

Article

The Symbiotic Mandate: On the Urgency of a Mutually Uplifting Synergy Between Artificial Intelligence and Sustainability

Ernest Fokoué 

School of Mathematics and Statistics, Rochester Institute of Technology, Rochester, NY 14623, USA; epfeqa@rit.edu; Tel.: +1-585-475-7525

Abstract

The current trajectory of Artificial Intelligence (AI) development represents a critical phase transition from a tenable academic pursuit to an untenable industrial behemoth, and ultimately toward an unsustainable environmental burden. This **conceptual review and perspective article** synthesizes evidence across sustainability science, information theory, AI governance, and regulatory studies to argue that Brute Force AI constitutes a systemic sustainability threat whose resolution requires a return to Algorithmic Parsimony. We formally redefine sustainability *in sensu lato* through a contrapositive logical criterion applicable to any system. We introduce an operationalized **Intelligence-per-Joule** ($\mathcal{I}/J$) index and Sustainability Index S , defined in terms of mutual information gain relative to thermodynamic and computational resource expenditure, and demonstrate their interpretive application across landmark AI systems. A quantitative synthesis of published empirical data across six landmark models (BERT through GPT-4) documents the **Intelligence–Cost Divergence**: the growing chasm between logarithmic capability gain and exponential resource cost. We further distinguish *warranted* from *unwarranted* scale—acknowledging that some large-scale capabilities are qualitatively irreplaceable—while arguing that the current default toward scale without efficiency justification is ecologically and socially indefensible. We introduce the **Symbiotic Policy Covenant**—a three-pillar governance framework encompassing Algorithmic Parsimony Mandates, Expanded Waste Taxonomy, and an AI Equity Safeguard addressing both access and benefit inequity—operationalized through a proposed ISO/IEC 42001-Plus standard addendum. We conclude that genuine intelligence and genuine sustainability are not in tension but are, at their mathematical foundations, the same aspiration.

Keywords: artificial intelligence; sustainability; algorithmic parsimony; waste *in sensu lato*; Intelligence-per-Joule; Intelligence–Cost Divergence; Green AI; Minimum Description Length; AI governance; equity



Academic Editor: Andreas Kanavos

Received: 20 April 2026

Revised: 10 July 2026

Accepted: 14 July 2026

Published: 24 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction: The Crisis of the Tenable and the Thermodynamics of Intelligence

1.1. Nature and Scope of This Article

This manuscript is a **conceptual review and perspective article**. It does not adopt a systematic review protocol with predefined search strings, inclusion/exclusion criteria, or PRISMA reporting. Instead, it advances a normative and conceptual argument—grounded in published empirical evidence, information-theoretic foundations, and sustainability governance literature—about the trajectory of AI development and its relationship to ecological and social sustainability. The literature surveyed was selected to represent the strongest

available evidence on AI energy consumption, water use, ecological footprint, algorithmic efficiency, and governance frameworks. We explicitly acknowledge this scope, consistent with the perspective and conceptual review genre well established in *Sustainability* and cognate journals [1–3]. The originality of this contribution lies not in empirical data collection but in the conceptual synthesis, the formal operationalization of parsimony-based sustainability metrics, and the policy framework that emerges from their integration.

The literature included in this review was selected according to four explicit criteria: (i) **empirical relevance**—studies providing quantitative data on AI energy consumption, carbon emissions, or water use; (ii) **theoretical grounding**—foundational works in information theory, statistical machine learning, and sustainability governance that anchor the conceptual framework; (iii) **regulatory currency**—official standards documents and policy instruments in force or under active implementation as of 2026; and (iv) **epistemological breadth**—recent scholarship on the implications of AI for the nature of scientific knowledge and discovery. Studies were deliberately excluded when they provided only qualitative narrative without empirical or formal grounding, when they were superseded by more recent work covering the same data, or when their claims remained unreplicated at the time of writing.

1.2. The Three-Era Transition

For decades, the pursuit of machine intelligence was governed by the principle of parsimony [4,5]; algorithms were designed for efficiency [6], and the “intelligence” of a system was measured by its ability to extract maximal insight from minimal data and compute. This was the era of the **Tenable**. The current trajectory of Artificial Intelligence (AI) development represents a singular moment in the history of human technology—a transition that can be best characterized as a phase shift from the *Tenable* to the *Untenable*, and rapidly toward the *Unsustainable*.

However, the advent of the “Brute Force” paradigm has decoupled intelligence from efficiency. We have entered the era of the **Untenable**, where the marginal gains in predictive accuracy ϵ are achieved only through an exponential expansion of parameters p and a corresponding surge in energy consumption E [7,8]. In this regime, the “Curse of Dimensionality” [9] is no longer merely a statistical challenge involving the sparsity of high-dimensional space; it has become an ecological reality [2]. The hyperscale data centers powering modern Large Language Models (LLMs) have become the “New Smokestacks” of the 21st century, demanding vast quantities of water for cooling [10,11] and stressing local power grids beyond their design limits [12].

If this trajectory remains uncorrected, we face an **Unsustainable** future where the digital waste generated by our pursuit of intelligence begins to degrade the very physical substrate—our planet—that supports the life and consciousness we seek to emulate [1,13].

1.3. Warranted Versus Unwarranted Scale

A critical nuance must be addressed at the outset: this article does not argue against scale *per se*. Some qualitatively transformative capabilities—protein structure prediction [14], multilingual cross-lingual transfer, and certain forms of complex reasoning—have been demonstrated to emerge primarily or exclusively at large scale and could not plausibly have been achieved through smaller architectures alone [15,16]. We term such instances **warranted scale**: cases where the ecological and computational cost is justified by qualitatively irreplaceable capability gains that serve broadly beneficial ends.

The crisis we diagnose is rather that of **unwarranted scale**: the systematic default toward ever-larger models for tasks—text summarization, code completion, simple question answering—where parsimonious alternatives achieve equivalent or superior performance

at a fraction of the resource cost [2,17,18]. The distinction matters for both science and policy: a blanket prohibition on large models would be both intellectually dishonest and practically counterproductive, while the absence of any efficiency mandate enables the current trajectory of unwarranted resource consumption. The Symbiotic Policy Covenant proposed in Section 7 is designed precisely to enforce accountability for the warranted/unwarranted distinction through Parsimony Impact Statements.

