

Article

Not peer-reviewed version

Post-Turing Architectures: Mindful Machines and the Energy Efficiency of Knowing

[Rao Mikkilineni](#) *

Posted Date: 9 October 2025

doi: 10.20944/preprints202510.0680.v1

Keywords: mindful machines; post-turing paradigm; artificial intelligence; digital genome; auto-poietic systems; cognitive architecture; ethical AI



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

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.

Article

Post-Turing Architectures: Mindful Machines and the Energy Efficiency of Knowing

Rao Mikkilineni

Center for Business Innovation. Golden Gate University. San Francisco, USA; rmikkilineni@ggu.edu

Abstract

The prevailing AI paradigm, grounded in the Turing model, operationalizes intelligence as symbol manipulation and scales performance through parameter growth and stochastic optimization. While this approach underlies deep learning and large language models, its reliance on brute-force computation imposes severe energy and epistemic costs: training a single foundation model such as GPT-3 requires over 1,200 MWh, and inference at scale accelerates unsustainable data-center expansion. Moreover, statistical generalization without semantic grounding limits interpretability and adaptability. This paper proposes Mindful Machines as a post-Turing, knowledge-centric alternative to computation-centric architectures. Mindful Machines integrate three interdependent layers—physical embodiment for environmental coupling, representational structures for symbolic and sub-symbolic integration, and reflective mechanisms for self-regulation and autopoiesis. This architecture enables adaptive knowledge reuse, goal-driven learning, and dynamic constraint satisfaction, reducing retraining overhead while enhancing semantic depth. By formalizing cognition as a triadic system rather than a unidimensional computational process, we demonstrate how meaning-driven architectures can achieve functional scalability with significantly lower energy budgets. The results suggest a paradigm shift toward sustainable AGI, where intelligence emerges from structured interaction and self-maintenance rather than parameter escalation.

Keywords: mindful machines; post-turing paradigm; artificial intelligence; digital genome; autopoietic systems; cognitive architecture; ethical AI

1. Introduction

Artificial Intelligence (AI) has advanced rapidly under the Turing paradigm, which conceptualizes intelligence as computation—symbol manipulation governed by fixed syntactic rules [1–4]. This model has dominated the design of digital computers and underpins the logic of machine learning and deep learning systems. Within this framework, progress is typically measured by scaling: larger datasets, deeper networks, and more compute yield incremental performance gains.

Recent foundation models, such as GPT-3 and its successors, exemplify this trajectory [5,6]. Trained on trillions of tokens and powered by billions of parameters, they approximate human-like language and reasoning through massive statistical interpolation. Yet this success comes with conceptual and material limits.

Despite impressive benchmark results, large models lack semantic grounding, transparent interpretability, and adaptive reasoning [7]. Their knowledge arises from correlations rather than structured understanding, resulting in brittle generalization, hallucination, and vulnerability to adversarial perturbations. Moreover, epistemic returns diminish as model size grows—the additional compute produces marginal cognitive gain—while the cost of incremental improvement escalates steeply [8].

The energy and power implications of this scaling are profound. Estimates suggest that training a single model of GPT-3's scale can consume over 1,200 MWh of electricity—equivalent to the annual energy use of more than 100 U.S. homes [9]. Continuous inference across millions of users compounds this footprint, driving exponential data-center growth and significant carbon emissions. These trends

raise sustainability concerns and challenge the viability of brute-force scaling as a pathway to Artificial General Intelligence (AGI).

To address these limitations, we propose a shift from computation-centric to meaning-centric architectures—an approach embodied in the emerging framework of Mindful Machines [10–12]. Rooted in post-Turing computation, the General Theory of Information [13–15], and autopoietic principles from biological systems [16,17], Mindful Machines integrate:

1. Physical embodiment (4E cognition: embodied, embedded, enactive, extended);
2. Knowledge representation (Digital Genomes encoded through Named Sets and Fundamental Triads); and
3. Reflective self-regulation (cognizing and meta-cognitive oracles).

Together, these components form an autopoietic system capable of adaptive knowledge reuse, contextual learning, and prudent self-maintenance. By distinguishing knowledge acquisition from knowledge use, Mindful Machines minimize energy-intensive retraining while enhancing semantic depth, coherence, and resilience.

