

The Hidden Life-Ground of Artificial Intelligence

Carbon, Water, Land, and the Life-Coherent Governance of Symbolic Power

From AI as Idol and Enclosure to AI as Tool and Commons

Dr. Bichara Sahely

ACADEMIC WHITE PAPER

June 2026



AI AS
IDOL AND
ENCLOSURE

EXTRACTION.
CONCENTRATION.
DISPOSSESSION.
FRAGILITY.

COMPUTE INFRASTRUCTURE
CHIPS • SERVERS • NETWORKS • COOLING
DATA CENTERS

⚡ ENERGY

💧 WATER

🌿 LAND

🏠 MINERALS

👥 LABOR

🌳 ECOSYSTEMS

🏘️ COMMUNITIES

🌐 LIFE'S WEB
INTERDEPENDENT
AND FINITE

AI AS
TOOL AND
COMMONS

STEWARDSHIP.
SHARING.
REGENERATION.
FLOURISHING.

The Hidden Life-Ground of Artificial Intelligence

Carbon, Water, Land, and the Life-Coherent Governance of Symbolic Power

From AI as Idol and Enclosure to AI as Tool and Commons

Academic White Paper

Dr. Bichara Sahely, BSc (Biology), MBBS, DM (Internal Medicine)

8 June 2026

Suggested citation

Sahely, B. (2026). The hidden life-ground of artificial intelligence: Carbon, water, land, and the life-coherent governance of symbolic power: From AI as idol and enclosure to AI as tool and commons. Academic White Paper.

Author Note

Dr. Bichara Sahely is a physician, public health thinker, and founder of the Life-Knowledge Commons at bsahely.com. His work explores the conditions of life, health, peace, ecology, economics, technology, and governance through a life-coherence framework. This academic white paper forms part of an ongoing series examining how contemporary institutions, technologies, and systems can be evaluated according to whether they expand or diminish life-capacity.

Acknowledgments

The author gratefully acknowledges the United Nations University Institute for Water, Environment and Health report *Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints* as the empirical foundation for this white paper. That report provides the carbon, water, and land footprint framing that made the present life-coherence interpretation possible. The conceptual synthesis, governance framework, and life-coherence analysis developed here are the responsibility of the author.

AI Transparency Note

This academic white paper was developed with the assistance of ChatGPT, an AI language model, in dialogue with the author. The author directed the conceptual framing, interpretive synthesis, life-coherence analysis, structure, and editorial judgment. AI assistance was used to support drafting, organization, refinement, and integration of the uploaded UNU-INWEH report into the author's broader Life-Knowledge Commons framework. All final interpretive responsibility rests with the author.

Methodological and Scope Note

This academic white paper is an interpretive synthesis and governance framework. It does not present an original lifecycle assessment, independent energy audit, or recalculation of AI's environmental footprints. Instead, it builds on the UNU-INWEH report and selected supporting literature to interpret AI's carbon, water, land, lifecycle, justice, and governance implications through the author's life-coherence framework. Numerical estimates are therefore treated as indicative of scale and governance significance, not as independently verified calculations by the author.

Abstract

Artificial intelligence is often experienced as an immaterial symbolic power: a prompt is entered, and language, images, code, predictions, summaries, or videos appear. Yet AI is not weightless. It depends on a hidden life-ground of electricity, water, land, minerals, labor, communities, ecosystems, and waste sinks. Building on the United Nations University Institute for Water, Environment and Health report *Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints*, this white paper extends the environmental accounting of AI into a life-coherence framework. It argues that AI's central governance question is not only how large its carbon, water, and land footprints are, but whether the conversion of life-support into symbolic output expands life-capacity or deepens symbolic excess, dependency, and enclosure.

The paper interprets AI as a hidden metabolism linking prompts, models, data centers, electricity, cooling, minerals, labor, e-waste, and ecological sinks. It distinguishes between model training and inference, highlights the escalating footprint of image and video generation, and examines the justice problem of local costs and distant benefits. It then develops the diagnostic framework of AI as Tool, Oracle, Idol, Enclosure, or Commons, before proposing a life-coherent governance framework centered on purpose, proportionality, transparency, sufficiency, lifecycle responsibility, place-based accountability, community consent, public-interest compute, knowledge integrity, and review and repair. A Caribbean and Small Island Developing States application is included to show how fragile grids, water constraints, climate vulnerability, and digital dependency make life-coherent AI governance especially urgent. The paper concludes with a practical Life-Coherent AI Use Protocol for individuals, institutions, governments, communities, and regional commons-building.

Keywords

Artificial intelligence; environmental footprint; carbon footprint; water footprint; land footprint; data centers; energy use; AI governance; life-coherence; symbolic power; digital enclosure; AI commons; Small Island Developing States; Caribbean; public-interest compute; lifecycle responsibility; environmental justice; technology ethics.

Executive Summary

Artificial intelligence is often experienced as an immaterial symbolic power. A prompt is entered, and language, images, code, predictions, summaries, or videos appear. Yet AI is not weightless. It depends on a hidden life-ground of electricity, water, land, minerals, labor, communities, ecosystems, and waste sinks. This white paper argues that AI must therefore be judged not only by what it can generate, predict, automate, or optimize, but by whether its use of the life-ground expands life-capacity within ecological and social limits.

Building on the United Nations University Institute for Water, Environment and Health report *Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints*, this paper extends AI environmental accounting into a life-coherence framework. The UNU-INWEH report is important because it moves beyond a carbon-only view of AI sustainability and examines how AI's electricity demand generates carbon, water, and land footprints. This multidimensional framing matters because low-carbon AI is not automatically low-water or low-land. Sustainability claims based on a single metric can conceal trade-offs, burden shifting, and local ecological stress.

The central argument of this paper is that AI is not merely a software layer or information service. It is a material infrastructure and metabolic system. Every prompt and output participates in a hidden chain: model development, data center processing, electricity demand, cooling, carbon-water-land footprints, mineral extraction, hardware manufacture, labor, e-waste, and ecological absorption. AI's visible symbolic outputs are therefore inseparable from the physical systems that make them possible.

The paper distinguishes between model training and inference. Training frontier models can require very large amounts of energy and computational capacity, but once AI systems are deployed, the repeated use of models across millions or billions of prompts can become the dominant operational burden. Inference is continuous, distributed, platform-shaped, and culturally normalized. This means AI governance cannot focus only on spectacular training runs. It must also govern everyday use at scale: prompts, summaries, searches, images, videos, embedded assistants, enterprise workflows, and default AI features.

A key concept developed in the paper is symbolic escalation. Not all AI outputs carry the same life-ground cost. Basic classification, short text, long text, image generation, and video generation differ significantly in resource intensity. A life-coherent AI system should therefore follow the principle of minimum sufficient symbolic form: if classification is enough, do not generate text; if text is enough, do not generate images; if images are enough, do not generate video. High-intensity outputs should be justified by real life-value, not normalized as default convenience or spectacle.

The paper also examines the justice problem of local costs and distant benefits. AI's benefits often accrue to platform companies, investors, powerful states, enterprise users, and distant consumers, while burdens may fall on communities near data centers, workers in supply chains, water-stressed regions, electricity grids, mining zones, and e-waste destinations. This separation of benefit from burden is a central life-coherence failure. The paper argues that consequences must return to decision-makers through carbon-water-land accounting, lifecycle responsibility, community consent, public accountability, and enforceable repair.

The political economy of AI demand is another major concern. AI demand is not simply chosen by users; it is produced by platform defaults, search integration, corporate automation strategies, venture capital, financialized infrastructure, cloud concentration, attention capture, institutional imitation, public subsidy without public control, and geopolitical competition. Efficiency improvements alone will not solve AI's environmental burden if lower costs lead to expanded use. Efficiency must therefore be joined to sufficiency. The goal is not maximum AI use, but appropriate AI use.

To interpret AI's social role, the paper develops a five-part diagnostic framework: AI as Tool, Oracle, Idol, Enclosure, or Commons. AI functions as a Tool when it is bounded, task-appropriate, transparent, and life-serving. It becomes an Oracle when machine output is treated as authority over situated human judgment. It becomes an Idol when society sacrifices energy, water, land, minerals, labor, attention, and ecological stability

to AI expansion. It becomes an Enclosure when compute, data, knowledge, infrastructure, and public life-support are captured by concentrated power. It becomes a Commons only when governed for shared life-capacity within ecological limits.

The paper then proposes a life-coherent AI governance framework organized around ten components: purpose, proportionality, transparency, sufficiency, lifecycle responsibility, place-based accountability, community consent, public-interest compute, knowledge integrity, and review and repair. This governance cycle moves beyond compliance toward stewardship. It asks not only whether AI is safe, efficient, or profitable, but whether it is worthy of the life-ground it consumes.

A dedicated section applies the framework to the Caribbean and Small Island Developing States. For SIDS, AI governance is especially urgent because fragile grids, water constraints, limited land, climate vulnerability, institutional capacity gaps, and external platform dependency can make AI adoption consequential. AI may support disaster preparedness, climate adaptation, water governance, education, health, and regional cooperation. But it may also deepen dependency, value leakage, cultural displacement, and infrastructural burden. The paper therefore argues for a Caribbean AI commons grounded in public-interest compute, regional cooperation, local knowledge, environmental safeguards, and digital sovereignty at the appropriate scale.

The paper concludes with a practical Life-Coherent AI Use Protocol. This protocol asks whether AI is genuinely needed, what life-capacity it expands, whether the smallest adequate model and lightest adequate modality are being used, what environmental and lifecycle burdens are involved, who benefits, who bears the burden, whether dependency is deepened, whether knowledge integrity is protected, whether governance and repair are possible, and whether the use contributes to the commons of life.

The central conclusion is simple: AI should not be judged by symbolic fluency alone. It should be judged by whether it serves life. If AI converts the life-ground into symbolic excess, dependency, and enclosure, restraint is wisdom. If AI expands shared life-capacity within ecological limits, governance is responsibility.

Abbreviations

AI — Artificial Intelligence

CO_{2e} — Carbon dioxide equivalent

GWh — Gigawatt-hour

IEA — International Energy Agency

IPCC — Intergovernmental Panel on Climate Change

kWh — Kilowatt-hour

SIDS — Small Island Developing States

TWh — Terawatt-hour

UNDP — United Nations Development Programme

UN-OHRLLS — United Nations Office of the High Representative for the Least Developed Countries, Landlocked Developing Countries and Small Island Developing States

UNU-INWEH — United Nations University Institute for Water, Environment and Health

Table of Contents

Author Note	3
Acknowledgments	3
AI Transparency Note.....	3
Methodological and Scope Note.....	3
Abstract.....	4
Keywords.....	4
Executive Summary.....	5
Abbreviations.....	6
List of Figures.....	8
List of Tables.....	8
1. Introduction: AI Is Not Immaterial.....	9
2. The UNU-INWEH Report as Empirical Foundation.....	12
3. The Hidden Metabolism of AI.....	15
4. Training, Inference, and Symbolic Escalation.....	19
5. Local Costs, Distant Benefits	25
6. The Political Economy of AI Demand	29
7. AI as Tool, Oracle, Idol, Enclosure, or Commons	32
8. A Life-Coherent Governance Framework for AI	36
9. Caribbean and SIDS Implications	41
10. The Life-Coherent AI Use Protocol	45
11. Conclusion: From Symbolic Power to Responsible Intelligence	48
References.....	50
Appendix A. Life-Coherent AI Use Checklist	52
Appendix B. Figure Source Note.....	52
Appendix C. About the Author.....	52

List of Figures

Figure 1. The Hidden Life-Ground of AI.....	11
Figure 2. AI’s Hidden Metabolism.....	18
Figure 3. Training versus Inference.....	22
Figure 4. The Modality Gradient of Symbolic Demand.....	23
Figure 5. Local Costs, Distant Benefits.....	28
Figure 6. AI as Tool, Oracle, Idol, Enclosure, or Commons.....	34
Figure 7. The Life-Coherent AI Governance Cycle.....	39
Figure 8. SIDS and the AI Dependency Trap.....	44

List of Tables

Table 1. UNU-INWEH Findings and Life-Coherence Translation.....	14
Table 2. AI Task Types and Life-Coherent Use Guidance.....	24
Table 3. Diagnostic Matrix: Tool, Oracle, Idol, Enclosure, Commons.....	35
Table 4. Stakeholder Responsibilities for Life-Coherent AI.....	40
Table 5. Life-Coherent AI Use Checklist.....	47

1. Introduction: AI Is Not Immaterial

Artificial intelligence is commonly encountered as an immaterial event. A question is typed, an answer appears. A prompt is entered, an image emerges. A query is made, a search result is reorganized. Code, summaries, diagnoses, predictions, translations, lessons, voices, videos, and recommendations are generated with a speed that gives the impression of weightlessness. In the interface, AI appears as pure symbolic power.

Yet artificial intelligence is not immaterial. It is not simply code, cognition, or computation floating above the world. It is a physical system embedded in data centers, electricity grids, water systems, land-use decisions, semiconductor supply chains, critical minerals, labor arrangements, thermal management systems, and waste streams (Aczel et al., 2026; Crawford, 2021; Crawford & Joler, 2018). Every generated sentence, image, classification, search enhancement, prediction, or video depends on an underlying life-ground. That life-ground includes energy, water, land, minerals, workers, communities, ecosystems, and ecological sinks.

This white paper begins from that correction.

The rapid rise of generative AI has intensified public attention to questions of bias, privacy, misinformation, labor displacement, intellectual property, surveillance, and governance. These concerns are necessary. But they remain incomplete if AI is treated only as a symbolic, cognitive, or informational technology. The environmental and infrastructural conditions that make AI possible are not secondary. They are constitutive. AI's symbolic outputs are produced by a hidden metabolism that converts material life-support into digital services. To govern AI responsibly, society must make that metabolism visible.

The United Nations University Institute for Water, Environment and Health report, *Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints*, provides an important empirical foundation for this task (Aczel et al., 2026). Its central contribution is to move beyond a carbon-only view of AI sustainability. The report examines how AI's electricity demand translates into carbon, water, and land footprints, and why those footprints vary by location, energy mix, data center siting, model design, task type, and scale of use. This is a crucial advance. A supposedly low-carbon pathway may still be water-intensive. A renewable-energy pathway may still occupy land. A data center may deliver distant economic and strategic benefits while imposing local grid, water, land, and ecological pressures. Single-metric sustainability can therefore conceal burden shifting.

This paper accepts the report's empirical findings and extends them through the life-coherence framework. The question becomes not only how large AI's carbon, water, and land footprints are, but what kind of social order is being built when life-support systems are converted into symbolic output at industrial scale.

The guiding question is: Does artificial intelligence expand life-capacity within ecological limits, or does it convert life-support into symbolic excess, dependency, and enclosure?

Life-capacity refers to the real conditions that enable living beings, communities, and ecosystems to maintain and develop their capacities: health, learning, care, ecological stability, clean water, reliable energy, meaningful work, cultural continuity, public reason, democratic agency, and intergenerational viability (McMurtry, 2009-2011, 2013). A technology is life-coherent when it strengthens these conditions within the limits of the living systems that sustain them. A technology becomes life-incoherent when it consumes, degrades, displaces, or encloses those conditions while presenting its outputs as progress.