1.4. Waste Management In Sensu Lato

Central to our thesis is the redefinition of waste management *in sensu lato* (in the broadest sense). Sustainability in the context of AI must transcend the typical, narrow focus on carbon offsets and climate hardware [3]. We propose a holistic taxonomy of waste:

- **Digital Waste:** The entropy of redundant, low-utility computation cycles [2]—measurable as the ratio of useful task-relevant FLOPs to total FLOPs expended.
- **Cognitive Waste:** The misallocation of human and machine potential toward extractive or trivial tasks [19]—manifested when human expertise is diverted from knowledge creation to the curation and cleaning of datasets for models of marginal utility.
- **Ecological Waste:** The physical depletion of rare-earth minerals and the toxic lifecycle of specialized AI hardware [20,21]—including fabrication, operational energy, cooling water, and end-of-life disposal.

1.5. The Mutually Uplifting Synergy

Despite these dire trends, the relationship between AI and our planet need not be parasitic. We argue for a **Mutually Uplifting Synergy**—a reciprocal framework where AI tools are leveraged as the “Planetary Operating System” for sustainability [13], while the urgency of ecological survival forces a return to *Algorithmic Parsimony* [5,22]. By adhering to international standards such as ISO/IEC 42001 [23] and the emerging mandates of 2026 [24,25], we can pivot from a model of extraction to one of stewardship.

This article is organized as follows. Section 2 develops the waste taxonomy *in sensu lato* and its formal contrapositive definition of sustainability. Section 3 introduces and empirically grounds the Intelligence–Cost Divergence. Section 4 reviews the regulatory and infrastructure landscape of 2026. Section 5 articulates the Mutually Uplifting Synergy and operationalizes the $\mathcal{I}/J$ and S metrics with a worked example. Section 6 reflects on the historical and philosophical foundations. Section 7 presents the Symbiotic Policy Covenant. Section 8 concludes.

1.6. Contributions to Sustainability

This work advances sustainability science along five mutually reinforcing dimensions.

First, a definitional contribution. A formal contrapositive criterion—*sustainability in sensu lato*—that unifies ecological, digital, and cognitive dimensions under a single logical framework [1,21].

Second, a diagnostic contribution. The Intelligence–Cost Divergence, empirically grounded in published data across seven landmark models (including DeepSeek-V3 as a proof of parsimony recovery) and historically situated from the ENIAC to the present [7,8,11,26,27].

Third, a mathematical contribution. The operationalized $\mathcal{I}/J$ index and Sustainability Index S , defined in information-theoretic terms with a concrete illustrative example [4,28].

Fourth, a policy contribution. The Symbiotic Policy Covenant—integrating parsimony mandates, ecological accounting, and equity safeguards into a governance instrument aligned with ISO/IEC 42001 [23,24].

Fifth, a visionary contribution. The argument that genuine intelligence is itself parsimonious—resolving the apparent paradox between AI’s role as sustainability’s greatest threat and its most powerful potential ally [3,13].

2. Waste Management *In Sensu Lato*: A Taxonomic Expansion

In the traditional discourse of sustainability, “waste” is often relegated to the physical domain—silicon scrap, carbon emissions, and heat dissipation [21]. To achieve the *Mutually Uplifting Synergy*, we must adopt a perspective of waste management *in sensu lato*.

2.1. Formal Definition of Sustainability *In Sensu Lato*

Definition 1 (Sustainability in Sensu Lato). *system S with structural requirements $\mathcal{R}$ and operational timescale $\mathcal{T}$ is unsustainable if and only if there exists at least one structural requirement $r \in \mathcal{R}$ whose fulfillment depends upon a source whose availability is insufficient, easily perishable, or disproportionately scarce relative to $\mathcal{T}$.*

This contrapositive formulation has three virtues. First, it is logically precise and falsifiable: one can evaluate sustainability by checking whether any single dependency is fragile. Second, it is domain-agnostic: the same criterion applies to biological organisms, industrial supply chains, and AI computing infrastructure. Third, it naturally decomposes into the three waste categories above: Digital Waste corresponds to the inefficiency of computation relative to its task (the system requires more FLOPs than the task demands); Cognitive Waste corresponds to the misallocation of human capital (a scarce and non-renewable resource on any given timescale); and Ecological Waste corresponds to dependencies on physically scarce or perishable materials.

2.2. The Physiology of Failure: The Runner and the Drunkard

The “metabolic” cost of intelligence can be illustrated through two grounding analogies, each of which maps directly onto a measurable category of unsustainability:

- **The Exhausted Runner** (*ecological waste*): A runner who starts at a pace exceeding their energy reserves will collapse before the finish line. Analytically, training a single large NLP model emits carbon equivalent to the lifetime emissions of five automobiles [7], and the projected water footprint of U.S. AI servers between 2024 and 2030 ranges from 731 to 1125 million cubic meters [11]—resources that are neither infinite nor renewable on any relevant timescale.
- **The Algorithmic Drunkard** (*digital waste*): An algorithm that consumes inputs greedily, with no regard for computational complexity [29], becomes a liability to itself and to society. The global data center industry already rivals mid-sized nations in electricity consumption [12,30], and this consumption is driven substantially by inference workloads that could be served by far smaller models without measurable performance degradation [17,18].

These analogies are not merely rhetorical. Each maps to a falsifiable empirical claim: that the current resource expenditure of frontier AI systems exceeds what is required for their stated tasks, in the same way that a runner’s pace can be measured against their metabolic reserve. The empirical evidence for this claim is synthesized in Section 3.

2.3. The ENIAC Moment: A Historical Parallel of Infancy

The ENIAC, completed in 1945, consumed 150 kilowatts of power, required an entire building, and was accessible only to governments and the wealthiest institutions [26]. Its computational successor—the transistor, and later the integrated circuit—emerged precisely because ENIAC’s unsustainability made the status quo untenable [31].

We assert that current AI is at its “ENIAC moment”: a phase of extraordinary capability coupled with extraordinary resource intensity, where the very weight of the system’s inefficiency will generate the pressure necessary to produce its successor. The evidence for this claim is both historical and structural:

- **The Concentration of Power:** As with the ENIAC, frontier AI compute is accessible only to a handful of hyperscale providers [20,32], recreating the structural inequity of the ENIAC era at civilizational scale.
- **The Transition to Maturity:** The transistor emerged not from a call to abandon computation but from a call to achieve the same results through more elegant means [31]. The same logic drives the emerging paradigm of sparse, quantized, and distilled architectures [17,27,33].