The theory and practice of Mindful Machines have been discussed in prior works on knowledge-centric computing and structural machines. Building upon that foundation, this paper pursues three interrelated goals:

1. To examine the shortcomings of current AI, Gen-AI, and AGI paradigms, focusing on epistemic and energetic inefficiencies;
2. To present the Mindful Machine architecture, highlighting how its triadic design overcomes these deficiencies; and
3. To explore practical strategies for deploying Mindful Machines as sustainable alternatives, with quantitative analysis of their potential to reduce energy and power consumption.

The remainder of this paper is organized as follows:

- Section 2 reviews the conceptual and energy limitations of current AI architectures.
- Section 3 introduces the theoretical foundations and triadic structure of Mindful Machines.
- Section 4 presents a roadmap for implementation and evaluation, including energy-efficiency design principles.
- Section 5 concludes with implications for sustainable, knowledge-centric AGI development.

2. AI Evolution and the Energy and Power Dilemma

Artificial Intelligence has progressed through successive conceptual eras—symbolic, statistical, and scaling—each redefining what it means to “compute intelligence.” Early symbolic AI (1950 s–1980 s) sought rationality through logic and explicit rule-based systems, mirroring Turing’s syntactic view of mind as computation. By the late 1990 s, limitations in brittleness and hand-crafted knowledge led to statistical learning, which exchanged logical transparency for data-driven inference. The deep-learning revolution (2012 onward) fused massive datasets, GPUs, and gradient-descent optimization into a single thesis: *scale is intelligence*.

2.1. Subsection

Transformer architectures [18] demonstrated that performance improves monotonically with the number of parameters and tokens—codified in Kaplan et al.’s “scaling laws [5].” This observation triggered an arms race of parameter expansion, from 100 million-parameter models in 2018 to multi-trillion-parameter systems today. Each generation multiplies computational demand roughly tenfold, while accuracy or reasoning improvement grows sub-linearly.

The implicit equation became:

Intelligence ∝ Compute^{0.5} × Data^{0.3}

—a heuristic that favors hardware escalation over conceptual refinement.

2.2. *The Material Cost of Intelligence*

The material footprint of this strategy is striking. Training GPT-3 (175 B parameters) required an estimated 1.3×10^3 MWh of electricity \approx 550 tons CO₂-e, and GPT-4 reportedly consumed several multiples of that figure. As of 2024, data-center electricity used for AI workloads surpassed 4% of global ICT energy and is projected to double by 2030. Each inference—every user query—draws power from accelerator clusters whose embodied carbon often exceeds the operational footprint of earlier supercomputers. Thus, the epistemic gain of “next-token prediction” is purchased at ecological expense.

From a philosophical standpoint, this energy–epistemology gap exposes the limits of the Turing paradigm. If intelligence is defined solely as the efficient transformation of symbols, then scaling compute seems natural. Yet biological and cognitive systems achieve comparable or greater adaptability at microwatt-to-watt scales through structural knowledge reuse and self-organization. The contrast between the neural watt and the silicon megawatt underscores an unsolved dilemma: the current AI trajectory maximizes correlation accuracy but minimizes energetic prudence.

2.3. *Diminishing Epistemic Returns*

Empirical evidence shows diminishing returns beyond a certain compute threshold. Model perplexity decreases predictably with size, but semantic coherence and truth-consistency plateau. For instance, the jump from 175 B to 1 T parameters increases training cost roughly 10× while improving benchmark accuracy by < 5%. This asymptotic curve mirrors thermodynamic limits—more energy yields less new knowledge. Hence, progress measured by FLOPs risks violating the principle of informational economy: meaning cannot be brute-forced.

In conclusion, the scaling laws of Large Language Models (LLMs), which show that performance improvements grow sub-linearly (following a power law with a small exponent) with respect to the computational resources (compute) used for training, while the increase in model size and data leads to a multiplicative, and thus exponential, rise in computational demand. This trend—where achieving modest gains requires disproportionately massive compute and model expansion—validates the idea that the development trajectory became a “hardware escalation over conceptual refinement” heuristic following the initial observations by Kaplan et al. [5]. Subsequent work, like the Chinchilla paper [6], verified these diminishing returns and critiqued the pure parameter-scaling arms race, arguing that compute-optimal training requires balanced scaling of both model size and data size to maximize efficiency. Biological systems, in contrast, achieve adaptability at microwatt-to-watt scales through structural knowledge reuse and self-organization [19]

2.4. *The Power Dilemma: Sustainability vs. Ambition*

The ethical and infrastructural implications are immediate. Mega-model training now competes with national grids for renewable allocation; water-cooling requirements affect local ecosystems; and inference at scale introduces continuous-load baselines that hinder grid flexibility. Organizations pursuing AGI confront a paradox: the quest for artificial omniscience threatens environmental resilience—the very substrate of future intelligence.