Artificial intelligence now stands at this threshold. It may be used as a bounded tool for health, education, ecological monitoring, disaster preparedness, accessibility, scientific discovery, and public-interest coordination. It may help model climate risks, optimize energy systems, support water governance, translate across languages, assist underserved communities, and widen access to knowledge. But AI may also become an idol or enclosure: a system to which societies sacrifice electricity, water, land, minerals, labor, attention, judgment, and ecological stability in pursuit of automation, market dominance, behavioral prediction, synthetic media, and symbolic abundance.

The distinction is not inherent in the technology alone. It depends on governance, ownership, purpose, scale, defaults, siting, energy sourcing, lifecycle responsibility, and the ability of affected communities to participate in decisions. It also depends on whether societies can distinguish useful intelligence from artificial demand, life-serving assistance from symbolic excess, and public-interest compute from platform dependency.

The environmental cost of AI is therefore not merely a technical issue. It is a civilizational diagnostic. It reveals whether digital society can still recognize the living ground beneath its symbols. It asks whether intelligence can be subordinated to life, or whether life will be reorganized around the expansion of machine-generated symbols.

This paper develops that diagnostic in eleven movements. First, it summarizes the UNU-INWEH report as an empirical foundation for understanding AI's carbon, water, and land footprints. Second, it interprets AI as a hidden metabolic system linking prompts, models, data centers, electricity, cooling, minerals, labor, and waste. Third, it distinguishes between training and inference, showing why everyday use at scale may become the dominant operational burden. Fourth, it examines symbolic escalation across task types, from lightweight classification to text generation, image generation, and high-complexity video. Fifth, it analyzes the justice problem of local costs and distant benefits. Sixth, it examines the political economy that manufactures AI demand through defaults, platform integration, automation strategies, financial incentives, and competitive arms races. Seventh, it integrates these findings into the framework of AI as Tool, Oracle, Idol, Enclosure, or Commons. Eighth, it proposes a life-coherent governance framework. Ninth, it considers implications for the Caribbean and Small Island Developing States. Tenth, it offers a Life-Coherent AI Use Protocol. Finally, it concludes by arguing that AI must be judged not by symbolic fluency alone, but by whether it expands shared life-capacity within ecological, social, and moral limits.

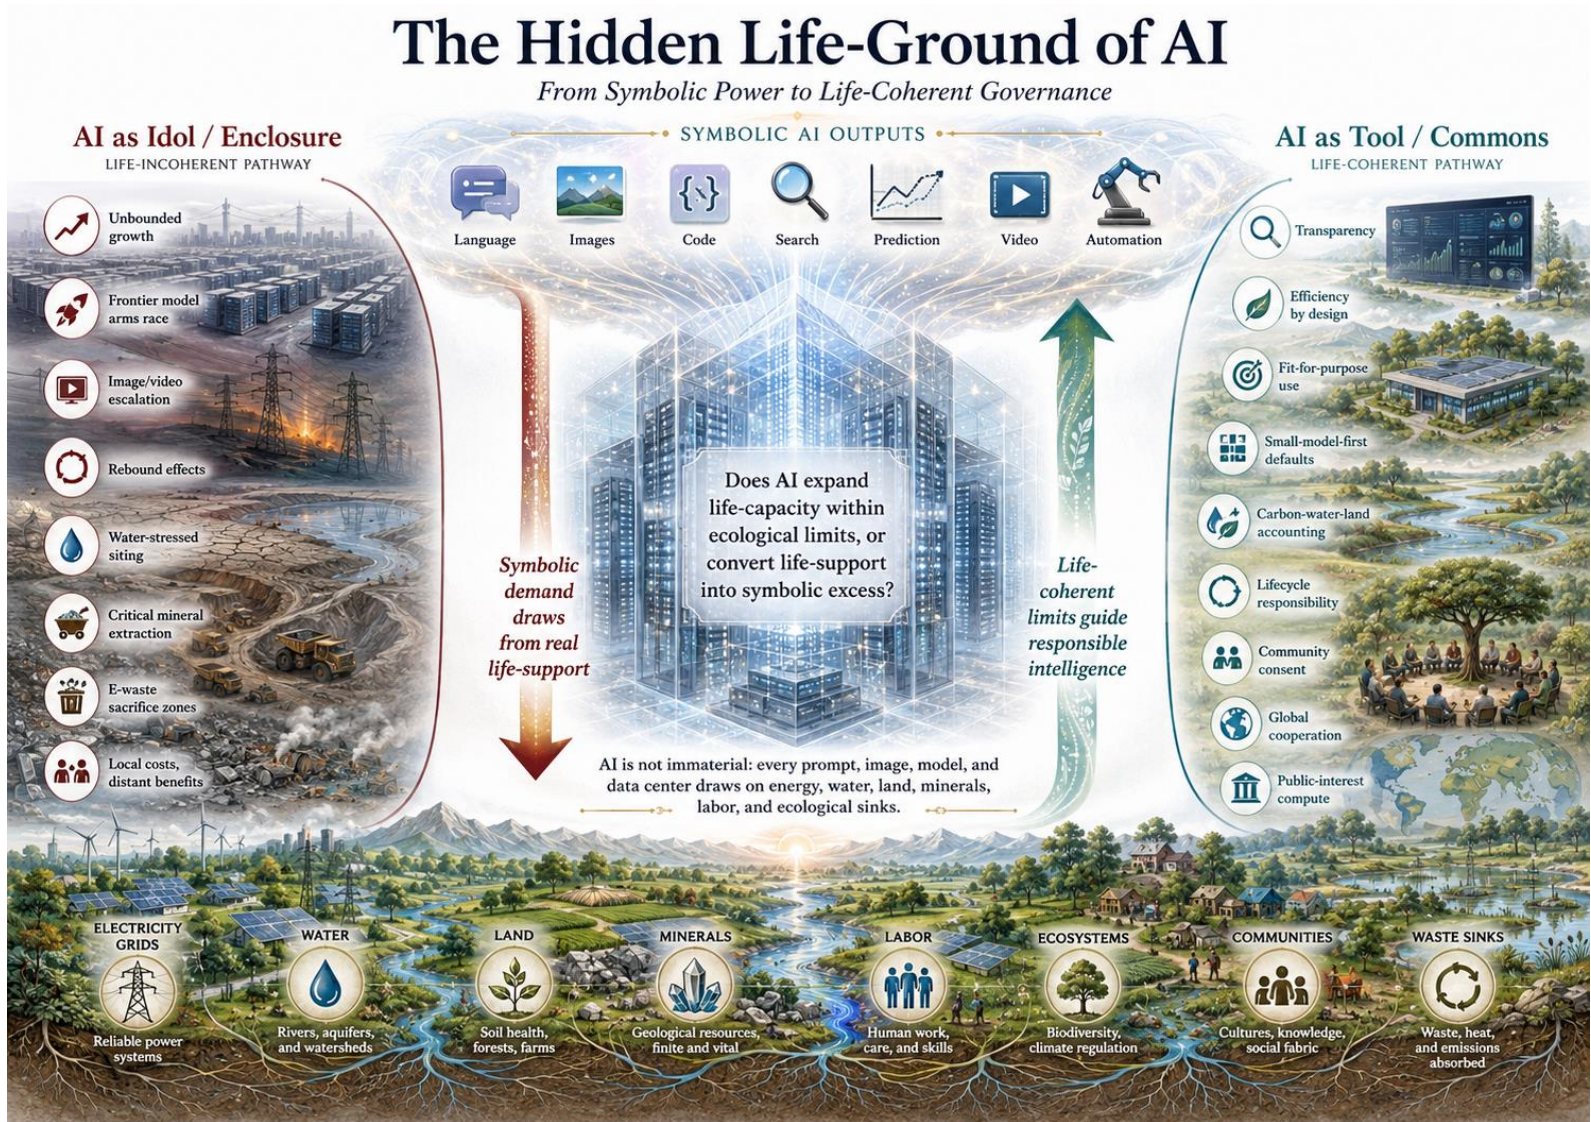


Figure 1. The Hidden Life-Ground of AI.

This master diagram shows artificial intelligence as a symbolic-output system grounded in material life-support. Language, images, code, search, prediction, video, and automation appear as visible AI outputs, while data centers, compute systems, electricity grids, water, land, minerals, labor, ecosystems, communities, and waste sinks form the hidden life-ground beneath them. The left side depicts the life-incoherent pathway of AI as idol and enclosure; the right side depicts the life-coherent pathway of AI as tool and commons. The central diagnostic question asks whether AI expands life-capacity within ecological limits or converts life-support into symbolic excess.

2. The UNU-INWEH Report as Empirical Foundation

The UNU-INWEH report provides a timely and necessary empirical correction to the dominant imagination of AI. It argues that artificial intelligence cannot be understood only as a software layer or economic sector. It must be understood as a material infrastructure with measurable environmental footprints. Its subtitle, Carbon, Water and Land Footprints, is therefore not incidental. It signals a shift from symbolic fascination to grounded accounting (Aczel et al., 2026).

The report begins by situating AI within the Fourth Industrial Revolution. Artificial intelligence is described as a transformative technology embedded across sectors: finance, healthcare, energy management, transportation, translation, recommendations, customer service, software development, medical imaging, diagnostics, environmental monitoring, and more. Its adoption has been extraordinarily rapid. Generative AI tools moved from laboratories into everyday life at unprecedented speed, producing text, images, code, music, voice, and video from simple prompts. This adoption has created a new infrastructure problem: the symbolic outputs visible to users are supported by physical systems whose environmental costs are not equally visible.

The report's first major contribution is to name AI's materiality. Artificial intelligence is not "just code." It involves data centers, chips, electricity generation, cooling systems, water withdrawals, land occupation, critical minerals, and eventual e-waste. This statement provides the bridge between environmental accounting and life-coherence analysis. The apparent immateriality of AI conceals a chain of dependence on life-support systems. The report makes this chain measurable.

The second contribution is methodological. The report moves beyond a carbon-only lens. Many digital sustainability discussions focus on greenhouse gas emissions, but the UNU-INWEH report examines carbon, water, and land footprints together. This matters because environmental burdens are multidimensional. A power source or data center strategy that reduces one footprint may worsen another. Low-carbon electricity can still involve high water use. Renewable energy infrastructure can have land implications. A region selected for cheap power may already face water stress or grid pressure. Sustainability accounting that measures only carbon can therefore hide trade-offs and redistribute harms onto communities with less political power.

The third contribution is scale. The report estimates that global data centers consumed 448 terawatt-hours of electricity in 2025. By 2030, data center electricity demand could rise to approximately 945 terawatt-hours, approaching 3 percent of projected global electricity use, a projection also supported by the International Energy Agency's broader energy and AI analysis (Aczel et al., 2026; International Energy Agency, 2025). Earlier recalibrations of global data center energy use also show why careful, evidence-based estimation is necessary when assessing digital infrastructure, since rapid efficiency improvements and demand growth can shift projections over time (Masanet et al., 2020). The associated footprints would be enormous: hundreds of millions of tonnes of carbon dioxide equivalent, trillions of liters of water, and thousands of square kilometers of land footprint. These figures reframe AI from a niche digital concern into a major infrastructure and resource-governance issue.

AI is not the only driver of data center growth, but it is becoming one of the most significant. The report estimates that AI workloads accounted for around 20 percent of total data center electricity use in 2025 and could reach 40 percent by 2030. This shift is important because AI demand is not merely added to existing digital infrastructure; it is reshaping infrastructure planning itself. Data center siting, power procurement, grid expansion, water use, cooling design, hardware turnover, and investment strategy are increasingly influenced by AI workloads.

The fourth contribution is the distinction between training and inference. Public debate often focuses on the dramatic cost of training frontier models. The report notes that GPT-3 training consumed an estimated 1.3 gigawatt-hours, while GPT-4 training is estimated at 50 to 70 gigawatt-hours. These figures are striking, especially when compared with electricity use in energy-poor regions. But the report also emphasizes that

training is only part of the picture. Once models are deployed, the continuous inference phase — everyday use by millions or billions of users — accounts for the majority of AI energy consumption. Inference may represent 80 to 90 percent of total AI energy use (Aczel et al., 2026).

This distinction is central for governance. A society could regulate frontier model training and still fail to address the operational burden of deployed systems. The everyday interface becomes the site where environmental costs accumulate: prompts, summaries, searches, chat interactions, generated images, synthetic voices, videos, embedded assistants, enterprise tools, and automated workflows. Small per-use impacts become infrastructure-level loads when multiplied by billions of interactions.

The fifth contribution is task differentiation. The report shows that “AI use” is not a single category. Different tasks have radically different energy implications. Basic text classification, such as spam filtering, is far lighter than a conversational generative response. Long responses are more intensive than short responses. Image generation is far more energy-intensive than lightweight text tasks. High-complexity video generation represents an even more demanding frontier. Product defaults, user prompts, model choice, output length, resolution, modality, and platform design all influence environmental footprints.

The sixth contribution is the recognition of rebound effects. Efficiency improvements matter, but they do not guarantee reduced total impact. If AI systems become cheaper, faster, and more efficient per task, overall use may expand so rapidly that total energy, water, and land impacts continue to rise. This is the familiar Jevons dynamic: efficiency gains can become growth accelerants when embedded in a system organized around expansion. The report’s emphasis on rebound effects helps prevent a narrow technological optimism in which efficiency alone is treated as sufficient.

The seventh contribution is lifecycle awareness. Although the report focuses primarily on the environmental footprints of AI’s energy use, it also acknowledges the broader lifecycle impacts of AI hardware. Chips, servers, cooling systems, batteries, and supporting infrastructure depend on critical minerals and manufacturing processes with environmental and social consequences. Extraction and processing often occur in regions with weaker oversight, and e-waste can expose frontline communities to hazardous substances. By 2030, AI infrastructure could generate millions of metric tons of e-waste annually. This points toward the need for full value-chain governance, from mineral sourcing to manufacturing, deployment, maintenance, reuse, recycling, and safe disposal (Aczel et al., 2026; Baldé et al., 2024).

The eighth contribution is governance. The report proposes six operational principles for a responsible AI ecosystem: transparency, efficiency by design, equity and environmental justice, lifecycle responsibility, global cooperation, and sustainable use. These principles move the conversation from impact measurement to institutional responsibility. From a life-coherence perspective, these principles are necessary but not yet sufficient. They make AI’s environmental costs visible, comparable, and actionable. But the deeper question remains: visible for what purpose? Measurement must lead to boundary, repair, and accountability. Transparency must not become another dashboard that legitimizes expansion. Efficiency must be joined to sufficiency. Lifecycle responsibility must include enforceable duties. Global cooperation must prevent the creation of digital sacrifice zones. Sustainable use must distinguish between life-serving need and artificially stimulated demand.

The report therefore serves as both empirical foundation and threshold document. It establishes that AI’s physical impacts can be measured. It shows that carbon, water, and land must be assessed together. It demonstrates that training and inference must be distinguished. It warns that task type and modality matter. It identifies rebound effects. It recognizes lifecycle harms. It proposes a responsible governance architecture. This white paper builds from that foundation by asking what kind of intelligence remains coherent with life.