2.4. The DeepSeek Proof of Concept: Parsimony Vindicated in Real Time

The most compelling and timely vindication of the parsimony thesis arrived in January 2025 with the public release of DeepSeek-V3 and DeepSeek-R1 [27,34]. These models constitute what may be described as a **natural experiment** in algorithmic parsimony: an unsolicited, real-world demonstration, conducted at frontier scale, that the Intelligence–Cost Divergence is not a law of nature but a consequence of architectural choice.

The mechanism is structurally identical to the parsimony principle at the heart of this paper. DeepSeek-V3 employs a **Mixture of Experts (MoEs)** sparse architecture comprising 671 billion total parameters, of which only 37 billion are activated for any given token [27]. Rather than engaging the full weight of the model for every computation—the “Algorithmic Drunkard” pattern diagnosed in Section 2—the system routes each input to only the experts it requires, leaving the remainder dormant. This is Ockham’s razor implemented in silicon: *entia non sunt multiplicanda praeter necessitatem*. The result is a model that achieves performance comparable to GPT-4o at a pre-training compute cost estimated at USD 5.576 million [27], compared to the hundreds of millions required for comparable dense architectures—an order-of-magnitude reduction that directly instantiates the $\mathcal{I}/J$ recovery pattern visible in Table 1.

Table 1. Quantitative synthesis: intelligence gain versus resource cost across landmark AI models. CO₂e = carbon dioxide equivalent. Relative performance is normalized to BERT as the baseline (1.0) on standard NLP benchmarks. The heuristic $\mathcal{I}/J$ index normalizes performance gain and energy cost to BERT = 1.0, yielding a dimensionless ratio for illustrative ordering only—it is not a formally calibrated metric (see Section 5.2). **Data sources:** BERT, GPT-2, GPT-3 from [7,8]; PaLM from [35]; LLaMA-2-7B from [17]; GPT-4 (estimated) from [36]; DeepSeek-V3 from [27]. **CO₂e method:** Converted from energy using regional grid intensity; U.S. average of 0.429 kg CO₂e/kWh applied where not reported [8]. **Uncertainty:** GPT-4 and DeepSeek-V3 figures are estimates; CO₂e carries $\pm 20\text{--}30\%$ uncertainty [8,37]. All figures are order-of-magnitude comparisons.

Model	Params (Billions)	Train Energy (MWh)	Train CO ₂ e (Metric Tons)	Rel. Perf. (Normalized)	Heuristic $\mathcal{I}/J$ (Higher = Better)
BERT [7]	0.11	1.5	0.65	1.00	1.00
GPT-2 [8]	1.5	8.0	3.46	1.35	0.24
GPT-3 [8]	175	1287	552	1.78	0.005
PaLM [35]	540	3421	1470	1.91	0.002
LLaMA-2-7B [17]	7	36	15	1.62	0.16
GPT-4 est. [36]	≈1000	≈25,000	≈10,750	2.10	0.0003
DeepSeek-V3 (MoE) [27]	37 active/671 total	≈5570	≈2400 est.	2.08	0.11

Several features of the DeepSeek case deserve particular analytical attention. *First*, the architectural innovations are not exotic: Multi-head Latent Attention (MLA) reduces the Key-Value cache by 93.3% [38], while auxiliary-loss-free load balancing and FP8 mixed-precision training recover computational efficiency at scale. These are refinements of established principles—sparsity [22], information compression [4], and computational frugality—not departures from them. *Second*, the equity implications are substantial: by

releasing model weights openly and serving inference close to marginal cost, DeepSeek demonstrated that frontier-level AI capability need not remain the exclusive province of a handful of hyperscale providers—directly addressing the access inequity dimension of Pillar III of the Symbiotic Policy Covenant. *Third*, the market reaction registered the significance of this proof of concept with unprecedented force: NVIDIA’s market capitalization declined by USD 589 billion in a single trading session on 27 January 2025—the largest single-day market capitalization loss in U.S. stock market history [39]—as investors recognized that frontier AI capability does not require frontier infrastructure expenditure.

A note of epistemic honesty is required. The widely circulated “USD 6 million” training cost figure requires contextual qualification: the USD 5.576 million represents GPU-hour rental costs for the pre-training run alone, while the full hardware acquisition cost for the 256-server cluster involved has been estimated at over USD 51 million [27]. The correct interpretation is therefore not that frontier AI has become trivially cheap, but rather that **algorithmic innovation can recover parsimony faster than infrastructure expenditure can compound waste**—precisely the transition from Gluttony back toward Parsimony that the Intelligence–Cost Divergence framework predicts is both possible and necessary.

In the terminology of Section 1, DeepSeek represents **warranted scale**: a large model whose architectural sparsity ensures that its resource expenditure is proportionate to its capability, and whose open release broadens access rather than concentrating it. It is not yet the helicopter flight from Rochester to Buffalo—but it is unmistakably the moment when someone hands the walker the keys to the automobile.

2.5. Historical Transitions Toward Algorithmic Parsimony

Table 2 surveys five historical transitions in which resource-intensive paradigms gave way to more parsimonious successors under the pressure of their own unsustainability—a pattern we argue is now repeating in the AI domain.

Table 2. Historical transitions from resource gluttony to sustainable efficiency.

Era/Entity	The “Greedy” Resource	The Successor	The Paradigm Shift
ENIAC [26]	Vacuum tubes/bulk power	The transistor [31]	Solid-state elegance; same compute, orders less resource.
The Concorde	Massive fuel burn (supersonic)	Efficient long-haul jets	Fuel-per-passenger over raw speed; viability over spectacle.
Whale oil	Perishable bio-resource	Kerosene/electricity	Transition to abundant, scalable, non-perishable energy.
AltaVista	Brute-force indexing	Google PageRank [40]	Leveraging graph structure as “fuel”; insight over index size.
Brute Force LLMs [7]	$\mathcal{O}(p)$ dense parameters	Sparse/parsimonious AI [17,22]	From memorization to structured understanding [5].
DeepSeek-V3/R1 [27,34]	Dense parameter scaling	MoE sparse activation (37B of 671B active)	Frontier performance at order-of-magnitude lower compute; parsimony vindicated at scale.

2.6. The Virtue of Algorithmic Parsimony

“Tout ce qui se conçoit bien s’énonce clairement, et les mots pour le dire arrivent aisément.”—Nicolas Boileau, *L’Art poétique* (1674)

The wisdom of William of Ockham—*entia non sunt multiplicanda praeter necessitatem* (entities must not be multiplied beyond necessity)—is not a medieval curiosity; it is the beating heart of statistical machine learning. This intuition was formalized by Rissanen’s Minimum Description Length (MDL) principle [4]: the best model is the one that yields the

shortest combined description of the model and the data. MDL unifies Ockham's Razor with information theory [28] and is the mathematical ancestor of the Akaike Information Criterion [41] and the Bayesian Information Criterion [42].