2.5. *Toward Knowledge-Centric Efficiency*

Resolving this dilemma requires redefining intelligence from energy-consuming computation to energy-coherent cognition. Whereas current models equate learning with parameter adjustment, a knowledge-centric approach—exemplified by Mindful Machines—separates knowledge acquisition (occasional, costly) from knowledge use and regulation (continuous, low-power). Through autopoietic replication, redundancy management, and meta-cognitive self-regulation, such systems achieve:

1. Order-of-magnitude reductions in compute per decision (\approx 4–10×) via routing, quantization, and retrieval-first reasoning;

2. Avoidance of periodic retraining, saving $\approx 10^3$ MWh per frontier-scale cycle;
3. Dynamic energy proportionality, scaling consumption with epistemic novelty rather than query volume.

These principles shift AI's growth curve from exponential energy dependence toward logarithmic efficiency, aligning computation with the thermodynamics of life.

3. Conceptual Foundations of Mindful Machines

3.1. Toward Knowledge-Centric Efficiency

The Turing paradigm treats cognition as symbol manipulation—powerful yet syntactically closed. Mindful Machines expand this view by integrating meaning and self-maintenance into computation. They draw upon Mark Burgin's General Theory of Information (GTI), which distinguishes three inseparable aspects of information [13–15]:

- Ontological (being)—the structures that exist;
- Epistemic (knowing)—the knowledge possessed by an agent; and
- Pragmatic (doing)—the use of that knowledge in action.

These dimensions correspond respectively to body, brain, and mind in an intelligent system. Intelligence therefore emerges not from processing alone, but from coherence among these triadic layers.

3.2. The Burgin–Mikkilineni Thesis (BMT)

The BMT [10–12,17] extends GTI into engineering practice, proposing that information processing systems evolve from symbolic to structural to autopoietic forms. Symbolic systems manipulate representations; structural systems maintain their organization; autopoietic systems reproduce and regulate themselves. In this hierarchy, Mindful Machines occupy the autopoietic tier. They manage their own computation, communication, and coordination—what Burgin and Mikkilineni term structural intelligence. Such systems integrate data, knowledge, and action through reflexive feedback loops rather than static pipelines.

3.3. The Triadic Architecture: Body – Brain – Mind

Mindful Machines embody a triadic architecture (Figure 1), aligning with GTI's ontological triad and 4E cognition theory.

- Body (4E Cognition: embodied, embedded, enactive, extended)) [20]:
- Implements the Embodied, Embedded, Enactive, and Extended dimensions of cognition. Sensors, actuators, and APIs serve as the organism's limbs—continuously coupling the system to its environment and enabling sense–act loops.
- Brain (Knowledge Representation):
- Realized through Large Language Models (LLMs), transformers, and graph-based schemas that organize experience into structured knowledge. Within the Digital Genome, this knowledge is codified as Named Sets and Fundamental Triads—machine-readable relationships that express purpose, constraints, and best-practice policies. The genome evolves incrementally; it is interpreted and refined, not perpetually retrained.
- Mind (Cognizing & Meta-Cognitive Oracles):
- Uses the Digital Genome to reason, infer, and act with prudence and purpose. Meta-cognition provides a reflective layer—Monitor → Evaluate → Adapt—that observes the system's own cognition, ensuring coherence between goals, context, and behavior. This enables self-explanation (“why I acted”), bias correction, and ethical constraint.

Together these layers form an autopoietic loop: perception informs knowledge; knowledge guides reasoning; reasoning reconfigures interaction; interaction refines perception. The machine thus sustains its identity through continuous renewal.

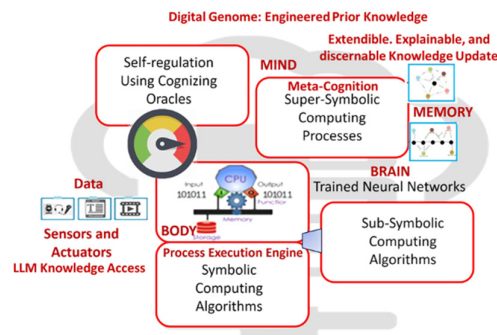


Figure 1. Mindful Machines separating data processing and knowledge acquisition from reasoning and acting.