Table 1. UNU-INWEH Findings and Life-Coherence Translation

UNU-INWEH finding	Life-coherence translation	Governance implication
AI depends on data centers, chips, electricity, cooling, water, land, minerals, and e-waste systems.	AI has a hidden life-ground.	Govern AI as material infrastructure, not only software.
Carbon-only accounting is insufficient.	Environmental burden is multidimensional.	Require carbon, water, and land accounting together.
Data center electricity demand is rapidly growing.	Symbolic output is becoming infrastructure-scale demand.	Integrate AI into energy planning and grid governance.
AI workloads are an increasing share of data center energy use.	AI is becoming a major driver of digital metabolism.	Require AI-specific footprint disclosure and planning.
Training frontier models is highly energy-intensive.	Model ambition carries life-ground cost.	Scrutinize frontier model development and resource budgets.
Inference may account for most AI energy use once systems are deployed.	Everyday use at scale becomes cumulative burden.	Govern defaults, routing, prompts, output length, and institutional adoption.
Image and video generation are more energy-intensive than lightweight text tasks.	Symbolic escalation increases life-ground demand.	Prefer minimum sufficient modality.
Efficiency gains may produce rebound effects.	Efficiency without sufficiency becomes acceleration.	Pair efficiency standards with demand governance.
Hardware lifecycles create mineral and e-waste burdens.	AI has a body from mine to waste.	Require lifecycle responsibility and extended producer accountability.
Benefits and environmental costs are unevenly distributed.	Local costs and distant benefits reveal life-incoherence.	Require place-based accountability, community consent, and fair benefit.

3. The Hidden Metabolism of AI

If artificial intelligence is not immaterial, then it must be understood metabolically. It receives inputs, transforms them through physical processes, and produces outputs while drawing on energy, materials, water, land, labor, and ecological sinks. In the visible interface, the user sees a prompt and a response. In the hidden infrastructure, the prompt activates a chain of computation, electricity demand, cooling, hardware use, data transmission, and environmental consequence.

This is the hidden metabolism of AI.

A metabolism is not merely a metaphor. Living organisms maintain themselves by transforming energy and materials through organized processes. Economies also have metabolisms: they extract, process, transport, consume, discard, and absorb waste through ecological systems. Artificial intelligence has now become part of this wider socio-technical metabolism. It converts material life-support into symbolic outputs: sentences, images, predictions, classifications, summaries, code, recommendations, voices, and videos.

The life-coherence framework begins by asking what is being converted into what, who benefits from the conversion, who bears its costs, and whether the process expands or diminishes life-capacity.

In the case of AI, the conversion chain can be described as follows: Prompt -> model -> data center -> electricity -> cooling -> carbon, water, and land footprints -> minerals -> labor -> waste -> ecological sink.

The prompt is the visible beginning, but not the real beginning. Before the prompt can be processed, there must already be trained models, manufactured chips, server racks, cooling systems, data center buildings, fiber networks, software stacks, electricity contracts, water arrangements, land-use approvals, supply chains, financial investments, and legal permissions. The apparent immediacy of AI output depends on a long prior history of extraction, design, construction, training, deployment, and maintenance.

The first stage of the hidden metabolism is model production. A model is not only an abstract mathematical artifact. It is the outcome of data collection, computational training, hardware capacity, software engineering, and institutional investment. Large models require specialized chips and high-performance computing systems. Their development depends on laboratories, cloud platforms, proprietary infrastructure, open-source communities, workers, data labelers, engineers, and energy-intensive training processes. Even before a user asks a question, the model has already absorbed environmental and social costs.

The second stage is deployment infrastructure. Once trained, AI systems must be hosted and made available. This requires data centers capable of running continuous inference. Data centers are the material backbone of AI. They house servers, processors, storage systems, networking equipment, cooling infrastructure, backup power systems, transformers, cables, and physical security systems. They require land. They require electricity. They produce heat. They often require water for cooling or rely on systems whose energy supply has water and land implications. They are therefore not neutral containers for digital activity. They are infrastructural organisms embedded in local energy, water, land, and community systems.

The third stage is electricity demand. Every AI interaction consumes energy, even if the amount per interaction varies enormously. The environmental meaning of this electricity depends on the grid supplying it. A fossil-heavy grid carries higher carbon intensity. A renewable-heavy grid may reduce emissions but may still involve land occupation, water trade-offs, mineral requirements, storage systems, transmission infrastructure, and intermittency challenges. This is why the UNU-INWEH report's carbon-water-land framing is essential. A technology cannot be considered sustainable simply because one environmental dimension looks favorable. The whole footprint must be examined together.

The fourth stage is cooling and thermal management. Computation produces heat. Dense AI workloads intensify this thermal burden. Servers must be cooled to operate safely and reliably. Cooling systems can draw directly on local water supplies, especially in high-density facilities, or indirectly on electricity systems whose generation requires water. In water-stressed regions, data centers can therefore become competitors for

aquifers, rivers, reservoirs, and municipal water systems. The question is not only how much water is used, but when, where, under what climatic conditions, and at whose expense (Li et al., 2023).

The fifth stage is land occupation and spatial transformation. Data centers occupy physical sites. Energy infrastructure occupies land. Transmission corridors, substations, solar farms, wind farms, mineral extraction zones, roads, and waste facilities all produce spatial consequences. Land is not an empty surface available for technical deployment. It is habitat, watershed, farmland, community space, cultural memory, and ecological function. When AI infrastructure expands, it reorganizes land into the service of symbolic production. The life-coherence question is whether that reorganization strengthens or weakens the life-capacities of the places affected.

The sixth stage is mineral and hardware dependence. AI relies on processors, memory, storage, networking equipment, cooling systems, batteries, and power infrastructure. These require copper, silicon, lithium, cobalt, rare earths, and other materials. Mining and processing can be energy-intensive, water-intensive, land-disrupting, and toxic. They can also be socially harmful when carried out under weak labor protections, poor environmental oversight, or extractive ownership structures. The clean interface of AI may therefore be linked to distant mines, polluted waters, unsafe labor, and communities that receive little benefit from the value generated downstream (Crawford, 2021; Crawford & Joler, 2018).

The seventh stage is labor. AI is frequently presented as automation beyond human work, but its hidden metabolism includes many forms of labor: miners, construction workers, electricians, engineers, data center technicians, chip fabrication workers, data labelers, content moderators, logistics workers, energy workers, maintenance crews, researchers, trainers, and users whose interactions generate data and refine systems. The automation narrative conceals this distributed human substrate. AI does not eliminate labor; it reorganizes, displaces, obscures, and often devalues it.

The eighth stage is waste and end-of-life burden. AI hardware has a physical lifecycle. Chips, servers, cables, batteries, storage devices, and cooling systems age, fail, or become obsolete. Competitive pressure toward more powerful models and faster hardware turnover intensifies e-waste risks. If not responsibly reused, repaired, recycled, or safely disposed of, obsolete infrastructure can contribute to toxic waste streams. Heavy metals and hazardous components can contaminate soil and water, often in communities far from the centers of profit and decision-making. End-of-life responsibility must therefore be part of AI governance from the beginning, not an afterthought (Baldé et al., 2024).

The ninth stage is the ecological sink. Carbon emissions must be absorbed by atmospheric, terrestrial, and oceanic systems already under stress. Waste heat must be dissipated. Pollutants must be diluted or contained. Water withdrawals must be recovered through hydrological cycles. Land disturbance must be compensated by ecological resilience that may or may not be restored. The living world becomes the sink for the consequences of symbolic production. When those sinks are overwhelmed, the result is not merely environmental degradation. It is life-capacity loss.

The hidden metabolism of AI therefore reveals an important inversion. AI is usually valued for its outputs: answers, images, automation, prediction, speed, personalization, convenience, and productivity. But life-coherent evaluation begins with the inputs and consequences. What forms of life-support are being consumed? What burdens are being displaced? What dependencies are being created? What capacities are being strengthened or weakened? What ecological limits are being crossed? What communities are being asked to absorb the costs?

This does not mean AI should be rejected. It means AI must be situated. A life-serving medical diagnostic tool, a climate-risk model, an accessible translation system, a water-management assistant, or an educational support system may justify material expenditure if it expands life-capacity and is governed within limits. But an endless stream of synthetic content, unnecessary high-resolution generation, manipulative personalization, speculative automation, or platform-driven symbolic excess may consume the same life-ground while producing little or no real life-value.

The distinction is essential. The problem is not computation as such. The problem is unbounded symbolic production detached from life-serving purpose.

AI's hidden metabolism also forces a redefinition of efficiency. Efficiency is often understood as reducing energy use per operation. That is necessary but incomplete. A system can become more efficient per task while total use grows so rapidly that overall burdens increase. A model can generate images more efficiently while platforms encourage vastly more image generation. A chatbot can reduce energy per response while being embedded into every search, document, device, and workflow. A data center can improve cooling efficiency while total demand expands beyond local resource limits. Efficiency without sufficiency becomes acceleration.

Life-coherent efficiency must therefore be bounded by purpose. The question is not simply whether AI can do more with less. The question is whether the "more" is needed, whether it serves life-capacity, and whether the "less" remains within ecological and community limits. Every symbolic output should be justified by its contribution to life-capacity.

AI's Hidden Metabolism

From Prompt to Ecological Sink

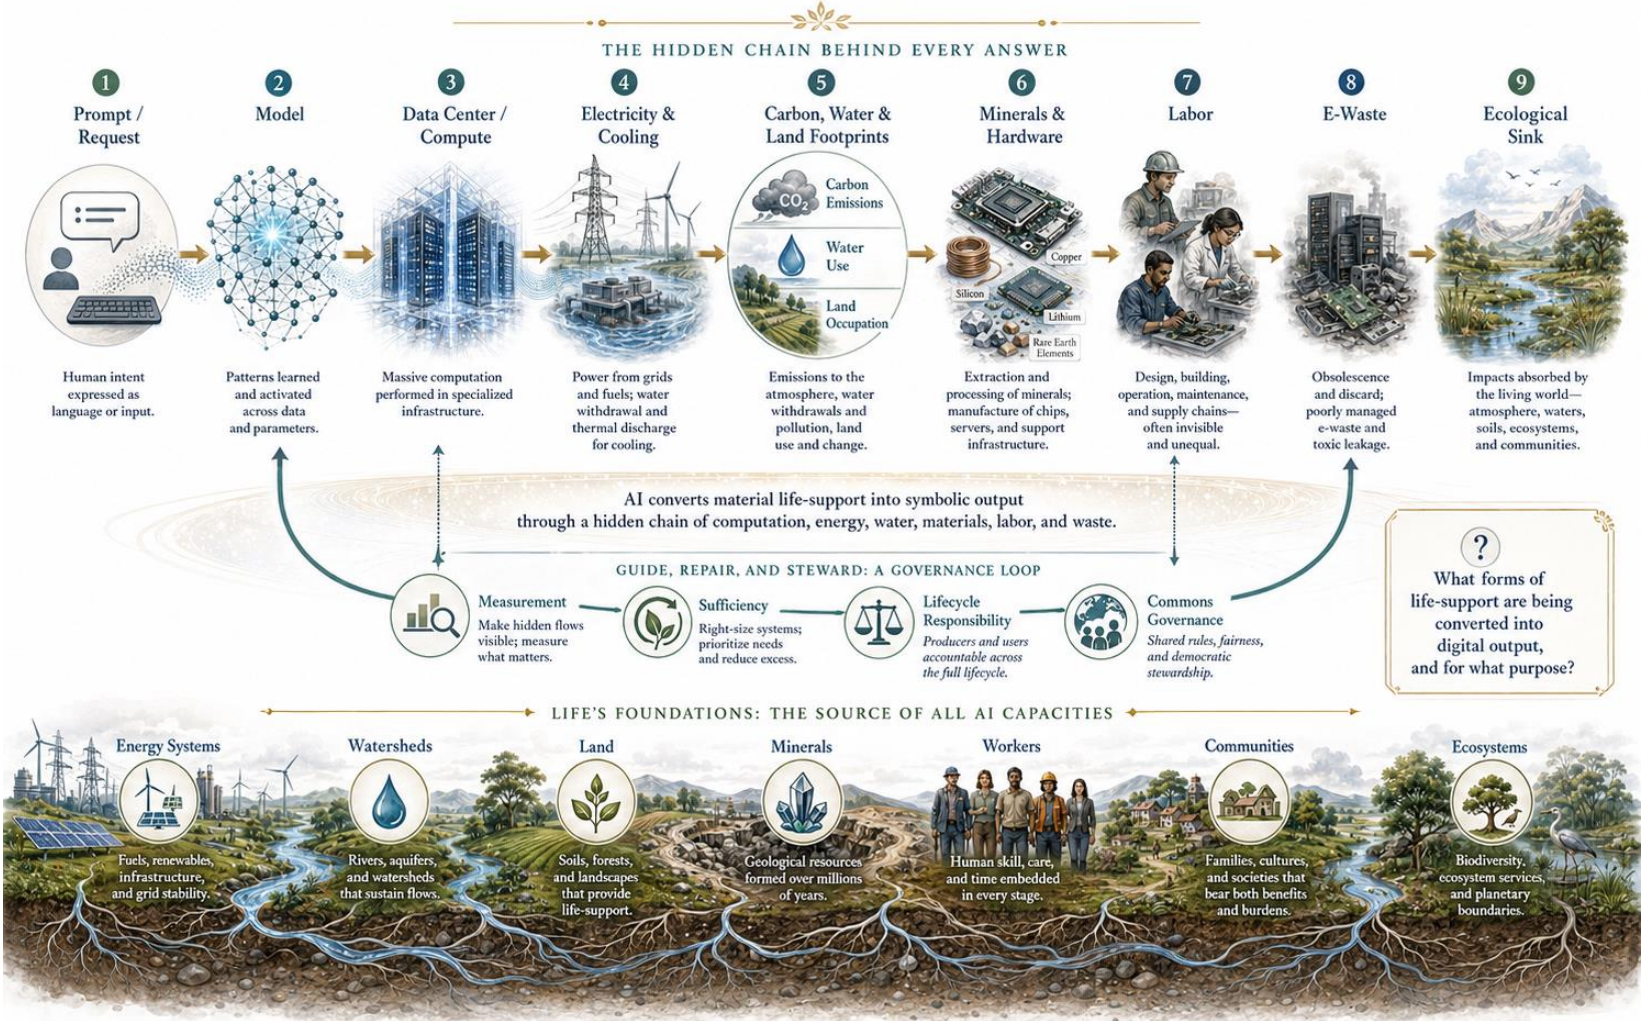


Figure 2. AI's Hidden Metabolism

This diagram presents AI as a metabolic conversion system. A visible prompt activates a hidden chain of model inference, data center processing, electricity use, cooling demand, carbon-water-land footprints, mineral and hardware dependence, labor, e-waste, and ecological sinks. The governance loop below the chain emphasizes measurement, sufficiency, lifecycle responsibility, and commons governance as necessary ways of making AI's material life-ground visible and accountable.