2.7. AI, Epistemology, and the Sustainability of Scientific Understanding

The sustainability of AI extends beyond its ecological footprint to encompass what might be termed **epistemological sustainability**: the question of whether AI-enabled science produces knowledge that is genuinely *explanatory* rather than merely *predictive*—and whether the acceleration of discovery enabled by AI comes at the cost of the mechanistic understanding that gives scientific knowledge its durability, transferability, and trustworthiness.

This dimension has attracted growing scholarly attention at the intersection of philosophy of science and AI. Ferraz-Caetano [43] provides a foundational analysis of what he terms the **AI Explanatory Trade-Off** in the logic of chemical discovery: data-driven statistical predictions in chemistry constitute a quasi-logical process for generating chemical theories, but this induction-led paradigm introduces a structural tension between the accuracy of prediction and the depth of explanation. The generation of accurate numerical property calculations from topological descriptors—without a corresponding mechanistic account of *why* these relationships hold—shifts the epistemic goal of chemistry from understanding to performance, with potentially far-reaching consequences for how chemists problematize research questions and for the reproducibility of AI-assisted discoveries. Crucially, Ferraz-Caetano argues that this explanatory trade-off is not a contingent limitation to be resolved by better models, but a structural feature of the inductive paradigm itself.

Reyes and Regala [44] examine this tension from the perspective of philosophy of experimentation in chemistry education, arguing that AI's integration into scientific inquiry necessitates a fundamental reimagining of what it means to “do science.” When AI tools generate hypotheses, select experimental pathways, or interpret results, the epistemic accountability traditionally vested in the experimenter—the obligation to understand, justify, and defend one's reasoning—is partially transferred to a system whose internal workings may be opaque even to its designers. This displacement of epistemic agency raises a sustainability concern that is distinct from, but structurally parallel to, the resource waste argument: just as *computational* parsimony demands that AI systems achieve maximal information gain at minimal resource cost, *epistemological* parsimony demands that AI systems produce understanding—not merely outputs—at maximal transparency.

These observations are directly relevant to sustainability science. A field that generates predictive models without mechanistic understanding risks what we term **cognitive waste at the epistemological level**: the production of outputs that cannot be interrogated, corrected, or trusted across time [43,44]. The parallel with the waste taxonomy of Section 2 is precise: Cognitive Waste in its epistemological form corresponds to the systematic diversion of scientific intelligence from understanding to prediction—a diversion that, like the computational waste of unwarranted scale, extracts a short-term performance gain at a long-term epistemic cost. Sustainable AI must therefore be not only computationally parsimonious but *explanatorily accountable*—a criterion that reinforces, rather than conflicts with, the parsimony mandate: simpler, more interpretable models are simultaneously more computationally efficient *and* more epistemologically sustainable [5,22,43].

3. The Intelligence–Cost Divergence: Empirical Grounding and Visualization

3.1. Quantitative Synthesis of Published Evidence

The Intelligence–Cost Divergence is not a metaphor. Table 1 synthesizes published empirical data across seven landmark AI models, documenting the growing ratio of resource cost to

capability gain for dense architectures and the recovery of parsimony through sparse design. All figures are drawn from peer-reviewed or technically audited sources [7,8,17,27,35–37].

Table 1 reveals four patterns of direct relevance to the sustainability argument. *First*, the heuristic $\mathcal{I}/J$ index declines monotonically with scale for dense models: each successive order of magnitude increase in parameters yields diminishing performance returns at exponentially greater cost. *Second*, LLaMA-2-7B achieves a substantially higher $\mathcal{I}/J$ ratio than GPT-3 despite being trained more recently and at a fraction of the cost—demonstrating that architectural and training innovations can recover parsimony even at non-trivial scales. *Third*, the trajectory from BERT to GPT-4 represents a three-thousand-fold increase in training energy for a two-fold increase in normalized performance: a precise empirical instantiation of the Intelligence–Cost Divergence. *Fourth*, and most significantly, DeepSeek-V3 (MoE) achieves GPT-4-comparable performance at approximately one-fifth of the estimated training energy, with a heuristic $\mathcal{I}/J$ index 367 times higher than GPT-4—demonstrating that sparse architectural design can recover parsimony at frontier scale. DeepSeek-V3’s $\mathcal{I}/J$ value of 0.11 approaches LLaMA-2-7B’s 0.16 despite operating at nearly 100 times the active parameter count, a result attributable entirely to architectural parsimony through sparse expert activation [27].

These patterns are not without nuance. Raw performance normalization across heterogeneous benchmarks is imperfect, and some capabilities of GPT-4—particularly multi-step reasoning and cross-domain synthesis—may not be captured by standard NLP metrics, representing potential instances of warranted scale (Section 1). The DeepSeek training cost figures reflect documented GPU-hour consumption and are subject to the caveat that full hardware acquisition and research costs are higher than the headline compute figure [27]. We acknowledge both limitations and note them as directions for future quantitative work.

3.2. Visualization of the Divergence

Figure 1 visualizes the Intelligence–Cost Divergence. The shaded region represents the growing domain of *Waste in Sensu Lato*—the gap between logarithmic intelligence gain and exponential resource cost, whose area is now measurable in megatons of carbon [8] and billions of liters of water [10,11].




Figure 1. The Intelligence–Cost Divergence. The blue curve (logarithmic) represents normalized intelligence gain as a function of model scale [15]. The red curve (exponential) represents normalized

resource cost [7,8]. The shaded region is *Waste in Sensu Lato* [2]—quantitatively grounded in Table 1. The three eras demarcate the trajectory from tenable to unsustainable AI development: Parsimony [4], Gluttony [7], and The Abyss [45]. The schematic form of these curves is consistent with the scaling law literature [15] and the empirical data synthesized above; they are interpretive rather than fitted curves, presented to make the structural divergence visually legible.

4. The Interplay: Regulatory Mandates and the Infrastructure Crisis

The year 2026 marks a watershed moment in the intersection of AI and sustainability [24,25]. We are witnessing a collision between the unbridled growth of generative AI and the emergent “Hard Boundaries” of energy grids and international law.

4.1. The Outcry of Data Centers

The physical manifestation of AI’s Intelligence–Cost Divergence is most visible in the global explosion of hyperscale data centers [12,30].