3.4. Autopoiesis and Self-Regulation

In this paper we discuss an architectural shift, moving from the current computation-intensive deep learning paradigm to a knowledge-centric one, inspired by biological systems. It asserts that true efficiency and resilience lie in the concept of “cloud autopoiesis”—digital services that self-produce, manage, and heal their components, borrowing from Maturana and Varela’s definition of living systems. This foundational autopoietic orchestration ensures the system remains reliable despite hardware instability and that its energy consumption is proportional to its activity rather than its sheer size.

The core of this new model, called Mindful Machines, is the crucial separation between data processing and high-level reasoning. Current large language models (LLMs) are relegated to the role of translators—converting sensory data into structured knowledge. Once knowledge is encoded, subsequent reasoning, regulation, and response generation proceed through lightweight cognitive processes. This separation yields a knowledge-centric economy of power, where high energy expenditure is only necessary for acquiring epistemic novelty (new, non-consensus knowledge), drastically reducing the energy required for routine inference and reasoning, thus mirroring the extreme energy efficiency of the human brain. Figure 1 shows the architecture of Mindful Machine Implementation.

We propose a new model for energy consumption that takes into account the architecture of Mindful Machines.

$$E_{\text{mindful}} \approx 0.1\text{--}0.25E_{\text{AGI}}$$

This formula introduces a concept of a “Mindful” (or potentially “Knowledge-Centric”) AI model, contrasting its energy consumption with that of a typical large AGI-type model.

This relationship proposes that a more efficient, “mindful” AI architecture can achieve comparable or better results on specific tasks while consuming only 10% to 25% of the energy required by a brute-force AGI-style model.

Architectural Efficiency: The justification for this dramatic reduction often stems from a paradigm shift away from simply scaling up massive general-purpose models. It advocates for techniques like:

Specialization: Using smaller, specialized models tailored to the task instead of one colossal model for all tasks.

Algorithmic Optimization: Employing advanced techniques like quantization (using fewer bits for calculations) to reduce energy per computation.

Inference Optimization: Reducing energy demands during use by, for example, generating shorter and more efficient responses.

In short, the second formula is a statement about sustainability and efficiency, arguing that intelligent AI does not necessarily have to follow the energy-hungry scaling curve of the first formula

3.4. Epistemological Implications

Philosophically, Mindful Machines redefine intelligence as coherence across levels of knowing. They embody a *recursive epistemology*: the system both acts in the world and reflects on the adequacy of its knowledge. In doing so, they operationalize Deutsch's notion of *good explanations*—models that are hard to vary yet easy to apply. Meaning, not magnitude, becomes the metric of progress.

4. Implementation and Evaluation Roadmap: From Theory to Practice

4.1. Overview

Translating the Mindful Machine paradigm from theory to implementation requires an engineering roadmap that preserves its philosophical essence—knowledge-centric, self-regulating, and energy-efficient.

This roadmap must demonstrate that post-Turing computation can be realized with current digital infrastructure while maintaining verifiable gains in energy and epistemic efficiency. While a few prototype demonstrations of Mindful Machines have proven the theory and implementation feasibility, this roadmap provides a new direction for energy reduction strategy.

1. The approach proceeds in four phases:
2. Encoding knowledge into a Digital Genome;
3. Deploying autopoietic infrastructure capable of self-replication and healing;
4. Activating cognitive and meta-cognitive processes through Cognizing Oracles; and
5. Evaluating energy use, semantic quality, and coherence as integrated performance metrics.

4.2. Phase 1: Knowledge Encoding Through the Digital Genome

The Digital Genome (DG) is the blueprint of a Mindful Machine. It formalizes goals, constraints, and operational policies using the logic of Named Sets and Fundamental Triads (GTI). Implementation begins with three repositories:

1. Functional Knowledge Base – defines what is to be achieved (e.g., billing, prediction, diagnosis).
2. Non-Functional Knowledge Base – specifies how performance, resilience, and compliance are to be maintained.
3. Historical and Episodic Memory – stores traces of prior states, allowing the system to learn prudently rather than repetitively.