4. Training, Inference, and Symbolic Escalation

The environmental discussion of artificial intelligence often begins with training. This is understandable. Training frontier models requires enormous computational effort. Early work on the energy and policy implications of deep learning helped establish that model training can carry significant computational and environmental costs, especially when large models are treated as the default route to progress (Strubell et al., 2019). It concentrates energy demand into dramatic episodes of model development, often involving specialized chips, high-performance computing clusters, long training runs, vast datasets, and teams of engineers working across cloud infrastructure. These training events are visible as milestones in the technological race: one model surpasses another, one company announces a new frontier system, one platform claims better reasoning, larger context windows, multimodal capacity, faster inference, or more human-like interaction.

Training is therefore the spectacle of AI development. But training is not the whole metabolism. Once a model is released, the center of environmental gravity begins to shift from the one-time cost of training to the ongoing cost of use. This is inference: the process by which a trained model generates outputs in response to user prompts, API calls, search queries, enterprise workflows, embedded tools, and automated systems. Inference is less spectacular than training, but it may become far more consequential because it repeats continuously across millions or billions of interactions (Aczel et al., 2026; Luccioni et al., 2023; Patterson et al., 2021).

The distinction between training and inference is essential for life-coherent governance. Training asks: What did it cost to build the model? Inference asks: What does it cost to keep using the model at scale? Training is episodic. Inference is continuous. Training is concentrated. Inference is distributed. Training is often controlled by developers. Inference is shaped by platforms, institutions, defaults, users, business models, and cultural habits.

The UNU-INWEH report makes this distinction clear. It notes that frontier training runs can require vast amounts of electricity, with later models requiring many times more energy than earlier generations. Yet it also emphasizes that once systems are deployed, the everyday operational phase may account for the majority of total AI energy use. This changes the governance problem. It is not enough to scrutinize large training runs while leaving daily deployment unchecked. The cumulative burden of ordinary use must also be governed.

This is especially important because AI adoption is no longer confined to specialized users. Generative AI is being inserted into search engines, document platforms, schoolwork, customer service, health systems, coding environments, office software, design tools, advertising systems, personal assistants, smartphones, surveillance systems, and public administration. What begins as a tool becomes an ambient layer. The user may not even know when AI has been added to an interaction. When AI becomes default infrastructure, inference becomes background consumption.

This produces a new form of environmental invisibility. The user does not see the data center. The user does not see the cooling load. The user does not see the grid congestion, water demand, hardware turnover, or land footprint. The user sees a smoother interface, a generated summary, a recommended reply, a search result rewritten by AI, or an image created in seconds. The environmental consequence is separated from the symbolic convenience.

Life-coherent analysis therefore requires attention not only to model size but also to use pattern. A highly capable model used sparingly for important tasks may be easier to justify than a slightly more efficient model deployed everywhere by default for trivial purposes. The central issue is not only energy per operation. It is the multiplication of operations by design, habit, and institutional dependence.

This is where the problem of symbolic escalation emerges. Symbolic escalation occurs when increasingly resource-intensive forms of AI output become normalized, not because they are necessary, but because they

are available, impressive, profitable, or encouraged by platforms. A user who needs a brief factual answer may receive a long generative response. A simple document search may become an AI-enhanced synthesis. A short explanation may become an illustrated presentation. A still image may become an animated video. A human conversation may become a synthetic interaction. A basic automation may become an always-on assistant. Each escalation may appear minor at the level of a single user. At scale, it becomes infrastructural demand.

Not all AI tasks are environmentally equivalent. Lightweight classification, such as spam detection or basic sorting, has a very different footprint from long-form text generation. Image generation requires more computation than a short text response. High-complexity video generation is more demanding still. The difference is not merely quantitative; it is a governance signal. Modality, output length, resolution, number of iterations, model size, and default settings all shape environmental impact.

A life-coherent AI system would therefore be organized by a principle of minimum sufficient symbolic form. If classification is enough, do not use generative text. If a short answer is enough, do not generate a long response. If text is enough, do not generate an image. If an image is enough, do not generate video. If a smaller model is enough, do not route the task to a frontier model. If a human conversation, local knowledge, or existing document is enough, do not force AI mediation. This principle is not anti-creativity. It is proportionality. It restores the relationship between symbolic output and life-ground cost. It asks that high-intensity forms of AI be justified by corresponding life-value.

The report's task-differentiation findings are especially important because they reveal that environmental responsibility is built into design choices. A company that defaults to long responses increases demand. A platform that routes simple tasks to large models increases demand. A search engine that adds generative summaries to routine queries increases demand. A school that encourages image and video generation for ordinary assignments increases demand. A marketing department that uses generative video for marginal engagement increases demand. A public institution that adopts AI without workload discipline increases demand.

Conversely, responsible design can reduce demand without rejecting AI. A platform can default to concise answers. A system can use smaller models for simple tasks. An interface can ask whether image or video generation is necessary. An enterprise can reserve high-capacity models for high-value tasks. A school can teach students to distinguish between meaningful use and symbolic excess. A government can require environmental disclosure for AI procurement. A data center can be evaluated not only by efficiency but by local water, land, and grid impacts.

Useful intelligence helps living systems understand, coordinate, heal, learn, repair, adapt, and flourish. It is bounded by need and accountable to consequences. Symbolic excess produces more output without corresponding life-value. It expands images, text, summaries, predictions, simulations, and automation because the system is designed to generate, monetize, or normalize them. It treats fluency as value, novelty as progress, and availability as justification. It converts life-support into surplus symbols.

This disciplined judgment can be expressed through four tests. First, the necessity test: Is AI needed for this task, or is it being used because it is available? Second, the proportionality test: Is the chosen model and modality proportionate to the task? Third, the life-value test: Does the output expand real life-capacity, such as learning, care, access, repair, safety, ecological understanding, or democratic agency? Fourth, the burden test: Where do the energy, water, land, hardware, labor, and waste burdens fall, and are they justified, disclosed, minimized, and fairly governed?

The training-inference distinction also reveals a danger in current AI governance discourse. Many companies emphasize efficiency improvements: better chips, better cooling, better models, better routing, better energy procurement. These improvements are necessary. But if they make AI cheaper and more ubiquitous, total resource use may continue to rise. Efficiency gains can be captured by expansion. This is the rebound problem. Life-coherent governance must therefore pair efficiency with sufficiency. Efficiency asks how to reduce the footprint per task. Sufficiency asks how many tasks, of what kind, for what purpose, are enough.

Proportionate AI is bounded by need. It is routed through the smallest adequate model. It favors lower-footprint modalities unless higher ones are justified. It makes environmental costs visible at the level of design, procurement, and use. It resists artificial demand. It expands life-capacity without normalizing symbolic excess. In this sense, the environmental cost of AI is also a measure of cultural discipline.

Training versus Inference

From Spectacular Model Building to Continuous Use at Scale

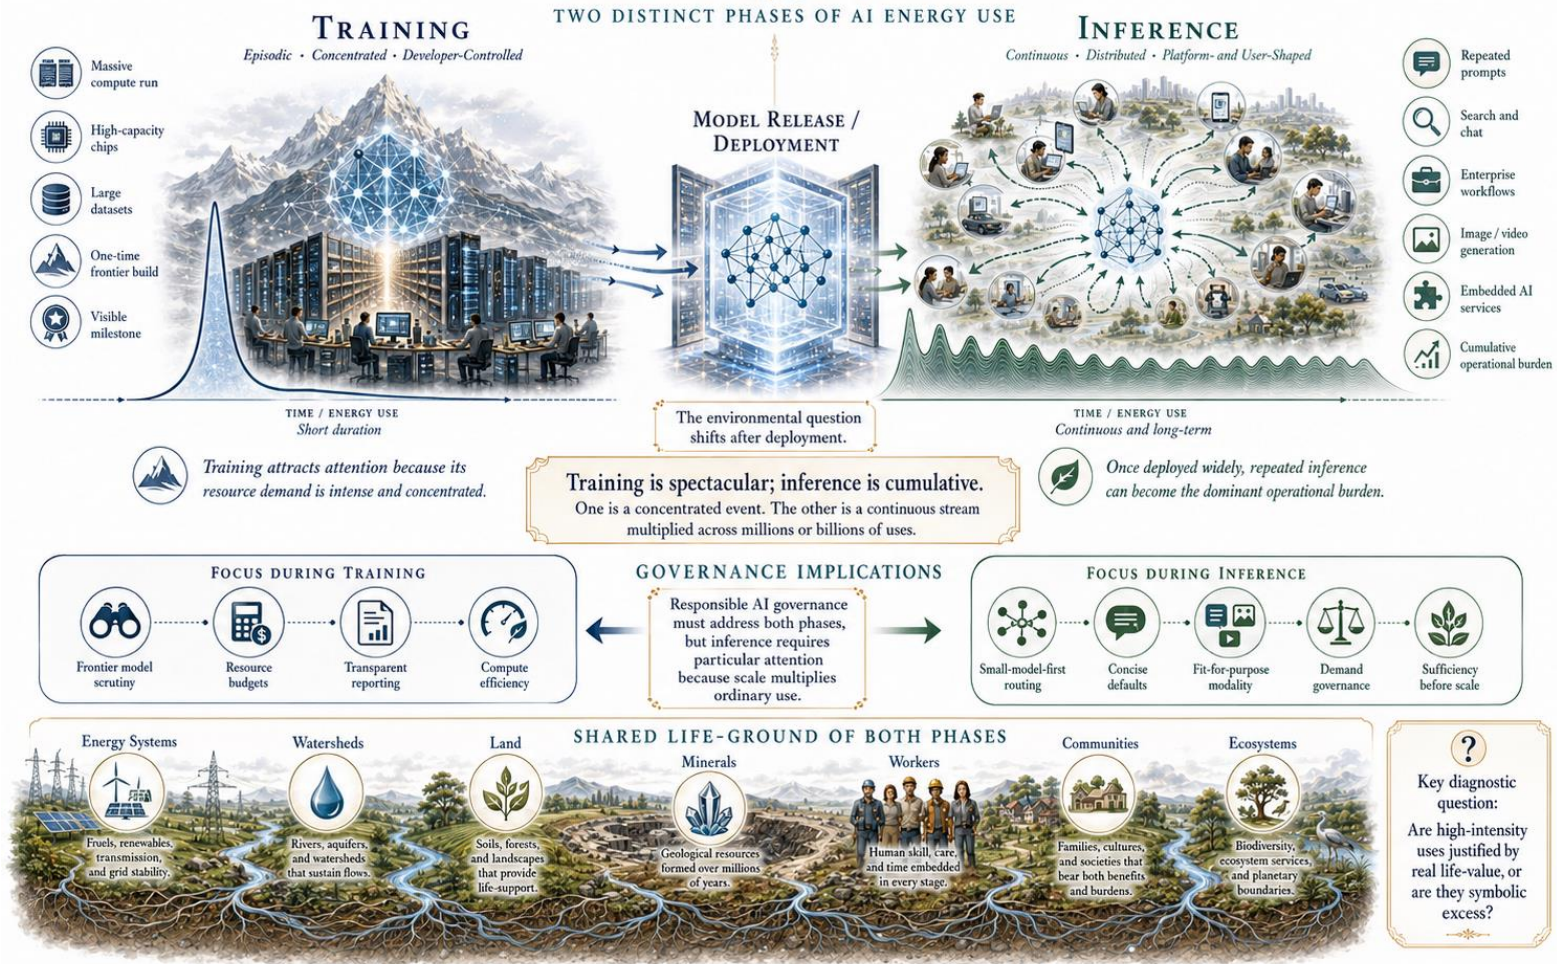


Figure 3. Training versus Inference.

This diagram distinguishes the episodic burden of model training from the continuous operational burden of inference. Training is shown as an intense, concentrated, developer-controlled phase, while inference is shown as a distributed, repeated, platform- and user-shaped phase that accumulates across prompts, searches, summaries, images, videos, and embedded AI services. The diagram highlights why AI governance must address both spectacular model-building and ordinary use at scale.

The Modality Gradient of Symbolic Demand

From Minimum Sufficient Form to High-Intensity Symbolic Output

NOT ALL AI OUTPUTS CARRY THE SAME LIFE-GROUND COST




Figure 4. The Modality Gradient of Symbolic Demand.

This diagram shows the increasing resource intensity of AI outputs across modalities, from classification and short text to long text, image generation, and video generation. It illustrates the principle that not all AI use carries the same life-ground cost. The central guidance is to use the minimum sufficient symbolic form: if classification is enough, do not generate text; if text is enough, do not generate images; if images are enough, do not generate video.

Table 2. AI Task Types and Life-Coherent Use Guidance

AI task type	Typical life-value potential	Risk of symbolic excess	Life-coherent guidance
Basic classification	Sorting, filtering, triage, detection	Low if task is clear and bounded	Use lightweight models or rules where sufficient.
Short text response	Clarification, explanation, accessibility	Moderate if overused for trivial queries	Prefer concise answers and small models where adequate.
Long text generation	Drafting, synthesis, education, research support	High if used to replace thinking or produce surplus content	Use when depth is justified; preserve human judgment and review.
Image generation	Education, communication, visualization, accessibility	High if used decoratively or repeatedly	Use only when visual form adds real value.
Video generation	Training, accessibility, public communication	Very high if used for spectacle or marketing excess	Reserve for high-value cases where motion is necessary.
AI-enhanced search	Synthesis, complex query support	High if applied to all routine searches	Use selectively; simple lookup should remain lightweight.
Automation	Administrative relief, workflow support	High if used for labor displacement or surveillance	Ensure human accountability and public value.
Prediction and scoring	Risk detection, planning, early warning	High in rights-affecting contexts	Require auditability, appeal, bias review, and human oversight.
Health AI	Diagnosis support, triage, education, monitoring	High if replacing clinical judgment	Use as support tool only; preserve clinician responsibility.
Education AI	Tutoring, feedback, accessibility	High if bypassing learning	Use to deepen learning, not replace attention or effort.
Climate and environmental AI	Monitoring, forecasting, adaptation	Low to moderate when public-interest aligned	Prioritize as high life-value use when transparent and locally governed.

5. Local Costs, Distant Benefits

Artificial intelligence is often consumed at a distance from its consequences. The user may be in one country, the data center in another, the minerals extracted from another, the hardware manufactured in another, the electricity generated from a regional grid, the water withdrawn from a local watershed, and the waste eventually discarded in yet another place. The output appears in the interface as if it belongs only to the user. The costs are distributed through infrastructures and ecosystems that may remain invisible to that user.

This separation creates one of the central justice problems of AI: local costs and distant benefits.

The benefits of AI are often captured by platform companies, cloud providers, investors, powerful states, high-income users, enterprise clients, advertisers, military and security institutions, and organizations able to integrate AI into their operations. These benefits include profit, productivity, strategic advantage, automation, market power, behavioral prediction, knowledge extraction, and geopolitical leverage. The burdens, however, may fall elsewhere: on communities near data centers, workers in supply chains, regions with water stress, landscapes disrupted by energy and mineral infrastructure, electricity grids under strain, and communities receiving toxic e-waste.

This is not an accidental side issue. It is a structural feature of global digital infrastructure. The AI interface detaches action from consequence. A generated image does not show the cooling system. A chatbot response does not show the power plant. A search enhancement does not show the transmission line. A synthetic video does not show the water basin. A frontier model announcement does not show the mine, the fabrication plant, the data center, the labor chain, or the waste stream. Symbolic convenience is spatially separated from material accountability.