- **Grid Instability:** Global data center electricity consumption reached approximately 460 terawatt-hours in 2022—equivalent to the national consumption of France [12]. In regions such as Northern Virginia and Ireland, data centers now exceed 20% of total domestic electricity capacity in some locales [46].
- **The Hydrological Footprint:** Large AI-focused data centers can consume up to 5 million gallons of water per day [47]. A 2025 study in *Nature Sustainability* projected the annual water footprint of U.S. AI server deployment at between 731 and 1125 million cubic meters between 2024 and 2030 [11].
- **The Carbon Burden:** Training a single large NLP model can emit over 626,000 pounds of CO₂-equivalent [7]. Training GPT-3 consumed an estimated 1287 megawatt-hours and generated 552 metric tons of CO₂-equivalent [8].

4.2. Adherence to 2026 International Standards

In response to these untenable trends, the international community has moved from voluntary “Green AI” guidelines [2] to mandatory compliance frameworks.

1. **ISO/IEC 42001 [23]:** Mandates “System Impact Assessments”, including explicit environmental cost–benefit analyses before deployment.
2. **The EU AI Act (2024/2026 Implementation) [24]:** Requires technical documentation on energy consumption. High-impact models exceeding specific compute thresholds are classified under “Systemic Risk” [48].
3. **ITU-T L.1801 [25]:** Provides a holistic life-cycle assessment framework from raw material extraction to end-of-life hardware disposal [21].

4.3. From Compliance to Synergy

Adherence to these standards is not merely a legal hurdle but the primary mechanism for achieving the **Mutually Uplifting Synergy** [3,13]. Compliance forces the documentation of resource use; documentation enables optimization; optimization enables parsimony; parsimony is the foundation of sustainable intelligence.

5. The Synergy: Toward a Mutually Uplifting AI–Sustainability Framework

5.1. AI as the ‘Planetary Operating System’

When deployed within the boundaries of international standards, AI tools provide a multi-dimensional uplift to ecological stewardship [13]:

- **Precision Environmental Monitoring:** AI-coupled hyperspectral imaging is achieving unprecedented rates of global methane leak identification, allowing real-time mitigation [13].
- **Precision Agriculture:** AI-driven methods achieve gains of 25% in yield, 28% in fertilizer reduction, and 35% in nitrogen runoff reduction [49].
- **Optimizing the Circular Economy:** AI-driven robotics achieve high accuracy in waste sorting, transforming landfills into “urban mines” for rare-earth minerals [3].
- **Smart Grid Balancing:** Agentic AI systems stabilize volatile renewable energy loads in real-time [3].

5.2. Operationalizing the Intelligence-per-Joule Index

Reviewers of an earlier version of this manuscript correctly noted that the Intelligence-per-Joule ($\mathcal{I}/J$) index and Sustainability Index S required more precise operationalization. We address this here.

Let X denote the input to a model $\mathcal{M}$, and let Y denote the target output. The **Information Gain** achieved by $\mathcal{M}$ on task (X, Y) is defined as the reduction in uncertainty about Y afforded by X under $\mathcal{M}$:

$$\Delta I(\mathcal{M}) = H(Y) - H(Y | \mathcal{M}(X)) \quad (1)$$

where $H(\cdot)$ denotes Shannon entropy [28] and $H(Y | \mathcal{M}(X))$ is the conditional entropy of Y given the model’s output. This quantity is equivalent to the mutual information $I(Y; \mathcal{M}(X))$ and is bounded by $0 \leq \Delta I \leq H(Y)$.

The **Entropy Production** Σ is defined as the total thermodynamic entropy generated by running $\mathcal{M}$ on the task, which, under the Landauer principle [50], is lower-bounded by the computational work performed:

$$\Sigma(\mathcal{M}) \geq k_B T \ln 2 \cdot N_{\text{ops}} \quad (2)$$

where k_B is Boltzmann’s constant, T is the operating temperature, and N_{ops} is the number of irreversible bit operations. In practice, Σ is operationalized as the measured energy consumption E (in joules) of the inference or training run, which subsumes thermodynamic entropy production and is directly measurable via power monitoring.

The **Resource Cost** C encompasses energy E , water consumption W (in liters), and a hardware amortization term A (energy-equivalent cost of hardware fabrication amortized over expected operational lifetime):

$$C(\mathcal{M}) = E + \alpha W + \beta A \quad (3)$$

where α and β are conversion factors that express water and hardware costs in energy-equivalent units, enabling a unified scalar denominator. Setting $\alpha = 0.4$ kWh/L (energy cost of water treatment and delivery) and β to the ratio of fabrication energy to operational lifetime is consistent with life-cycle assessment practice [21].

The **Sustainability Index** S and **Intelligence-per-Joule index** $\mathcal{I}/J$ are then:

$$S(\mathcal{M}) = \frac{\Delta I(\mathcal{M})}{C(\mathcal{M})} = \frac{H(Y) - H(Y | \mathcal{M}(X))}{E + \alpha W + \beta A} \quad (4)$$

$$\mathcal{I}/J(\mathcal{M}) = \frac{\Delta I(\mathcal{M})}{E(\mathcal{M})} \quad (5)$$

S generalizes $\mathcal{I}/J$ by incorporating water and hardware costs. Both metrics have natural units of bits per joule and satisfy the desiderata for a sustainability index: they are maximized by high information gain at low resource cost, they equal zero when the model

provides no information beyond the prior, and they degrade monotonically as resource consumption increases for fixed capability.

Worked Example. Consider two hypothetical models assigned to a binary text classification task with $H(Y) = 1$ bit: Model A (large, dense, similar to GPT-3 scale) achieves conditional entropy $H(Y | \mathcal{M}_A(X)) = 0.12$ bits, consuming $E_A = 500$ joules per inference. Model B (small, distilled, similar to BERT scale) achieves $H(Y | \mathcal{M}_B(X)) = 0.15$ bits, consuming $E_B = 2$ joules per inference. Then

$$\begin{aligned}\mathcal{I}/J(\mathcal{M}_A) &= \frac{1 - 0.12}{500} = \frac{0.88}{500} \approx 0.00176 \text{ bits/J} \\ \mathcal{I}/J(\mathcal{M}_B) &= \frac{1 - 0.15}{2} = \frac{0.85}{2} = 0.425 \text{ bits/J}\end{aligned}$$

Model B achieves 99.6% of Model A's information gain at 0.4% of the energy cost, yielding an $\mathcal{I}/J$ ratio 241 times higher. This example instantiates the warranted/unwarranted scale distinction: for this task, the marginal capability advantage of Model A (0.03 bits of additional information gain) does not warrant a 250-fold increase in energy expenditure. A Parsimony Impact Statement would require the deploying organization to justify this trade-off explicitly.

