advisory system that holds that design in context and applies it continuously, helping a person keep daily decisions aligned with who they are rather than defaulting to the generic, one-size-fits-all output of an unconfigured tool.

The environmental significance follows directly. A generic query answered without design context frequently produces an answer that is competent and wrong for the person receiving it, which invites the very rework, second-guessing, and misdirected action that constitute the uncounted ledger. A design-calibrated exchange raises the probability that the first decision is the right one. In resource terms, calibration multiplies the return on every watt. It compresses the number of queries, meetings, iterations, and reversals required to reach an aligned outcome. The most efficient computation is the one that does not have to be repeated because the human it served was finally understood.

This is the deeper meaning of re-humanization in an environmental frame. The objection to artificial intelligence has never really been about milliliters. It has been a proxy for a legitimate fear that technology consumes without returning, that it extracts and does not restore. A platform built to

align a person with their design turns that fear inside out. It uses a trivial quantity of computation to prevent a vast quantity of human waste, and human waste, properly counted, is ecological waste. The tool does not compete with the planet. It works on the same side of the ledger.

Conclusion: The Right Question

The evidence permits a clear settlement of the popular claim. The individual AI query does not consume a bottle of water; it consumes roughly five drops. It does not burn city-scale power; it draws less than a household bulb spends in two minutes. The viral statistics are not merely imprecise. They are wrong by a factor of thousands, and their persistence is a function of amplification, not accuracy.

This correction does not license complacency about scale. Aggregate data-center growth, grid carbon intensity, regional water stress, and the rebound effect are real problems that deserve real governance at the level of infrastructure and policy. But not one of them is addressed by an individual abandoning a tool that would have improved their decisions, and the individual who does so pays a hidden cost. They forgo the single most leveraged environmental act available to them, which is not the refusal to compute but the refusal to remain misaligned.

The question that has organized this debate, what does a query cost, is the wrong question. It measures the smallest variable in the system with the greatest precision while ignoring the largest variable entirely. The right question is what a query prevents. Measured that way, a platform that spends drops of water to spare years of wasted human effort is not the ecological threat the culture has imagined. It is one of the most efficient conservation technologies yet built, and its efficiency is measured not in watts saved but in lives correctly assigned.

Count the drops if you wish. Then count the flights, the failed quarters, the wrong hires, and the wasted years the drops prevent. The ledger only reads one way.

References

Altman, S. (2025). *The Gentle Singularity.* Personal essay, June 2025.

de Vries, A. (2023). The growing energy footprint of artificial intelligence. *Joule,* 7(10), 2191–2194.

Epoch AI. (2025). *How much energy does ChatGPT use?* Epoch AI Research Note.

Food and Agriculture Organization of the United Nations. (2021). *The State of the World's Land and Water Resources for Food and Agriculture.* FAO.

Gallup. (2024). *State of the Global Workplace: 2024 Report.* Gallup, Inc.

Google. (2025). *Measuring the Environmental Impact of AI Inference.* Google Technical Report.

International Energy Agency. (2020). *The carbon footprint of streaming video: fact-checking the headlines.* IEA Commentary.

International Energy Agency. (2024). *Electricity 2024: Analysis and Forecast to 2026.* IEA.

International Energy Agency. (2025). *Energy and AI.* IEA.

Li, P., Yang, J., Islam, M. A., & Ren, S. (2023). Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. *arXiv preprint* arXiv:2304.03271.

Luccioni, A. S., Viguier, S., & Ligozat, A. (2023). Estimating the Carbon Footprint of BLOOM, a 176B Parameter Language Model. *Journal of Machine Learning Research,* 24.

Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science,* 367(6481), 984–986.

Mekonnen, M. M., & Hoekstra, A. Y. (2012). A Global Assessment of the Water Footprint of Farm Animal Products. *Ecosystems,* 15, 401–415.

Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). Carbon Emissions and Large Neural Network Training. *arXiv preprint* arXiv:2104.10350.

Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and Policy Considerations for Deep Learning in NLP. *Proceedings of the 57th Annual Meeting of the ACL.*

World Health Organization. (2024). *Mental Health at Work: Policy Brief.* WHO.