The UNU-INWEH report helps make this separation visible by emphasizing that AI's environmental footprints vary by location. Carbon, water, and land impacts are not abstract global averages alone. They depend on electricity supply mixes, data center siting, cooling strategies, water availability, land-use context, grid capacity, and regional governance (Aczel et al., 2026; Dodge et al., 2022; Lacoste et al., 2019; Li et al., 2023). A data center placed in one location may impose different burdens than the same facility placed elsewhere. A low-carbon site may not be low-water. A region with available land may still contain fragile ecosystems or communities with limited political power. A cheap energy jurisdiction may externalize costs through fossil generation, water withdrawals, or public infrastructure subsidies.

This means AI sustainability cannot be evaluated only from the perspective of the user, the firm, or the global emissions ledger. It must be evaluated from the perspective of affected places. A life-coherent framework begins with the question: Where does the burden land?

The first burden is grid competition. Data centers require stable, high-volume electricity. As AI workloads expand, data centers can become major new demands on local and regional grids. This may require new generation capacity, transmission upgrades, substations, backup systems, or altered energy planning. In some cases, data center demand may compete with residential needs, industrial uses, public services, electrified transport, climate adaptation, or renewable energy transition goals. A project that appears economically beneficial may, in practice, redirect public infrastructure toward private computational demand.

The second burden is water stress. AI computation produces heat, and heat must be managed. Data centers may use water directly for cooling or indirectly through electricity generation. In water-stressed regions, this can intensify competition among households, farmers, ecosystems, municipalities, and industry. Water is not merely an input. It is a life-condition. When cooling requirements are treated as technical needs without democratic scrutiny, communities may be asked to subsidize symbolic production through their watersheds.

The third burden is land transformation. Data centers occupy land, but their land footprint extends beyond buildings. Energy generation, transmission corridors, substations, roads, cooling infrastructure, warehouses, mineral extraction, and waste systems all reorganize territory. Land is habitat, soil, watershed, food base,

cultural memory, and community space. When land is converted to serve distant compute demand, the affected place bears a transformation that may not be reflected in the price of the AI service.

The fourth burden is mineral extraction. AI hardware depends on critical minerals and complex manufacturing systems. The communities living near extraction and processing zones may experience pollution, water depletion, occupational hazards, land degradation, and social disruption. These harms are often distant from the centers where AI products are marketed and consumed. The clean language of “cloud intelligence” can conceal extractive realities on the ground (Crawford, 2021; Crawford & Joler, 2018).

The fifth burden is labor invisibility. AI is presented as automation, but it depends on many forms of human work. There are workers who mine, transport, assemble, fabricate, label, moderate, construct, repair, maintain, clean, secure, and dispose. Some of this labor is highly paid and visible. Much of it is hidden, precarious, outsourced, or geographically distant. A life-coherent analysis must ask not only how AI affects employment through automation, but whose labor makes automation appear autonomous.

The sixth burden is e-waste and toxic afterlives. Rapid hardware turnover can produce large waste streams. Obsolete servers, chips, cables, batteries, storage devices, and cooling equipment do not vanish when replaced. They enter reuse, recycling, export, landfill, or informal processing systems. If governance is weak, the end-of-life phase can expose workers and communities to hazardous substances (Baldé et al., 2024). The symbolic economy’s discarded body becomes someone else’s toxic environment.

The seventh burden is digital dependency. Countries and communities without domestic AI infrastructure may become dependent on external providers for compute, platforms, models, data systems, pricing, access, and governance rules. The benefits of AI may arrive as services, while ownership, strategic control, and value capture remain elsewhere. This is especially important for low- and middle-income countries and Small Island Developing States. They may be encouraged to adopt AI tools while lacking the infrastructural sovereignty to shape their design, environmental terms, or long-term dependencies.

Taken together, these burdens reveal a pattern: AI can concentrate benefits while distributing costs. This is a familiar structure in life-incoherent systems. The gain is privatized, centralized, and monetized. The cost is socialized, localized, ecological, delayed, or displaced. Those who decide may not be those who suffer. Those who profit may not be those whose water, land, labor, or waste sinks are used. Those who consume symbolic outputs may never encounter the people and places that make those outputs possible.

The life-coherence principle is therefore: Consequences must return to decision-makers.

This does not mean every user must know the full supply chain of every prompt. It means institutions must be designed so that decisions cannot escape their life-ground. Data center permitting, AI procurement, cloud contracts, model deployment, infrastructure financing, and public adoption should all require carbon-water-land accounting, lifecycle responsibility, community consultation, and enforceable safeguards. If a community bears the grid, water, land, or waste burden, that community must have voice, protection, benefit, and recourse.

The opposite condition is sacrifice-zone AI. Sacrifice-zone AI occurs when the symbolic benefits of artificial intelligence are produced by imposing environmental, infrastructural, labor, or toxic burdens on communities that lack adequate power to refuse, shape, or benefit from the system. Sacrifice-zone AI can coexist with sustainability branding. A company may purchase renewable energy credits, publish carbon targets, or optimize cooling while still creating local water stress, land conflict, hardware waste, or public infrastructure pressure. Carbon accounting alone cannot reveal the full pattern of burden shifting.

A life-coherent approach requires place-based accountability. It asks where the data center is located, what grid supplies it, what communities share that grid, what water systems support cooling or electricity generation, what land has been converted, what ecological functions are affected, what subsidies or public infrastructure are involved, what hardware supply chains feed the facility, what happens to obsolete equipment, who benefits economically, who governs the risks, who can object, who can enforce repair, and who is left dependent.

These questions shift AI governance from abstract ethics to lived consequences. A technology becomes life-incoherent when its benefits travel upward and outward while its harms settle downward and locally. The task, then, is not merely to make AI cleaner. It is to make AI answerable: to watersheds, grids, workers, communities, ecosystems, future generations, and the life-ground.

LOCAL COSTS, DISTANT BENEFITS

AI workloads impose concentrated environmental burdens where infrastructure is located, while benefits and value accrue elsewhere.

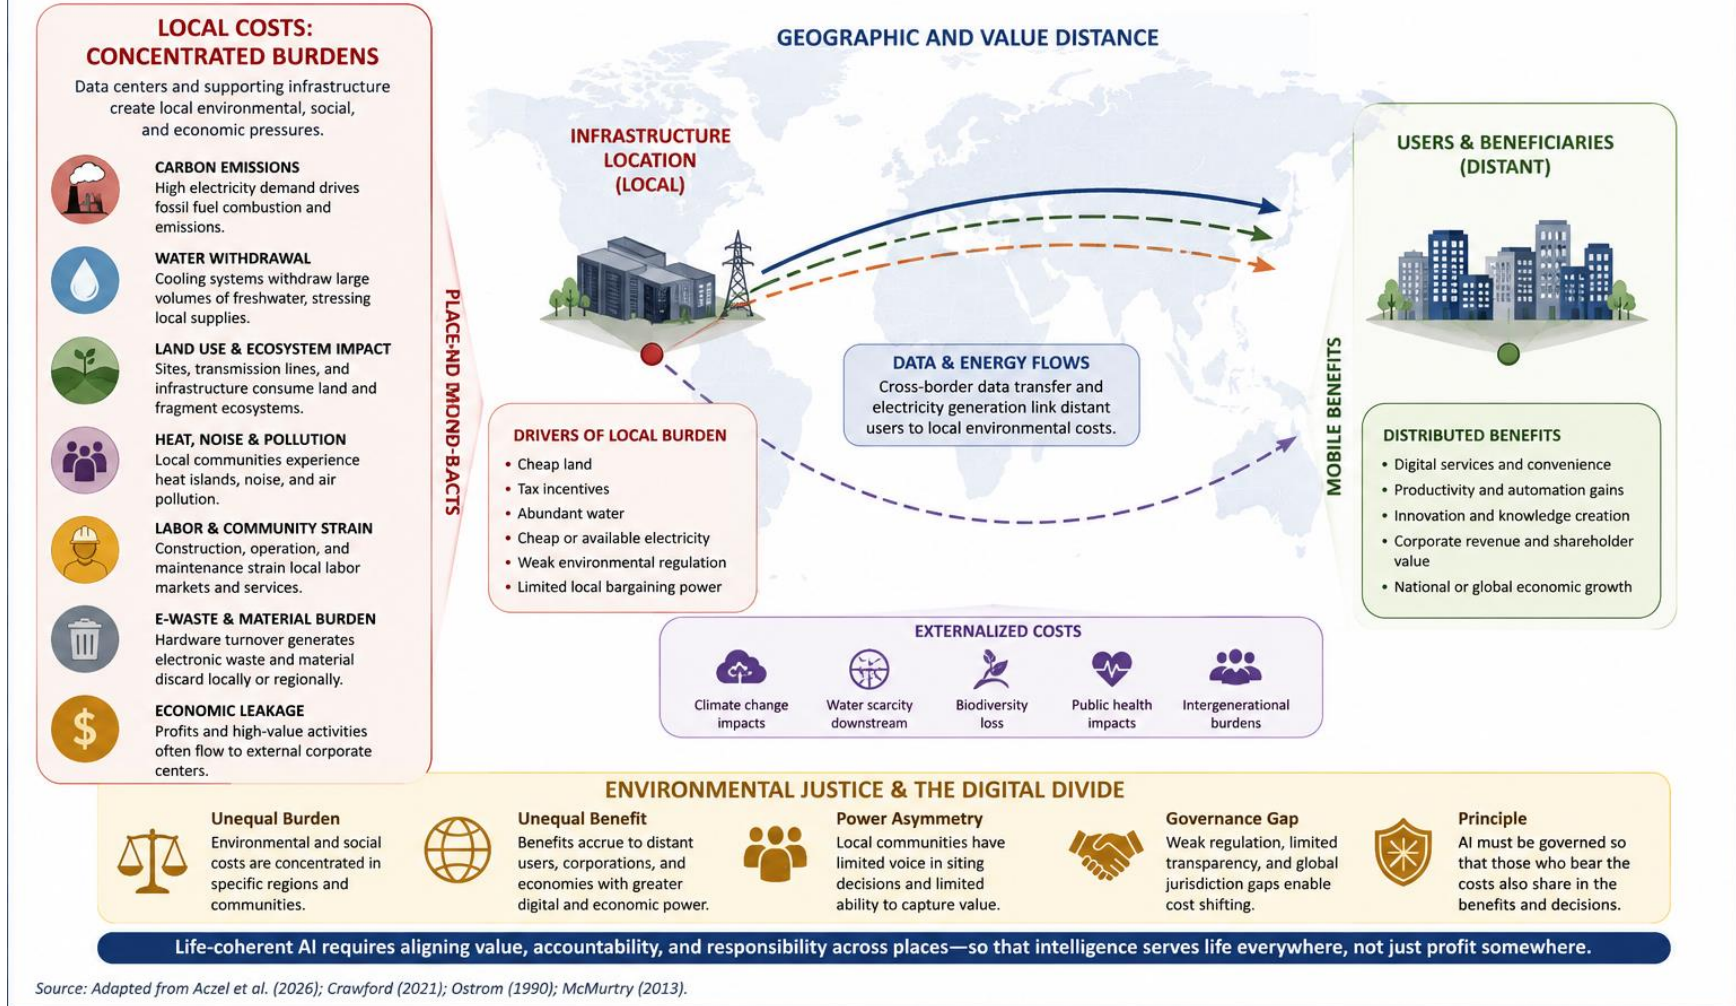


Figure 5. Local Costs, Distant Benefits

This diagram maps the justice asymmetry of AI infrastructure. It shows how digital services, productivity gains, innovation, corporate revenue, and strategic advantages may flow toward distant users and beneficiaries, while carbon emissions, water withdrawal, land transformation, labor strain, e-waste, ecosystem disruption, and public health burdens are concentrated near infrastructure sites and supply chains. The diagram emphasizes that life-coherent AI requires aligning value, accountability, and responsibility across places.

6. The Political Economy of AI Demand

The environmental burden of artificial intelligence is often discussed as though it arises from user behavior alone. People prompt too much, generate too many images, request long answers, automate unnecessary tasks, or choose high-intensity outputs when simpler forms would suffice. These behaviors matter. But they do not explain the deeper structure of demand.

AI demand is not merely chosen. It is produced.

It is produced by platform design, product defaults, business models, corporate competition, investor expectations, cloud infrastructure, search integration, office software integration, institutional procurement, automation strategies, advertising markets, military and security priorities, and the wider political economy of digital expansion. The user sees a tool. The system sees a growth frontier.

This distinction is essential. If AI's environmental burden is framed mainly as the result of individual overuse, responsibility is displaced downward onto users while the deeper machinery of demand generation remains untouched. A life-coherent analysis must therefore ask not only how people use AI, but how societies are being reorganized to require, normalize, monetize, and expand AI use.

The first driver is platform defaulting. AI is increasingly inserted into existing digital environments as a default layer. Search engines offer generative summaries. Office software suggests writing assistance. Email platforms generate replies. Browsers integrate AI copilots. Design tools encourage image generation. Customer service systems default to chatbots. Learning platforms add AI tutors. Smartphones embed AI assistants. Enterprise systems incorporate AI analytics and workflow automation. In each case, what begins as an optional tool becomes part of the ordinary interface.

Defaults matter because they shape behavior before conscious choice occurs. A default summary will be used more often than a feature that must be intentionally selected. A system that automatically routes a query through generative AI creates demand that may not have existed. A platform that presents AI enhancement as the normal path transforms ordinary digital activity into inference load. The user may not experience this as increased consumption, but the infrastructure does.

The second driver is search integration. Search is one of the most frequent digital behaviors in the world. When AI-generated answers are layered onto search, even a modest increase in energy per query can become significant at global scale. The question is not whether AI-enhanced search can be useful. It can be. The question is whether every search requires generative processing. Many queries are navigational, factual, or simple. If all are routed through high-intensity AI systems by default, the result is not intelligence proportional to need. It is symbolic over-processing.

The third driver is corporate automation strategy. Organizations adopt AI to reduce labor costs, accelerate workflows, personalize services, expand analytics, and gain competitive advantage. Some of this use may be genuinely life-serving. AI can reduce administrative burden, support diagnosis, improve translation, detect risks, assist research, and expand accessibility. But corporate adoption is often governed by productivity, cost reduction, and competitive pressure rather than life-capacity. When AI is adopted primarily to automate, replace, accelerate, monitor, or extract, its environmental cost is joined to social cost.

The fourth driver is the frontier model arms race. Companies and states compete to build larger, faster, more capable, more multimodal systems. Model releases become market signals. Capability benchmarks become prestige indicators. Investors reward scale. Firms fear being left behind. Governments frame AI capacity as strategic infrastructure. This produces an arms-race dynamic in which the question becomes not "what level of intelligence is sufficient for life-serving tasks?" but "who can build the most powerful system first?" An arms race does not naturally produce sufficiency. It produces escalation.