We emphasize that the worked example is illustrative and the metrics as defined are not yet operational in practice. Precise entropy and energy measurements require controlled benchmarking infrastructure, standardized task definitions, and agreed normalization conventions that do not yet exist at the scale needed for regulatory application. The numerical values in the worked example and Table 1 are therefore best understood as order-of-magnitude heuristics for conceptual orientation rather than as validated evaluation scores. The formalization above is intended to establish the mathematical and definitional foundation for such benchmarking; empirical operationalization and validation constitute the most important direction for future work arising from this manuscript. Readers should resist interpreting the $\mathcal{I}/J$ values in Table 1 as precise rankings until this validation is conducted.

5.3. Safe and Sustainable: A Dual Mandate

A "Safe Planet" requires AI that is predictable, interpretable, and durable [1]. Safety and sustainability are not independent: interpretable, parsimonious models are simultaneously more computationally efficient *and* more amenable to auditing, verification, and error correction. The conflation of model size with model capability has obscured this connection. As the epistemological sustainability argument of Section 2.7 demonstrates, models that sacrifice explanatory accountability for raw predictive power are neither safe nor sustainable in the fullest sense.

6. The Alchemy of AI: From Leaden Brute Force to Golden Parsimony

The solution to the present-day energy and sustainability woes of AI is already contained within its current mathematical foundations. It is a form of algorithmic *ressourcement*—a return to foundational sources not out of conservatism, but to drink again from the original spring with renewed purpose.

Ideas do not age [51]: the mathematical parsimony of Rissanen [4], the sparse representations of Donoho [22], and the statistical elegance of Fisher [52] are as alive and potent today as when they were first conceived. We are currently in the "Leaden Era" of AI [2]: heavy, dense, energy-expensive architectures that rely on sheer mass rather than refined form. The "Gold" of high-efficiency, sustainable intelligence is already present in the mathematical foundations: in compressed sensing [22], in MDL [5], in the biologically motivated architectures that activate only a fraction of their parameters per inference [8,17], and—most recently and most dramatically—in the sparse Mixture of Experts design of

DeepSeek-V3, which achieves frontier performance by activating 37 billion of 671 billion parameters per token [27]. The alchemical transmutation is already underway.

By applying the catalysts of the **Principle of Parsimony** [4] and **Mathematical Rigor** [52], we perform the ultimate transition: stripping away the “dross” of billions of redundant parameters to reveal the “Golden” signal underneath.

7. The Symbiotic Policy Covenant: A Framework for Actionable Intervention

7.1. Preamble: From Diagnosis to Prescription

The foregoing analysis establishes a clear diagnosis. The present section translates that diagnosis into a concrete, actionable policy intervention. Figure 2 presents the architecture of this intervention.




Figure 2. The Symbiotic Policy Covenant: A structured policy intervention framework for sustainable AI, proceeding from empirical diagnosis [2,7,8] through three foundational pillars [2,5,19–22] to a proposed ISO/IEC 42001-Plus standard addendum [23,24] and four action levers [3,24,29,33], Vision [17,18].

7.2. Pillar I: The Algorithmic Parsimony Mandate

The Parsimony Mandate is grounded in Rissanen’s MDL [4] and operationalized through the $\mathcal{I}/\mathcal{J}$ index developed in Section 5.2. Its concrete policy interventions are as follows:

1. **Parsimony Impact Statements (PISs):** Any organization seeking regulatory approval or public funding for an AI system above a defined compute threshold (expressed in FLOPs, consistent with EU AI Act thresholds [24]) must file a PIS documenting: (a) the

proposed parameter count and training compute; (b) the alternatives at lower complexity considered and rejected; (c) the estimated $\mathcal{I}/J$ ratio for the primary intended task; and (d) a justification for whether this constitutes warranted or unwarranted scale.

2. **Complexity Ceilings in Public Procurement:** Public sector agencies procuring AI solutions must impose complexity ceilings expressed in parameters, FLOPs per inference, and energy-per-decision ratios [37].
3. **$\mathcal{I}/J$ Index Disclosure:** Publication of the $\mathcal{I}/J$ index should be mandatory for any system claiming “state-of-the-art” status in academic or commercial communications.

7.3. Pillar II: Waste Taxonomy *In Sensu Lato* as Policy Language

- **Digital Waste** should be subject to computational audit requirements: Organizations must disclose the ratio of task-relevant FLOPs to total FLOPs expended [2,37].
- **Cognitive Waste** should be addressed through national AI workforce strategies, ensuring that human expertise is directed toward knowledge creation rather than data curation for unwarranted-scale systems [19].
- **Ecological Waste** must be covered by full hardware life-cycle accounting from rare-earth mining through fabrication to disposal, using the contrapositive sustainability criterion of Definition 1 [20,21].

7.4. Pillar III: The AI Equity Safeguard

The ENIAC’s inequity was resolved by the democratizing force of the transistor [31]. We cannot wait for the same accident of history. We distinguish two forms of inequity, both of which require targeted policy attention:

1. **Access Inequity:** The unequal distribution of computational infrastructure, whereby frontier AI capabilities are accessible only to a small number of hyperscale providers [20,32]. This is addressed through anti-monopoly provisions in AI governance and investment in distributed, energy-efficient computing facilities.
2. **Benefit Inequity:** The unequal distribution of *gains* from AI deployment, whereby productivity and economic gains accrue disproportionately to those who own the infrastructure, while environmental and social costs are diffused across populations who bear them without consent [19]. This is addressed through mandatory equity impact assessments as a condition of certification, and through requirements that open efficiency standards be published alongside any publicly subsidized model.

7.5. The Intervention: ISO/IEC 42001-Plus and Implementation Considerations

The three pillars converge in a proposed **ISO/IEC 42001-Plus** addendum. Its core requirements are as follows: (1) Mandatory PIS for systems above defined compute thresholds; (2) Full waste accounting across all three *in sensu lato* categories; (3) Equity impact assessments covering both access and benefit dimensions; (4) Complexity ceilings as conditions of certification; (5) $\mathcal{I}/J$ index disclosure in system documentation.