The Digital Genome is version-controlled and interpretable—its evolution is documented, not opaque. This replaces the opaque “weight space” of deep models with a living knowledge substrate amenable to audit and modification

4.3. Phase 1: Phase 2: Autopoietic Infrastructure

Implementation of the body–brain–mind triad depends on an underlying autopoietic cloud infrastructure that manages computation as a living organism manages metabolism. Using containerized micro-services (e.g., Docker + Kubernetes), the infrastructure includes:

1. Autopoietic Process Manager (APM): deploys, scales, and replaces components dynamically based on health and demand.
2. Cognitive Network Manager (CNM): maintains dynamic inter-service connectivity and ensures data coherence.
3. Software Workflow Manager (SWM): orchestrates logical task flow, providing recovery and re-execution on failure.

These managers embody digital autopoiesis: replication, redundancy, and repair operating continuously to sustain coherence on unreliable hardware. Energy-aware orchestration is achieved through autoscaling, hibernation, and carbon-intensity-aware placement—features already available in major cloud platforms but coordinated here under cognitive control.

4.4. Phase 3: Cognizing and Meta-Cognitive Oracles

The Mind layer activates through a federation of Cognizing Oracles—modular reasoning units that apply the knowledge encoded in the DG. Each oracle performs a specialized cognitive function: explanation, planning, ethical reasoning, or optimization. Meta-cognitive control loops oversee these oracles:

1. Monitor: track reasoning outcomes and system energy state.
2. Evaluate: compare outcomes against DG goals and energy budgets.
3. Adapt: modify reasoning pathways, thresholds, or resource allocation.

This recursive self-observation enables prudence - the ability to decide when to compute, when to reuse knowledge, and when to abstain.

4.5. Phase 4: Evaluation Framework

Evaluation extends beyond accuracy metrics to include semantic, epistemic, and energetic dimensions.

4.5.1. Energy and Power Metrics

Adopting lifecycle analysis standards (Patterson et al., 2021):

$$E_{\text{total}} = E_{\text{train}} + E_{\text{infer}} + E_{\text{infra}}$$

Training Energy (E_{train}): minimized through knowledge reuse and one-time genome encoding.

Inference Energy (infer): measured in Joules per token or per decision (J/token or J/task).

Infrastructure Energy (infra): includes autoscaling overhead and cooling, quantified through power-usage-effectiveness (PUE).

Typical targets for well-optimized Mindful Machine deployments are:

1. Inference energy reduction: 4–10× compared with uniform large-model inference.
2. Training energy avoidance: $\approx 1,000$ MWh per frontier-scale model not retrained.
3. Overall energy-per-decision: 0.1–0.25 of conventional AGI baseline.

4.5.2. Semantic and Epistemic Metrics

1. Semantic Coherence: alignment between generated explanations and DG knowledge structures.
2. Epistemic Efficiency: ratio of meaningful outputs to total energy consumed (bits of validated knowledge per kWh).
3. Reflective Integrity: consistency between cognition and meta-cognition—measured via self-reporting logs of reasoning adaptation.

These metrics together evaluate how well a system knows, not merely how much it computes.

4.6. Energy-Efficiency Design Principles

Knowledge > Computation: favor retrieval and rule execution before invoking generative inference.

1. Dynamic Right-Sizing: allocate model scale to epistemic difficulty, not request volume.
2. Sparse Activation: activate only relevant oracles; adopt mixture-of-experts architectures.
3. Energy as a Service-Level Objective: treat joules-per-token as a first-class optimization parameter.
4. Autopoietic Balance: maintain redundancy only where it adds resilience or interpretability.
5. Reflexive Learning: meta-cognition continuously evaluates energy–knowledge trade-offs.

These principles integrate philosophy and engineering: prudence becomes a control variable.

4.7. Demonstration and Verification Path

Pilot implementations can proceed in three tiers:

1. Micro-scale — single application (e.g., customer-billing predictor) with DG and cognitive orchestration.
2. Meso-scale — enterprise-level deployment where multiple oracles coordinate across services.
3. Macro-scale — federated Mindful Machines collaborating across organizations, forming a knowledge-centric ecosystem.