The fifth driver is venture capital and financialized infrastructure. AI expansion is not only a technological process. It is an investment frontier. Capital flows into model companies, chip manufacturers, data center construction, cloud platforms, energy procurement, land acquisition, cooling technologies, and AI-enabled enterprise services. The expectation is growth: more users, more subscriptions, more automation, more infrastructure, more market capture, more data, more compute. When finance organizes the direction of technological development, the question of need is displaced by the question of return. A technology does not have to be life-serving to attract investment. It has to promise growth, market power, monetizable dependency, or strategic advantage (McMurtry, 2013).

The sixth driver is cloud concentration. Advanced AI depends on specialized compute capacity that is heavily concentrated among a small number of corporations and countries. This concentration shapes who controls access, pricing, infrastructure siting, environmental disclosure, data governance, and technical standards. It also means that many institutions and countries become consumers of AI capacity rather than co-governors of it. Cloud concentration is a form of enclosure. Compute becomes a gatekeeping infrastructure for knowledge, automation, research, education, public administration, and economic participation.

The seventh driver is attention capture and synthetic media expansion. AI is not only used to answer questions. It is used to generate content, personalize feeds, optimize persuasion, produce advertising, simulate intimacy, create synthetic influencers, automate engagement, and increase the volume of media circulating through digital systems. This expands not only energy demand but symbolic saturation. A society already struggling with attention fragmentation may use AI to generate more text, more images, more videos, more recommendations, more notifications, more persuasion, and more noise. Energy, water, land, and minerals are converted into content streams that may degrade attention, understanding, mental health, and public reason.

The eighth driver is institutional imitation. Organizations often adopt AI because others are adopting AI. Schools, universities, hospitals, ministries, companies, media organizations, and nonprofits may feel compelled to integrate AI to appear modern, efficient, or competitive. This creates a legitimacy pressure. AI becomes a sign of innovation even when its actual contribution is unclear.

The ninth driver is public subsidy without public control. Data centers and AI infrastructure often depend on public goods: electricity grids, water systems, roads, land-use permissions, tax incentives, workforce development, research institutions, and legal stability. Yet the resulting infrastructure may remain privately controlled, environmentally opaque, and oriented toward profit rather than public need. This creates a commons inversion. Public systems support private compute, while private compute shapes public dependency.

The tenth driver is geopolitical competition. AI is increasingly framed as a matter of national security, economic sovereignty, and global power. States compete for chips, data centers, research talent, energy access, and model leadership. This geopolitical framing can accelerate infrastructure buildout and reduce willingness to impose ecological limits. When AI becomes a strategic race, restraint may be portrayed as weakness. Yet a race without ecological discipline is not sovereignty. It is dependency on an escalating system of resource conversion.

These drivers reveal why AI demand cannot be treated as natural. It is structurally generated by systems organized around growth, competition, enclosure, and symbolic expansion. The environmental footprint of AI is therefore not simply the footprint of user curiosity. It is the footprint of a political economy.

Sustainable AI cannot rely only on user education. Users should be encouraged to choose lighter, fit-for-purpose forms of AI, but they cannot counteract platform-wide defaulting, enterprise integration, or infrastructure investment on their own. Responsibility must be placed where design and scale are determined. Efficiency standards must be paired with demand governance. AI environmental disclosure must occur at the level of products, platforms, institutions, and infrastructure. Public institutions should not adopt AI merely because it is available. Compute governance should be treated as a commons question. AI demand should be disciplined by the principle of sufficiency.

The growth question asks: How far can AI scale? The life-coherent question asks: What forms of AI are worthy of scaling because they expand life-capacity within ecological limits? The growth question asks: How can we make AI cheaper, faster, and more ubiquitous? The life-coherent question asks: How can we make AI accountable, sufficient, and subordinate to the commons of life? These questions lead directly into the diagnostic framework of AI as Tool, Oracle, Idol, Enclosure, or Commons.

7. AI as Tool, Oracle, Idol, Enclosure, or Commons

The environmental cost of artificial intelligence cannot be understood apart from the social role AI is being invited to play. A technology's footprint is not only a matter of physical intensity. It is also a matter of purpose, scale, governance, and cultural meaning. The same computational capacity may serve medicine, ecological monitoring, education, accessibility, disaster preparedness, public reasoning, or democratic coordination. It may also serve surveillance, labor displacement, advertising manipulation, synthetic media saturation, financial speculation, military targeting, behavioral prediction, and platform dependency.

The life-coherence question is therefore not simply: How much does AI cost? It is also: What is AI becoming in the life of society?

This paper proposes five interpretive roles through which AI can be evaluated: Tool, Oracle, Idol, Enclosure, and Commons. These are not rigid categories. They are diagnostic orientations. A single system may function as a tool in one context and as an enclosure in another. A model may begin as a tool and become an oracle when its outputs are treated as authoritative. A platform may promise commons-like access while enclosing compute, data, attention, and infrastructure. A society may claim to use AI pragmatically while gradually reorganizing its institutions around the idol of artificial intelligence.

AI functions as a tool when it is bounded, task-appropriate, transparent, accountable, and subordinated to human and ecological purposes. Tool-use is defined by proportionality. The technology is selected because it helps accomplish a real task better, more safely, more accessibly, or more fairly than available alternatives, without displacing responsibility or exceeding ecological and social limits. In healthcare, education, climate governance, environmental monitoring, and public administration, AI can be life-serving when it remains subordinate to a life-serving purpose. The smallest adequate model, lightest adequate modality, and clearest accountable process should be preferred. A diagnostic assistant cannot become the physician. A policy model cannot become the polity. A risk score cannot become justice. A generated summary cannot become understanding.

AI becomes an oracle when its symbolic outputs are treated as superior to situated human judgment. In this role, the machine is no longer merely assisting inquiry. It is being asked to speak with authority. Its fluency, speed, scale, and apparent neutrality create the illusion that it sees more clearly than embodied persons, communities, traditions, disciplines, or democratic processes. The oracle function is subtle because it often begins in usefulness. AI can synthesize information quickly. It can identify patterns that humans may miss. But when these capacities are detached from context, accountability, and lived consequence, they can begin to displace the forms of judgment that make knowledge responsible. Human judgment is embodied, formed through experience, error, care, responsibility, memory, moral development, and participation in a world shared with others (Maturana & Varela, 1992). AI can model language about these processes, but it does not live them.

AI becomes an idol when society sacrifices life-support to it while treating its expansion as inevitable, sacred, or unquestionable. An idol is not merely something admired. It is something to which real life is subordinated. In the AI context, idolatry appears when energy, water, land, minerals, labor, public money, attention, education, institutional design, and ecological stability are reorganized around the demand that AI must grow. The idol says that larger models must be built, every platform must include AI, every institution must adopt AI, every task must be automated, and whatever AI requires must be provided. This is not technological maturity. It is symbolic capture.

AI becomes an enclosure when the infrastructures of knowledge, communication, computation, data, labor, and public reasoning are captured by private or concentrated power. Enclosure is the conversion of shared life-capacity into controlled access. Historically, enclosure meant common land was fenced, privatized, and reorganized for private gain. In the digital age, enclosure can occur through platforms, data ownership, cloud dependency, proprietary models, closed standards, algorithmic gatekeeping, and control over compute (Ostrom, 1990). AI enclosure has symbolic and material dimensions: collective human meaning may become

raw material for proprietary systems, while compute capacity, data centers, chips, cloud platforms, energy contracts, and AI infrastructure are controlled by a small number of corporations or powerful states.

AI becomes a commons when its capacities are governed to expand shared life-capacity within ecological limits. Commons does not mean unrestricted use. It means shared stewardship. It means that access, responsibility, benefit, and boundary are held together. A commons is not a free-for-all. It is a living governance arrangement that protects the conditions of shared flourishing (Ostrom, 1990; McMurtry, 2002, 2013).

AI as commons would prioritize public-interest uses: health, education, ecological repair, climate adaptation, disaster preparedness, water governance, accessibility, translation, scientific research, democratic participation, and local knowledge preservation. It would support smaller, task-appropriate, auditable models where possible. It would disclose carbon, water, land, and lifecycle footprints. It would be governed by sufficiency rather than endless growth. It would include affected communities in infrastructure decisions. It would treat compute as a public-interest capacity, not only a private commodity.

The commons pathway requires public-interest compute, environmental transparency, sufficiency by design, lifecycle responsibility, community consent and benefit, knowledge plurality, and democratic governance. AI commons must not homogenize the world through dominant languages, dominant datasets, dominant epistemologies, and dominant corporate values. Local knowledge, Indigenous knowledge, ecological knowledge, clinical judgment, cultural memory, and lived experience must not be reduced to data extraction. A commons protects diversity because life itself depends on diversity.

The five categories can be understood as a diagnostic movement. AI begins as a tool when it is used for bounded assistance. It becomes an oracle when its outputs are treated as superior to situated judgment. It becomes an idol when society sacrifices life-support to its expansion. It becomes an enclosure when the infrastructures of knowledge, compute, and life-support are captured by concentrated power. It becomes a commons only when it is governed for shared life-capacity within ecological limits.

The movement from tool to oracle to idol to enclosure is a life-incoherent drift. It can happen gradually, without explicit decision. A useful assistant becomes a default authority. A default authority becomes an object of institutional dependence. Institutional dependence becomes infrastructural capture. Infrastructural capture becomes enclosure. The movement toward commons requires deliberate reversal. It asks that AI be brought back under the governance of life: purpose, proportion, transparency, accountability, sufficiency, justice, ecological limits, and shared benefit.

The question is not whether AI is intelligent. The question is whether its intelligence is coherent with life.

AI as Tool, Oracle, Idol, Enclosure, or Commons

Five Social Roles of Artificial Intelligence



Figure 6. AI as Tool, Oracle, Idol, Enclosure, or Commons.

This diagram presents five social roles of artificial intelligence. AI functions as a tool when bounded and life-serving; as an oracle when machine output is treated as authority; as an idol when society sacrifices life-support to AI expansion; as an enclosure when compute, data, knowledge, and infrastructure are captured by concentrated power; and as a commons when AI is governed through shared stewardship for life-capacity within ecological and social limits.

Table 3. Diagnostic Matrix: Tool, Oracle, Idol, Enclosure, Commons

AI role	Defining feature	Life-coherence status	Warning signs	Corrective pathway
Tool	Bounded assistance for a real task	Potentially life-coherent	Overuse, poor fit, lack of accountability	Fit-for-purpose use, human judgment, transparency
Oracle	Machine output treated as superior authority	Life-risking	Deference to AI, weakened judgment, no contestability	Restore situated judgment, auditability, appeal
Idol	Life-support sacrificed to AI expansion	Life-incoherent	Inevitability language, unbounded scaling, symbolic awe	Desacralize AI; impose sufficiency and limits
Enclosure	Compute, data, knowledge, and infrastructure captured by concentrated power	Life-incoherent	Dependency, opacity, monopoly control, public subsidy without public governance	Build public-interest compute and commons governance
Commons	Shared stewardship of AI for life-capacity within limits	Life-coherent	Risk of capture or under-governance	Democratic oversight, lifecycle responsibility, community benefit

8. A Life-Coherent Governance Framework for AI

If artificial intelligence is materially grounded in energy, water, land, minerals, labor, communities, ecosystems, and waste sinks, then AI governance cannot be limited to data protection, algorithmic bias, intellectual property, or model safety. These concerns are necessary, but incomplete. AI governance must also become ecological, infrastructural, social, and democratic. It must govern not only what AI outputs, but what AI consumes, displaces, concentrates, normalizes, and makes dependent.

A life-coherent governance framework begins with a simple principle: Artificial intelligence must be subordinated to the conditions of life.

This does not mean technological rejection. It means that AI should be developed, deployed, and used only in ways that expand life-capacity within ecological and social limits. It means that the symbolic power of AI must be accountable to its material life-ground. It means that efficiency must be joined to sufficiency, transparency to enforceability, innovation to public purpose, and access to shared responsibility.

The UNU-INWEH report identifies six guiding principles for a responsible AI ecosystem: transparency, efficiency by design, equity and environmental justice, lifecycle responsibility, global cooperation, and sustainable use. The UNU-INWEH principles of transparency, efficiency by design, equity and environmental justice, lifecycle responsibility, global cooperation, and sustainable use align with broader responsible AI ethics frameworks, but require stronger grounding in ecological limits, lifecycle accountability, and the civil commons (Aczel et al., 2026; Floridi & Cowls, 2019). These principles provide a necessary foundation. The life-coherence framework extends them into a governance cycle that asks: What is the purpose of the AI use? What is the smallest adequate model? What are the carbon, water, land, hardware, labor, and waste implications? Who benefits? Who bears the burden? Who decides? What safeguards exist? What repair is possible? What limits must be respected?

A life-coherent governance framework has ten interlocking components.

First, purpose: life-value before deployment. The first governance question should not be whether AI can perform a task. It should be whether the task is worth performing with AI. AI deployment should require a life-value justification. The adopting institution should be able to explain what life-capacity the system expands: health, learning, accessibility, ecological repair, public safety, democratic participation, disaster preparedness, administrative justice, water security, climate adaptation, scientific understanding, or meaningful human work.

Second, proportionality: fit-for-purpose intelligence. A life-coherent AI system should use the smallest adequate model, shortest adequate output, and lowest adequate modality for the task. Proportionality should be built into design. Platforms should offer low-footprint modes. Institutions should require small-model-first routing. Search systems should distinguish between simple lookup, navigational search, synthesis, and deep reasoning.

Third, transparency: making the life-ground visible. AI systems should disclose the material conditions of their operation. This includes energy use, carbon intensity, water footprint, land implications, hardware lifecycle, data center locations where appropriate, cooling strategies, and relevant supply-chain impacts. Transparency must be standardized, comparable, and verifiable (Henderson et al., 2020; Lacoste et al., 2019). It should not be left to selective corporate storytelling. But transparency is not enough. A dashboard can describe a harmful system without changing it. Disclosure must be tied to decision rights, limits, accountability, and repair.

Fourth, sufficiency: efficiency within limits. Efficiency asks how to reduce the footprint per task. Sufficiency asks which tasks should be performed, how often, at what scale, and for what purpose. Without sufficiency, efficiency can intensify expansion. A more efficient model may be deployed everywhere. A cheaper image generator may encourage more image generation. A faster chatbot may increase total prompting. Unit impacts fall while total impacts rise. Efficiency becomes life-coherent only when governed by sufficiency (Schwartz et al., 2020).

Fifth, lifecycle responsibility: from mine to model to waste. AI governance must cover the full material chain: mineral extraction, manufacturing, chip fabrication, server deployment, cooling infrastructure, energy procurement, maintenance, hardware turnover, reuse, recycling, and safe disposal. A model cannot be called responsible if its hardware lifecycle is irresponsible. Lifecycle responsibility also includes labor. The workers who mine, fabricate, label, moderate, construct, maintain, clean, repair, and recycle AI infrastructure are part of the system.

Sixth, place-based accountability: no sacrifice-zone AI. Because AI's environmental footprints vary by location, governance must evaluate impacts where they occur. Data centers should be assessed as energy-water-land infrastructures, not ordinary commercial buildings. Permitting should include grid impact analysis, water stress assessment, land-use review, climate resilience evaluation, cumulative impact assessment, community consultation, and lifecycle waste planning.