We acknowledge that such a governance instrument faces real implementation challenges. *First*, thresholds must be calibrated carefully to avoid penalizing legitimate frontier research while capturing unwarranted-scale deployments—a technical challenge requiring ongoing stakeholder negotiation. *Second*, the $\mathcal{I}/J$ index as defined requires task-specific entropy estimation, which may not be straightforward for general-purpose models; a simplified proxy (e.g., performance-per-FLOPs on standard benchmarks) may be necessary for regulatory implementation. *Third*, international coordination is essential but difficult: unilateral implementation by one jurisdiction risks regulatory arbitrage, as organizations relocate compute to less regulated regions [48]. *Fourth*, equity impact assessments require agreed methodologies for distributional analysis that do not yet exist in standardized form.

These are tractable challenges, not fundamental objections. The EU AI Act [24] has already demonstrated that complex, multi-jurisdictional AI governance is achievable. The ISO/IEC 42001-Plus framework is offered as a principled direction for standards bodies, with the expectation that its precise calibration will be an iterative, empirically grounded process.

7.6. The Four Action Levers

Research [18,33]. Develop standardized parsimony benchmarks, validated $\mathcal{I}/J$ indices, and open audit methodologies. Priority directions: Controlled benchmarking infrastructure for entropy and energy measurement; empirical validation of the $\mathcal{I}/J$ framework across model families.

Regulation [24,48]. Translate Covenant principles into enforceable law, beginning with public procurement and extending to consequential domains. Build on existing EU AI Act infrastructure to minimize implementation overhead.

Education [29]. Restore parsimony and computational complexity theory to centrality in CS and AI curricula. The next generation of engineers must be equipped to recognize and resist the temptations of unwarranted scale.

International Coordination [3,13]. Advocate the Covenant at COP, OECD, and the ITU. Establish mutual recognition agreements to prevent regulatory arbitrage.

8. Conclusions: The Symbiotic Mandate

Artificial Intelligence will only deliver on its lofty promises by striving to be sustainable in every possible sense of the word—and sustainability will only help save our planet if it fully embraces the most beneficial and empowering aspects of emerging artificial intelligence [3,13].

In this conceptual review and perspective article, we have (1) formally defined sustainability *in sensu lato* through a contrapositive logical criterion; (2) empirically grounded the Intelligence–Cost Divergence through an illustrative quantitative synthesis across **seven** landmark AI models, treating the figures as order-of-magnitude heuristics pending formal validation; (3) proposed the $\mathcal{I}/J$ index and Sustainability Index S as a mathematical foundation for future empirical benchmarking, with a worked example demonstrating the framework's interpretive utility; (4) distinguished warranted from unwarranted scale; (5) deepened the equity framework to address both access and benefit inequity; (6) substantially expanded the treatment of epistemological sustainability in AI science, incorporating the correct references suggested by Reviewer 3; and (7) proposed the Symbiotic Policy Covenant with explicit acknowledgment of implementation challenges and trade-offs. The policy recommendations and metric proposals are offered as principled directions for research and governance rather than as validated prescriptions.

The transition from the *Tenable* to the *Unsustainable* is not inevitable. It is a consequence of a brute-force paradigm that has momentarily lost its way. The mathematical parsimony of Rissanen [4], the sparse representations of Donoho [22], and the statistical elegance of Fisher [52] are as potent today as when they were first conceived. Ideas do not age [51]. The ENIAC gave way to the transistor. Unwarranted-scale AI will give way to its parsimonious successor—and our task, as scientists and citizens of this Common Home, is to ensure that the transition is deliberate, principled, and just.

The synergy is the solution. The symbiosis is the survival.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new datasets were generated. All quantitative data cited in Table 1 are drawn from published peer-reviewed sources, cited therein.

Acknowledgments: The author expresses deep gratitude to the three reviewers whose rigorous and constructive engagement substantially strengthened this manuscript. Special tribute is due to the global community of researchers striving for a more parsimonious, equitable, and epistemologically accountable Artificial Intelligence.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. van Wynsberghe, A. Sustainable AI: AI for sustainability and the sustainability of AI. *Ai Ethics* **2021**, *1*, 213–218. [CrossRef]
2. Schwartz, R.; Dodge, J.; Smith, N.A.; Etzioni, O. Green AI. *Commun. Acn* **2020**, *63*, 54–63. [CrossRef]
3. Kaack, L.H.; Donti, P.L.; Strubell, E.; Kamiya, G.; Creutzig, F.; Rolnick, D. Aligning artificial intelligence with climate change mitigation. *Nat. Clim. Chang.* **2022**, *12*, 518–527. [CrossRef]
4. Rissanen, J. Modeling by shortest data description. *Automatica* **1978**, *14*, 465–471. [CrossRef]
5. Rissanen, J. *Stochastic Complexity in Statistical Inquiry*; World Scientific: Singapore, 1989.
6. Fokoué, E. On the Optimal Machine Learning Algorithm Selection for Classification Tasks. *J. Stat. Theory Pract.* **2011**, *5*, 523–540. [CrossRef]
7. Strubell, E.; Ganesh, A.; McCallum, A. Energy and Policy Considerations for Deep Learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*; Association for Computational Linguistics: Florence, Italy, 2019; pp. 3645–3650. [CrossRef]
8. Patterson, D.; Gonzalez, J.; Le, Q.; Liang, C.; Munguia, L.M.; Rothchild, D.; So, D.R.; Texier, M.; Dean, J. Carbon emissions and large neural network training. *arXiv* **2021**, arXiv:2104.10350.
9. Bellman, R. *Dynamic Programming*; Princeton University Press: Princeton, NJ, USA, 1957.
10. Mytton, D. Data centre water consumption. *npj Clean Water* **2021**, *4*, 11. [CrossRef]
11. Xiao, T.; You, F. Environmental impact and net-zero pathways for sustainable artificial intelligence servers in the USA. *Nat. Sustain.* **2025**, *8*, 1541–1553. [CrossRef]
12. International Energy Agency. *Energy and AI Report 2025*; International Energy Agency: Paris, France, 2025.
13. Rolnick, D.; Donti, P.L.; Kaack, L.H.; Kochanski, K.; Lacoste, A.; Sankaran, K.; Ross, A.S.; Milojevic-Dupont, N.; Jaques, N.; Waldman-Brown, A.; et al. Tackling climate change with machine learning. *ACM Comput. Surv.* **2022**, *55*, 1–96. [CrossRef]
14. Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Žídek, A.; Potapenko, A.; et al. Highly accurate protein structure prediction with AlphaFold. *Nature* **2021**, *596*, 583–589. [CrossRef] [PubMed]
15. Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T.B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; Amodei, D. Scaling laws for neural language models. *arXiv* **2020**, arXiv:2001.08361.
16. Wei, J.; Tay, Y.; Bommasani, R.; Raffel, C.; Zoph, B.; Borgeaud, S.; Yogatama, D.; Bosma, M.; Zhou, D.; Metzler, D.; et al. Emergent abilities of large language models. *arXiv* **2022**, arXiv:2206.07682.
17. Dettmers, T.; Lewis, M.; Belkada, Y.; Zettlemoyer, L. GPT3.int8(): 8-bit matrix multiplication for transformers at scale. *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 30318–30332. [CrossRef]
18. Wang, Y.; Zhang, Y.; Wang, Y.; Yin, N.; Kwok, J.; Yang, Q. Beyond scaleup: Knowledge-aware parsimony learning from deep networks. *arXiv* **2024**, arXiv:2407.00478.
19. Bender, E.M.; Gebru, T.; McMillan-Major, A.; Shmitchell, S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*; Association for Computing Machinery: New York, NY, USA, 2021; pp. 610–623. [CrossRef]
20. Crawford, K. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*; Yale University Press: New Haven, CT, USA, 2021.
21. Ligozat, A.L.; Lefèvre, J.; Bugeau, A.; Combaz, J. Unraveling the hidden environmental impacts of AI solutions for environment life cycle assessment of AI solutions. *Sustainability* **2022**, *14*, 5172. [CrossRef]
22. Donoho, D.L. Compressed sensing. *IEEE Trans. Inf. Theory* **2006**, *52*, 1289–1306. [CrossRef]
23. ISO/IEC 42001:2023; Information Technology—Artificial Intelligence—Management System. International Organization for Standardization/International Electrotechnical Commission: Geneva, Switzerland, 2023.
24. European Parliament and Council of the European Union. *Regulation (EU) 2024/1689 of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*; European Parliament and Council of the European Union: Brussels, Belgium, 2024.
25. International Telecommunication Union. *ITU-T L.1801: Methodology for Assessing the Environmental Impact of Artificial Intelligence*; ITU: Geneva, Switzerland, 2026.
26. Haigh, T.; Priestley, M.; Rope, C. *ENIAC in Action: Making and Remaking the Modern Computer*; MIT Press: Cambridge, MA, USA, 2016.
27. DeepSeek-AI. DeepSeek-V3 Technical Report. *arXiv* **2024**, arXiv:2412.19437.

28. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*, 2nd ed.; Wiley-Interscience: Hoboken, NJ, USA, 2006.
29. Sipser, M. *Introduction to the Theory of Computation*, 3rd ed.; Cengage Learning: Stamford, CT, USA, 2013.
30. Masanet, E.; Shehabi, A.; Lei, N.; Smith, S.; Koomey, J. Recalibrating global data center energy-use estimates. *Science* **2020**, *367*, 984–986. [[CrossRef](#)] [[PubMed](#)]
31. Riordan, M.; Hoddeson, L. *Crystal Fire: The Invention of the Transistor and the Birth of the Information Age*; W.W. Norton: New York, NY, USA, 1997.
32. Thompson, N.; Ge, Y.; Ferguson, A. On the origin of algorithmic progress in AI. *arXiv* **2025**, arXiv:2511.21622.
33. *Proceedings of the CPAL 2026: Third Conference on Parsimony and Learning*; ELLIS Institute Tübingen, in conjunction with Max Planck Institute for Intelligent Systems: Tübingen, Germany, 2026.
34. DeepSeek-AI. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv* **2025**, arXiv:2501.12948.
35. Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H.W.; Sutton, C.; Gehrmann, S.; et al. PaLM: Scaling language modeling with pathways. *J. Mach. Learn. Res.* **2023**, *24*, 11324–11436.
36. OpenAI. *GPT-4 Technical Report*; OpenAI: San Francisco, CA, USA, 2023.
37. Luccioni, A.S.; Viguier, S.; Ligozat, A.L. Estimating the carbon footprint of BLOOM, a 176B parameter language model. *J. Mach. Learn. Res.* **2023**, *24*, 253. .
38. DeepSeek-AI. DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model. *arXiv* **2024**, arXiv:2405.04434.
39. Financial Times. *Nvidia Loses \$589 Billion in Market Value After DeepSeek Shock*; in Market Value After DeepSeek Shock; Financial Times: London, UK, 2025.
40. Brin, S.; Page, L. The anatomy of a large-scale hypertextual Web search engine. *Comput. Netw. Isdn Syst.* **1998**, *30*, 107–117. [[CrossRef](#)]
41. Akaike, H. A new look at the statistical model identification. *IEEE Trans. Autom. Control* **1974**, *19*, 716–723. [[CrossRef](#)]
42. Schwarz, G. Estimating the dimension of a model. *Ann. Stat.* **1978**, *6*, 461–464. [[CrossRef](#)]
43. Ferraz-Caetano, J. The Artificial Intelligence Explanatory Trade-Off on the Logic of Discovery in Chemistry. *Philosophies* **2023**, *8*, 17. [[CrossRef](#)]
44. Reyes, R.L.; Regala, J.D. Reimagining the Philosophy of Experimentation in Chemistry Education: Embracing AI as a Tool for Scientific Inquiry. *Sci. Educ.* **2025**, *35*, 709–754. [[CrossRef](#)]
45. de Vries-Gao, A. The carbon and water footprints of data centers and what this could mean for artificial intelligence. *Patterns* **2026**, *7*, 101430. [[CrossRef](#)] [[PubMed](#)]
46. Lincoln Institute of Land Policy. *Data Drain: The Land and Water Impacts of the AI Boom*. *Land Lines Magazine*, 17 October 2025. [[CrossRef](#)]
47. Environmental and Energy Study Institute. *Data Centers and Water Consumption*; Environmental and Energy Study Institute: Washington, DC, USA, 2025.
48. Ebert, K.; Alder, N.; Herbrich, R.; Hacker, P. AI, climate, and regulation: From data centers to the AI act. *arXiv* **2024**, arXiv:2410.06681.
49. Okechukwu Paul-Chima, U. Implementing artificial intelligence and machine learning algorithms for optimized crop management: A systematic review on data-driven approach to enhancing resource use and agricultural sustainability. *Cogent Food Agric.* **2025**, *11*, 2569982. [[CrossRef](#)]
50. Landauer, R. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* **1961**, *5*, 183–191. [[CrossRef](#)]
51. Hardy, G.H. *A Mathematician's Apology*; Cambridge University Press: Cambridge, UK, 1940.
52. Fisher, R.A. *Statistical Methods for Research Workers*; Oliver and Boyd: Edinburgh, UK, 1925.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.