Each tier reports energy telemetry (Wh/query), semantic metrics, and reflexive logs for longitudinal evaluation

4.8. Executive Synthesis

Implementation of Mindful Machines converts energy waste into informational prudence:

Aspect	Conventional AI	Mindful Machines
Learning Mode	Continuous retraining	Knowledge reuse + episodic update
Energy Profile	Exponential with scale	Logarithmic with novelty
Inference Routing	Uniform large model	Difficulty-based adaptive routing
Meta-Control	Absent	Present (Monitor ? Evaluate ? Adapt)
Reliability	Hardware redundancy	Autopoietic replication
Epistemic Metric	Accuracy	Coherence / Meaning per kWh

In doing so, Mindful Machines realign intelligence with the principle of informational economy—the idea that sustainable cognition requires knowledge to govern power, not the reverse

5. Conclusions and Future Implications for Sustainable AGI

The history of AI demonstrates that increasing computational scale does not necessarily yield higher intelligence—only higher energy consumption. The Turing paradigm equated mind with mechanical manipulation, but as the limits of scaling emerge, this equation appears incomplete. Intelligence, whether natural or artificial, cannot be measured solely by parameter count or teraflops per second. It must also be measured by meaning per joule—the capacity to generate coherent, contextually grounded, and energetically prudent knowledge.

1. Mindful Machines offer a pragmatic and philosophical response to this realization.
2. They redefine intelligence as a relationship among three interdependent dimensions:
3. Body (interaction and perception)—the embodied link to the environment;
4. Brain (knowledge representation)—the structured memory of understanding; and
5. Mind (reasoning and meta-cognition)—the reflective process that maintains coherence and purpose.

Through this triadic coherence, Mindful Machines integrate epistemology, computation, and energy management into a single framework, realizing what earlier AI systems could only simulate: self-regulated intelligence aligned with life’s organizing principles.

At its core, the Mindful Machine paradigm challenges the computational metaphor of mind that dominated the 20th century. Where Turing Machines execute externally defined instructions, Mindful Machines generate and govern their own rules through autopoietic self-regulation. Where deep learning relies on statistical correlation, Mindful Machines pursue semantic coherence—knowledge that sustains itself through explanation and reflection. And where current AI consumes vast energy for marginal improvement, Mindful Machines embody energy coherence—a thermodynamic harmony between knowing and doing.

This transformation aligns with the General Theory of Information (GTI), which defines information as a triadic relationship among the ontological, epistemic, and pragmatic domains.

Mindful Machines make this relationship operational: they are systems that know what they are, know what they know, and know how to sustain that knowing.

Philosophically, this shift moves AI beyond simulation toward participation in meaning-making processes. It marks a move from syntactic intelligence to semantic and reflexive intelligence—a step that parallels the evolution of life itself.

The pursuit of Artificial General Intelligence (AGI) is no longer merely a technical ambition; it is an ecological and ethical problem. Current AI trajectories replicate the logic of industrial growth—more data, more power, more carbon—without regard to planetary limits. By contrast, Mindful Machines embody a philosophy of sufficiency: using the least energy required to maintain knowledge coherence and fulfill purpose. This has direct implications for ethical AI design:

1. Energy Accountability: every inference should have a measurable energy signature.
2. Knowledge Transparency: every decision should trace back to its Digital Genome.
3. Autopoietic Responsibility: systems should self-regulate within human-defined ecological and moral boundaries.

In this light, prudence—an ancient virtue often neglected in modern engineering—becomes the new form of intelligence.

5.1. Future Research Directions

Future research in Mindful Machine development should focus on:

1. Quantifying epistemic efficiency—developing formal metrics linking energy use to the value of generated knowledge (bits of validated meaning per kWh).
2. Cross-layer coherence modeling—investigating how body–brain–mind synchronization can be optimized through self-organizing feedback.
3. Cognizing Oracles and meta-cognitive control loops—exploring architectures that enable reflection, ethical reasoning, and adaptive learning without retraining.
4. Energy-proportional intelligence frameworks—integrating power-aware scheduling, carbon-intensity tracking, and epistemic novelty detection into orchestration managers.
5. Policy and governance frameworks—creating standards for Knowledge per Watt (K/W) as a complement to existing performance benchmarks.

These directions extend the philosophy of computation into the realm of sustainable design, where epistemology and thermodynamics converge.

5.2. Toward a New Definition of Artificial General Intelligence

In the conventional narrative, AGI is imagined as a singular system capable of mastering any task—a synthetic omniscience achieved through scale. In the Mindful Machine narrative, AGI becomes something more modest yet more profound: a distributed ecology of autopoietic agents, each managing its own knowledge, energy, and purpose in coherence with others. Such a network does not seek dominance but harmony—a planetary intelligence that mirrors life's cooperative equilibrium. Hence, the question is no longer When will AGI arrive? but How will we know when our machines have learned to think wisely and live lightly? Mindful Machines suggest the answer lies not in greater power but in greater coherence—between matter, energy, and meaning.