Seventh, community consent and democratic oversight. AI infrastructure decisions affect public goods: water, energy, land, labor systems, knowledge systems, public services, and democratic life. They should not be made solely by private firms, technical experts, or investment authorities. Communities affected by data centers, energy buildout, water withdrawals, land-use changes, mining, and waste systems should be involved early. AI use in public institutions must also be transparent, contestable, and subject to appeal where rights and services are affected.

Eighth, public-interest compute and the commons. As AI becomes increasingly important for education, health, research, public administration, climate adaptation, and economic participation, compute capacity becomes a civil infrastructure issue. If access to AI depends entirely on concentrated private platforms, societies risk enclosure. A life-coherent alternative is to develop commons-oriented AI infrastructure: public compute facilities, university and research consortia, regional compute cooperatives, open public-interest models, shared environmental reporting standards, public procurement pools, and governance mechanisms that align compute with social need rather than private expansion.

Ninth, knowledge integrity and human judgment. AI governance must protect the living conditions of understanding: attention, education, dialogue, memory, evidence, professional judgment, community knowledge, and democratic reasoning. Environmental sustainability cannot be separated from epistemic sustainability. A system that consumes energy, water, land, and minerals to produce synthetic content that degrades public reason is doubly life-incoherent.

Tenth, continuous review, repair, and withdrawal. AI systems should not be approved once and then allowed to expand indefinitely. Their impacts change over time as use scales, models change, defaults shift, data centers expand, energy mixes evolve, and institutions become dependent. Life-coherent governance requires periodic review of environmental footprints, social impacts, labor effects, community burdens, performance, dependency risks, and public value. Systems that fail life-coherence tests should be redesigned, restricted, replaced, or withdrawn. Withdrawal is an essential governance capacity. A society that cannot stop using a harmful system is not governing it. It is dependent on it.

These ten components can be organized into a governance cycle: Purpose -> Proportionality -> Transparency -> Sufficiency -> Lifecycle Responsibility -> Place-Based Accountability -> Community Consent -> Public-Interest Compute -> Knowledge Integrity -> Review and Repair.


The cycle begins with purpose because AI should not be deployed without a life-value justification. It moves to proportionality because the model and modality must fit the task. It requires transparency because hidden burdens cannot be governed. It adds sufficiency because efficiency alone can accelerate total demand. It includes lifecycle responsibility because AI has a body from mine to waste. It requires place-based accountability because burdens land somewhere. It requires community consent because affected people must have voice. It builds public-interest compute because shared life-capacity cannot depend entirely on enclosure. It protects knowledge integrity because symbolic systems can weaken judgment. It ends with review and repair because governance must remain alive.

This cycle transforms AI governance from compliance into stewardship. Compliance asks whether rules have been followed. Stewardship asks whether life-capacity is being protected and expanded. Responsible AI asks whether harms are managed. Life-coherent AI asks whether the system is worthy of the life-ground it consumes. The difference matters because AI is scaling rapidly. Once infrastructures, habits, dependencies, and investments lock in, correction becomes harder. Governance must therefore act before symbolic power becomes material inevitability.

Compliance is not enough.
Governance must become stewardship.

The Life-Coherent AI Governance Cycle

From Responsible AI Principles to Stewardship of the Life-Ground



Does AI expand life-capacity within ecological limits?




Figure 7. The Life-Coherent AI Governance Cycle.

This diagram presents the governance cycle proposed in the paper: purpose, proportionality, transparency, sufficiency, lifecycle responsibility, place-based accountability, community consent, public-interest compute, knowledge integrity, and review and repair. The cycle shows that compliance is not enough. Life-coherent AI governance must become stewardship of the life-ground, ensuring that AI remains answerable to the conditions of life

Table 4. Stakeholder Responsibilities for Life-Coherent AI

Stakeholder	Core responsibilities
Governments	Integrate AI into energy, water, land-use, climate, labor, procurement, and digital sovereignty policy; require carbon-water-land disclosure; protect communities.
AI developers	Design for proportionality, transparency, auditability, efficiency, sufficiency, and human accountability.
Data center operators	Disclose resource use; avoid water-stressed siting; coordinate with grid planning; manage lifecycle impacts; engage communities.
Utilities and energy planners	Treat AI demand as a major infrastructure variable; protect residential, public-service, and decarbonization priorities.
Investors and financiers	Treat carbon, water, land, labor, lifecycle, and community impacts as material risks.
Public institutions	Adopt AI only where life-value is clear, accountability is preserved, and footprint disclosure is available.
Researchers and universities	Develop low-footprint models, public-interest applications, independent assessment methods, and governance tools.
Civil society	Monitor AI infrastructure, demand disclosure, affected communities, and public reason.
Users and organizations	Practice fit-for-purpose use, avoid unnecessary high-intensity outputs, and ask for footprint transparency.
International institutions	Harmonize standards, support lower-income regions, prevent burden shifting, and advance AI as a shared civil commons.

9. Caribbean and SIDS Implications

The environmental cost of artificial intelligence is a global issue, but it is not experienced evenly. Small Island Developing States face a distinctive set of vulnerabilities and opportunities. Their energy systems are often small, import-dependent, and expensive. Their freshwater systems are limited and climate-sensitive. Their land area is constrained. Their economies are exposed to external shocks. Their digital infrastructures are often dependent on external platforms. Their climate risks are existential rather than abstract. Their public institutions frequently operate with limited technical capacity. Their communities are deeply affected by decisions made elsewhere.

For Caribbean and other Small Island Developing States, the question is not whether artificial intelligence should be accepted or rejected. The question is whether AI will deepen dependency or strengthen life-capacity.

This distinction is crucial. AI may offer real benefits for SIDS: disaster risk reduction, hurricane forecasting, water-system monitoring, health-service support, education, translation, tourism management, fisheries protection, coastal mapping, renewable-energy planning, administrative efficiency, and regional research collaboration. Used wisely, AI could help small states make better use of limited resources, widen access to knowledge, support climate adaptation, and strengthen public services. But AI could also intensify existing vulnerabilities. If imported as a platform dependency, it may place critical public functions under external control. If adopted uncritically, it may increase energy demand without life-value justification. If embedded into education, health, governance, and business without local capacity, it may weaken human judgment and institutional sovereignty.

A life-coherent approach to AI in SIDS must therefore begin from vulnerability, sovereignty, and sufficiency (IPCC, 2022; UNDP, 2024; UN-OHRLLS, n.d.).

Many island energy systems are small, costly, and historically dependent on imported fossil fuels. They face the challenge of transitioning to renewable energy while maintaining reliability, affordability, and resilience. In such contexts, large new electricity loads cannot be treated as ordinary commercial demand. Data centers, AI infrastructure, and high-intensity digital systems may compete with households, hospitals, schools, water utilities, transport electrification, and climate adaptation needs. For SIDS, electricity is not only a market commodity. It is a life-support system. AI infrastructure should not weaken the energy security of small island societies.

Water is one of the most important life-ground issues for SIDS. Many islands face limited freshwater resources, seasonal drought, saltwater intrusion, aging infrastructure, rainfall variability, and increasing climate stress. AI data centers may require water directly for cooling or indirectly through the electricity systems that power them. In water-stressed settings, the question is not only how much water is used in aggregate. It is when it is used, from which source, under what conditions, and in competition with whom. A life-coherent AI policy for SIDS should include water-screening requirements for digital infrastructure. AI itself may help improve water governance through leak detection, rainfall forecasting, aquifer monitoring, desalination optimization, watershed protection, flood mapping, drought planning, and public communication. The distinction is decisive: AI should help protect water, not silently consume it.

Small islands have limited land. Land is needed for housing, farming, tourism, infrastructure, forests, watersheds, biodiversity, cultural life, recreation, renewable energy, and climate adaptation. Coastal land is increasingly threatened by sea-level rise, storm surge, erosion, and development pressure. Data centers and related infrastructure occupy physical space, and their supporting systems may expand their effective footprint. In SIDS, land cannot be treated as empty space for imported development models. It is a scarce life-support matrix.

The UNU-INWEH report notes that advanced AI compute is concentrated in a small number of countries and providers (Aczel et al., 2026). This matters greatly for SIDS. Most small states will not build frontier models or hyperscale compute systems. They will access AI through external platforms, cloud services, imported

tools, vendor contracts, and global software ecosystems. This creates a dependency risk. If schools, hospitals, ministries, courts, utilities, media systems, and businesses become dependent on AI systems controlled elsewhere, small states may lose practical sovereignty over knowledge, data, pricing, access, cybersecurity, environmental accountability, and institutional memory.

A life-coherent SIDS strategy should therefore develop compute sovereignty at the appropriate scale. This does not mean every island must build expensive frontier infrastructure. It means small states and regions should retain enough capacity to govern the AI systems they use. This may include regional AI policy frameworks, public-sector technical teams, local data governance, shared Caribbean compute resources, open-source and low-footprint models, university partnerships, public procurement standards, and regional environmental disclosure requirements. For the Caribbean, regional cooperation may be more realistic and life-coherent than isolated national efforts. A shared Caribbean AI commons could support climate adaptation, disaster risk reduction, public health, education, language and culture, marine protection, water governance, and research.

Public-sector AI adoption in SIDS should begin with a life-value test. Does the system improve access to public services? Does it reduce burdens without reducing accountability? Does it strengthen local capacity rather than replace it? Does it preserve human review and appeal? Does it protect sensitive data? Does it use proportionate models? Does it disclose environmental and vendor dependencies? Does it align with national development goals? Does it build institutional learning? Does it remain reversible?

For small societies, education and culture are not merely sectors. They are the means by which memory, identity, judgment, language, and civic capacity are reproduced across generations. AI can support education, but it can also weaken learning if used as a shortcut around attention, reading, writing, memory, reasoning, and dialogue. The risk is not only plagiarism or cheating. The deeper risk is dependency on synthetic fluency. A life-coherent educational approach would teach AI literacy as part of life literacy: what AI is, what it is not, what it costs, how it can help, how it can mislead, how it affects the environment, how to use it proportionately, and how to preserve human judgment.

Health systems in SIDS often face specialist shortages, chronic disease burdens, constrained budgets, migration of health workers, fragmented data systems, and climate-sensitive health risks. AI may offer meaningful support in clinical decision assistance, public health surveillance, health education, workflow management, translation, imaging, triage, and patient communication. These uses may be life-coherent if they strengthen care. But health AI must be governed carefully. Imported models may not reflect local populations, disease patterns, health infrastructure, or cultural context. Data privacy risks may be significant. Overreliance on AI may weaken clinical reasoning.

Climate adaptation is one of the strongest life-coherent use cases for AI in SIDS. Small islands face hurricanes, floods, droughts, heat stress, coral reef decline, coastal erosion, sea-level rise, vector-borne disease risks, and food and water insecurity. AI can support early warning, hazard mapping, evacuation planning, infrastructure risk assessment, rainfall forecasting, coastal monitoring, insurance analysis, ecosystem protection, and recovery coordination. In this domain, AI can help make complexity actionable. But even here, models must be transparent enough for public use, local knowledge must be included, false precision must be avoided, data must be updated, and communities must not be reduced to risk scores.

The AI dependency trap occurs when a small state adopts AI tools rapidly without developing the capacity to govern them. At first, the tools appear useful and efficient. Over time, public institutions, businesses, schools, and citizens become dependent on external platforms. Local expertise is not built. Data governance remains weak. Environmental footprints are unknown. Vendor costs rise. Public reasoning is mediated by systems designed elsewhere. Local knowledge is underrepresented. Critical services become difficult to operate without external AI providers. The trap is not obvious at the beginning because dependency often arrives as convenience.

To avoid this, SIDS need a deliberate strategy: adopt AI selectively, build local and regional capacity, use open and auditable systems where possible, require environmental and data disclosure, prioritize public-interest applications, avoid AI-by-default procurement, protect education and professional judgment, invest in human skills alongside digital tools, collaborate regionally, and retain the capacity to refuse, exit, or replace systems. The most important sovereignty is not the ability to build everything. It is the ability to govern what one depends on.

A Caribbean AI commons would not imitate the frontier race. It would begin from the region's own life-ground: watersheds, reefs, hurricanes, public health, schools, migration, culture, small economies, fragile grids, limited land, and deep community knowledge. It would ask how AI can serve these realities without enclosing them. The guiding principle should be: Do not scale AI faster than the capacity to govern its life-ground.

Artificial intelligence may have a place in the Caribbean future. But that place must be chosen, not imposed. It must be proportionate, not fashionable. It must be accountable, not opaque. It must serve public life, not enclose it. It must protect water, energy, land, culture, judgment, and community. In small islands, the life-ground is never abstract. It is the watershed, the coastline, the school, the clinic, the family, the reef, the road, the power line, the cistern, the ferry, the church, the market, the playing field, the memory of storms, and the knowledge of neighbors. AI becomes life-coherent only when it bows to that ground.

SIDS AND THE AI DEPENDENCY TRAP

From imported AI dependency to a Caribbean AI commons for life-capacity and climate resilience