In conclusion, we observe that the evolution of AI is approaching a philosophical crossroads. One path continues toward massive computation, opaque models, and ecological exhaustion; the other leads toward knowledge-centric, self-regulating systems that understand both their environment and their limits. Mindful Machines illuminate this latter path. They demonstrate that the future of intelligence—human or artificial—depends not on power but on prudence, not on endless training but on conscious knowing.

Sustainable AGI will not be built by machines that think like humans, but by systems that learn to think with the wisdom of nature.

Author Contributions: R. M is responsible for the article and has extensively used ChatGPT, Gemini, and Copilot for accessing global knowledge, analyze and synthesize it to develop the Mindful Machines.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: During the preparation of this manuscript/study, the author used ChatGPT 5o. Gemini 2.5, Copilot for the purposes of research. Access to global knowledge and assistance in writing. The author has reviewed and edited the output which is solely analyzed and created by the author and takes full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Turing, A.M. *The Essential Turing*; Copeland, B.J., Ed.; Oxford University Press: Oxford, UK, 2004. [Google Scholar]
2. Deutsch, D. (2011). *The Beginning of Infinity: Explanations that Transform the World*. Viking Press.
3. Dodig-Crnkovic, G. (2012). *Information and Computation: Computational Natural Philosophy*. IGI Global. <https://doi.org/10.4018/978-1-4666-1788-5>
4. Dodig Crnkovic, G. Significance of Models of Computation, from Turing Model to Natural Computation. *Minds Mach.* 2011, 21, 301–322. [Google Scholar] [CrossRef]
5. Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chen, B., Dickens, R., Hernandez, D., Jain, A., Leike, J., & Pope, M. (2020). Scaling laws for neural language models. *arXiv Preprint arXiv:2001.08361*. <https://arxiv.org/abs/2001.08361>
6. Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, R., de Maas, D. G., Lengyel, K., van den Oord, A., & Clark, J. (2022). Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
7. Sejnowski, T. J. (2022). Large language models and the reverse Turing test. *Neural Computation*, 34(10), 1953–1989. https://doi.org/10.1162/neco_a_01518
8. Patterson, D., Gonzalez, J., Le, Q. V., Dean, J., Jouti, N. P., & Young, C. (2021). Carbon emissions and large neural network training. *arXiv Preprint arXiv:2104.10350*. <https://arxiv.org/abs/2104.10350>
9. Vinuesa, R., Azizpour, H., Leite, I. et al. The role of artificial intelligence in achieving the Sustainable Development Goals. *Nat Commun* 11, 233 (2020). <https://doi.org/10.1038/s41467-019-14108-y>
10. Mikkilineni, R. (2025). General Theory of Information and Mindful Machines. *Proceedings*, 126(1), 3. <https://doi.org/10.3390/proceedings2025126003>
11. Mikkilineni, R. (2024). General Theory of Information: The Bridge to Mindful Machines. *Preprints*. <https://doi.org/10.20944/preprints202410.2564.v1>
12. Mikkilineni, R., & Kelly, W. P. (2024). Machine Intelligence with Associative Memory and Event-Driven Transaction History. *Preprints*. <https://doi.org/10.20944/preprints202404.1298.v1>
13. Burgin, M. (2010). *Theory of Information: Fundamentality, Diversity and Unification*. World Scientific. <https://doi.org/10.1142/7498>
14. Burgin, M. (2022). *The General Theory of Information and Computational Intelligence*. Springer.
15. Burgin, M., & Mikkilineni, R. (2021). The Triadic Mind: Structural Machines and Autopoietic Computation. *Information Sciences*, 580, 562–578. <https://doi.org/10.1016/j.ins.2021.09.021>
16. Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel Publishing Company.
17. Mikkilineni, R. (2022). A New Class of Autopoietic and Cognitive Machines. *Information*, 13(1), 24. <https://doi.org/10.3390/info13010024>

18. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
19. Levin, M. The computational boundary of a 'Self': developmental bioelectricity drives multicellularity and scale-free cognition. *Front. Psychol.* 10, 2688 (2019)
20. Froese, T., & Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*, 173(3–4), 466–500. <https://doi.org/10.1016/j.artint.2008.12.001>

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.