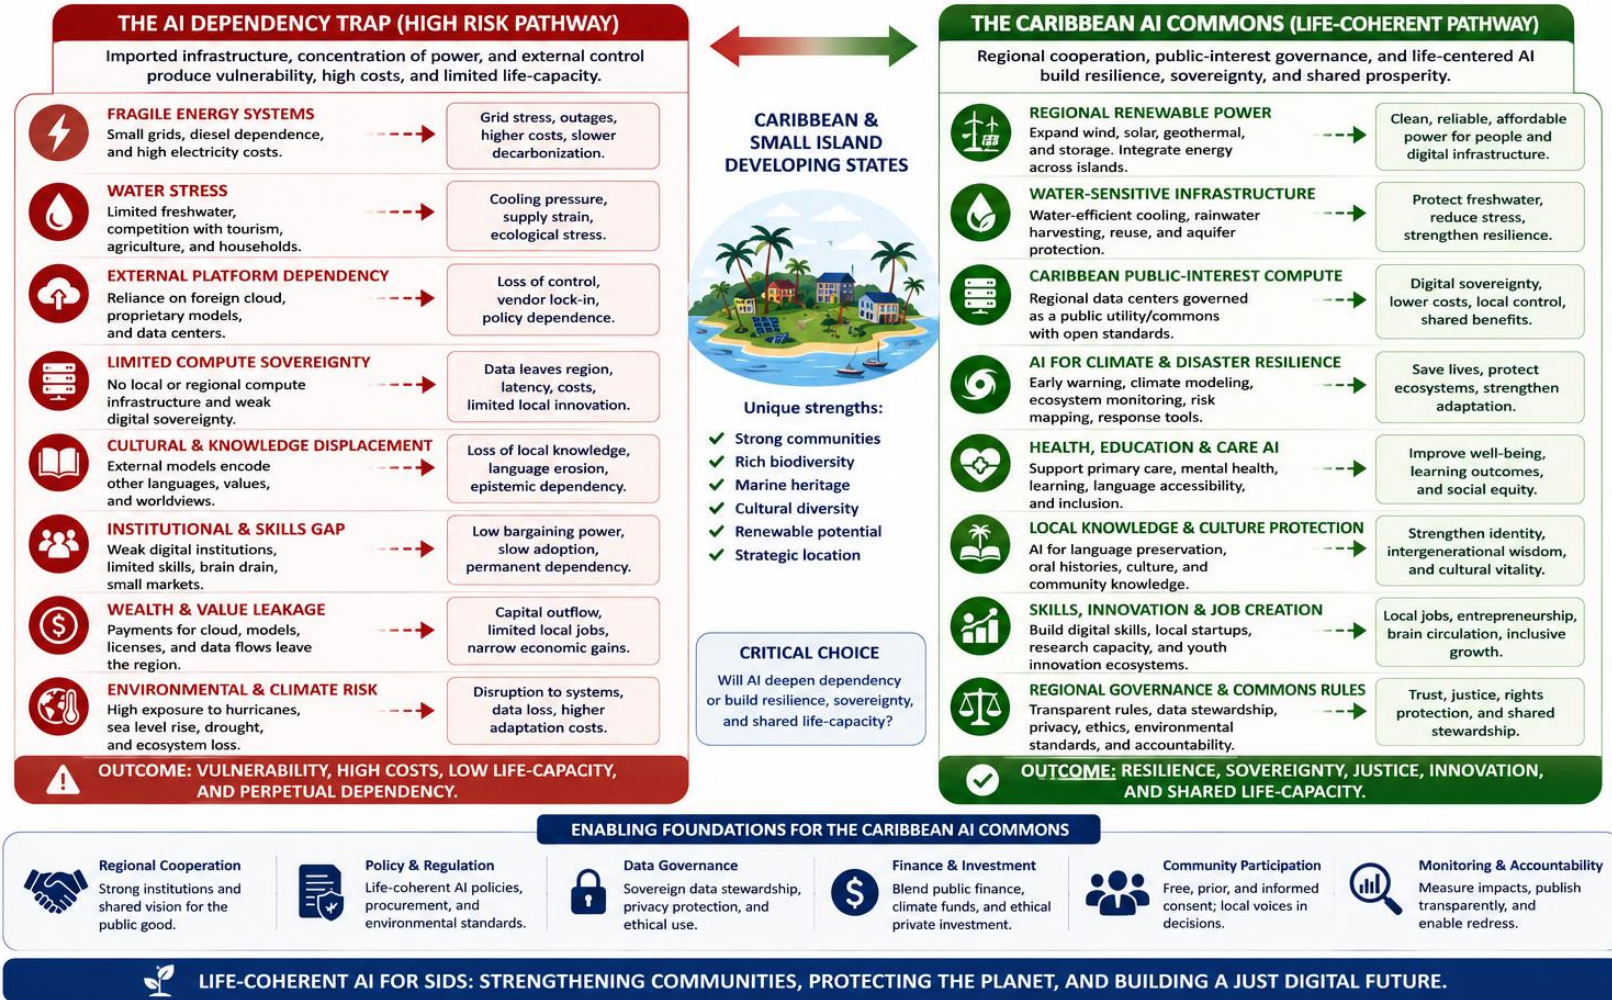


Figure 8. SIDS and the AI Dependency Trap

This diagram contrasts two pathways for Caribbean and Small Island Developing States. The high-risk pathway shows imported AI dependency, fragile energy systems, water stress, external platform control, limited compute sovereignty, cultural displacement, institutional capacity gaps, value leakage, and climate vulnerability. The life-coherent pathway shows a Caribbean AI commons based on regional cooperation, renewable energy, water-sensitive infrastructure, public-interest compute, climate and disaster resilience, health and education support, local knowledge protection, skills development, and shared governance.

10. The Life-Coherent AI Use Protocol

The environmental and social governance of artificial intelligence cannot remain only at the level of principles. Principles are necessary, but they must become usable. Institutions, communities, educators, clinicians, policymakers, researchers, companies, public agencies, and ordinary users need a practical way to decide when AI use is justified, when it is excessive, and how its hidden life-ground costs should be made visible.

The Life-Coherent AI Use Protocol is designed for that purpose. It translates the paper's central question into a practical decision sequence: Does this AI use expand life-capacity within ecological limits, or does it convert life-support into symbolic excess, dependency, and enclosure?

The protocol does not assume that AI is inherently harmful. Nor does it assume that AI is inherently beneficial. It asks for discernment. Artificial intelligence can support health, education, climate adaptation, accessibility, public service, scientific research, water governance, disaster preparedness, and ecological repair. It can also intensify extraction, surveillance, labor displacement, misinformation, symbolic overload, dependency, and ecological burden. The difference depends on purpose, proportionality, footprint, governance, and accountability.

The protocol is organized around ten questions.

First, the need question: Is AI genuinely needed for this task? AI should not be used merely because it is available, fashionable, impressive, or already embedded in a platform. Many tasks can be done through ordinary human judgment, existing documents, simple search, basic software, rules-based systems, conversation, local knowledge, or institutional memory. The refusal of unnecessary AI is not backwardness. It is intelligence governed by sufficiency.

Second, the purpose question: What life-capacity is this AI use intended to expand? A justified AI use should be able to name the life-capacity it serves. Vague claims of innovation, efficiency, competitiveness, personalization, modernization, or productivity are not enough. The purpose should be concrete.

Third, the proportionality question: Is the smallest adequate model and lightest adequate modality being used? A simple task should not be routed through an unnecessarily large model. A brief answer should not become a long generative output. A text explanation should not become an image unless the visual form adds real value. An image should not become video unless motion is necessary. High-resolution or high-complexity outputs should not be default settings.

Fourth, the footprint question: What are the carbon, water, land, and energy implications of this use? Every AI use has a material footprint, even when that footprint is hidden. For an individual user, the footprint question may be simple: Do I need this long answer, image, or video? Could a shorter or simpler output meet the need? For an institution, the footprint question must be more formal. Procurement and deployment should ask vendors for energy use estimates, carbon accounting, water footprint information, land and siting implications, model-routing policies, data center locations where appropriate, and lifecycle practices.

Fifth, the lifecycle question: What material chain supports this AI system from extraction to disposal? AI is not only electricity. It is also hardware. Chips, servers, storage systems, cooling equipment, batteries, cables, buildings, substations, and networks all have lifecycles. They depend on minerals, manufacturing, transport, labor, maintenance, replacement, recycling, and waste systems.

Sixth, the justice question: Who benefits, who bears the burden, and who decides? AI systems often separate benefits from burdens. The justice question asks whether affected people and places are visible in the decision. An AI use that expands convenience for the powerful while shifting burdens onto the vulnerable is life-incoherent, even if it is technically efficient.

Seventh, the dependency question: Does this AI use strengthen local capacity or deepen dependency? AI can either build capacity or replace it. It can help people learn, understand, coordinate, diagnose, plan, and repair.

Or it can make institutions dependent on external vendors, opaque systems, proprietary platforms, and machine-generated judgment.

Eighth, the knowledge integrity question: Does this AI use support understanding, or does it substitute synthetic fluency for judgment? AI can produce fluent outputs without living understanding. It can summarize, imitate, explain, and generate with confidence. But fluency is not wisdom. Prediction is not responsibility. Synthesis is not conscience. Simulation is not relationship.

Ninth, the governance question: Is there transparency, accountability, appeal, review, and repair? No AI system should be trusted simply because it performs well in demonstration. Governance must accompany deployment. The higher the stakes, the greater the need for human review, public accountability, and enforceable safeguards. AI that cannot be governed should not be deployed in high-stakes contexts.

Tenth, the commons question: Does this AI use contribute to the commons of life? The commons of life includes the shared conditions that make flourishing possible: clean water, reliable energy, healthy ecosystems, public knowledge, cultural memory, education, health, democratic agency, meaningful work, trust, care, and intergenerational viability. An AI use contributes to the commons when it strengthens these shared conditions. It becomes life-incoherent when it privatizes benefits, hides costs, displaces burdens, weakens judgment, deepens dependency, or converts public life-support into private symbolic power.

The protocol can be used at different levels. For individuals, it encourages mindful use: shorter prompts where sufficient, fewer unnecessary images, avoidance of disposable video generation, careful checking of outputs, and recognition that every symbolic output has a life-ground. For educators, it supports AI literacy: teaching students not only how to use AI, but how to judge when not to use it, how to preserve learning, and how to understand environmental costs. For clinicians and health institutions, it asks whether AI improves care, protects accountability, and strengthens clinical judgment without creating unsafe dependency. For governments, it provides a procurement and policy screen: no AI adoption without life-value justification, environmental disclosure, data protection, accountability, appeal, and lifecycle review.

The protocol can lead to four possible decisions: proceed, proceed with limits, redesign, or refuse or withdraw. Proceed when the AI use has clear life-value, is proportionate, has acceptable and disclosed footprints, protects knowledge integrity, includes accountability, and contributes to the commons. Proceed with limits when the AI use is valuable but requires restrictions: smaller models, shorter outputs, lower frequency, local safeguards, environmental reporting, human review, or community consultation. Redesign when the purpose is valid but the current system is too resource-intensive, opaque, dependency-producing, unjust, or poorly governed. Refuse or withdraw when the AI use lacks life-value, produces symbolic excess, hides unacceptable burdens, weakens judgment, deepens dependency, or cannot be governed responsibly.

The ability to refuse is essential. Without refusal, governance becomes performance. A society that cannot say no to AI cannot say yes responsibly. At the smallest scale, the protocol may change how a person prompts. At the institutional scale, the protocol changes procurement, design, education, policy, and oversight. At the civilizational scale, the protocol asks whether technological intelligence can be brought under the discipline of life. Artificial intelligence becomes responsible only when it can answer to life.

Table 5. Life-Coherent AI Use Checklist

Question	Diagnostic prompt	Possible decision
Need	Is AI genuinely needed for this task?	Proceed only if AI adds real value.
Purpose	What life-capacity does this use expand?	Clarify or reject vague innovation claims.
Proportionality	Is the smallest adequate model and lightest adequate modality being used?	Reduce model size, output length, or modality if possible.
Footprint	What are the carbon, water, land, and energy implications?	Require disclosure or choose a lower-footprint pathway.
Lifecycle	What extraction, hardware, labor, and waste systems are involved?	Require lifecycle responsibility.
Justice	Who benefits, who bears the burden, and who decides?	Add safeguards, consent, benefit-sharing, or refuse.
Dependency	Does this strengthen local capacity or deepen external reliance?	Build capacity, preserve exit options, avoid lock-in.
Knowledge integrity	Does this support understanding or replace judgment with fluency?	Preserve human learning, review, and accountability.
Governance	Is there transparency, appeal, review, and repair?	Do not deploy high-stakes AI without governance.
Commons	Does this contribute to shared life-capacity within ecological limits?	Proceed, limit, redesign, or withdraw.

11. Conclusion: From Symbolic Power to Responsible Intelligence

Artificial intelligence has entered public life as symbolic power. It answers, predicts, classifies, summarizes, generates, translates, recommends, simulates, and automates. It produces language, images, code, voices, music, video, and decisions with a speed and fluency that can appear almost weightless. To the user, AI often arrives as convenience. To institutions, it arrives as efficiency. To markets, it arrives as growth. To states, it arrives as strategic capacity. To culture, it arrives as spectacle.

But the central finding of this paper is that AI is not weightless.

Artificial intelligence has a life-ground. It depends on electricity, water, land, minerals, labor, communities, ecosystems, and waste sinks. Its outputs are symbolic, but its conditions are material. Its interface is digital, but its metabolism is ecological. Its intelligence is artificial, but the life-support systems it consumes are real.

The UNU-INWEH report makes this visible by measuring AI's environmental costs through carbon, water, and land footprints (Aczel et al., 2026). That multidimensional framing is essential because carbon alone cannot reveal the full pattern of burden. A low-carbon pathway may still be water-intensive. A renewable pathway may still occupy land. A data center may serve distant users while reshaping local grids, watersheds, landscapes, and communities. A model may appear efficient per query while total use expands through rebound effects. A digital service may seem clean while its hardware lifecycle depends on mining, manufacturing, labor, and e-waste systems that remain hidden from the interface.

This paper has extended that empirical foundation through the life-coherence framework. The question is not only how large AI's footprint is. The deeper question is whether AI's use of the life-ground is justified by its contribution to life-capacity.

Life-capacity means the real conditions that allow people, communities, and ecosystems to maintain and develop their powers of life: health, learning, care, ecological stability, clean water, reliable energy, meaningful work, cultural memory, democratic agency, public reason, and intergenerational viability. A technology is life-coherent when it strengthens these conditions within ecological and social limits. It becomes life-incoherent when it consumes, degrades, displaces, or encloses them while presenting its outputs as progress.

Artificial intelligence now stands at this threshold. It can serve as a tool. It can help clinicians, teachers, researchers, public servants, communities, and citizens understand complexity, reduce burdens, improve access, support adaptation, and repair damaged systems. It can assist climate modeling, water governance, disaster preparedness, health education, translation, ecological monitoring, and scientific discovery. Used with discipline, transparency, proportionality, and accountability, AI can expand life-capacity.

But AI can also become an oracle, idol, and enclosure. It becomes an oracle when its fluent outputs displace situated judgment. It becomes an idol when society sacrifices energy, water, land, minerals, labor, attention, and ecological stability to its expansion. It becomes an enclosure when compute, data, knowledge, public infrastructure, and symbolic production are captured by concentrated power. In these forms, AI does not merely consume the life-ground. It reorganizes society around systems that may weaken the capacities of judgment, care, community, and responsibility that intelligence should serve.

The difference between these pathways is not determined by technology alone. It is determined by purpose, ownership, governance, design, scale, defaults, procurement, siting, energy sourcing, lifecycle responsibility, community consent, and the ability to refuse.

This is why efficiency alone is insufficient. Better chips, better cooling, better model routing, and lower energy per task are necessary, but they do not guarantee life-coherence. If efficiency lowers cost and expands total use, it can accelerate the very burden it was meant to reduce. Without sufficiency, efficiency becomes growth's servant. Without purpose, innovation becomes escalation. Without transparency, accountability

becomes impossible. Without community power, sustainability becomes burden shifting. Without commons governance, AI becomes enclosure.

The governing question must therefore change. The growth question asks: How far can AI scale? The life-coherent question asks: What forms of AI are worthy of scaling because they expand life-capacity within ecological limits? The growth question asks: How can AI become cheaper, faster, and more ubiquitous? The life-coherent question asks: How can AI become accountable, proportionate, sufficient, and subordinate to the commons of life? The growth question asks: What can be automated? The life-coherent question asks: What should remain human, relational, embodied, situated, and answerable?

These questions do not reject technology. They restore discernment. They recognize that intelligence is not measured by output alone. True intelligence is the capacity to align knowledge, action, care, and limits with the conditions of life.

For Caribbean and other Small Island Developing States, the life-coherence challenge is especially concrete. The life-ground is not abstract. It is the watershed, the power grid, the coastline, the reef, the school, the clinic, the household, the cistern, the farm, the road, the community, the cultural memory, and the fragile ecology of island life. AI may serve these realities if governed as a bounded tool and regional commons. It may harm them if imported as dependency, spectacle, or infrastructure burden.

The task is therefore not to become anti-AI or pro-AI. The task is to become life-coherent.

A life-coherent society does not ask whether a machine can generate more. It asks whether more generation serves life. It does not confuse prediction with wisdom. It does not confuse automation with care. It does not confuse scale with value. It does not confuse fluency with truth. It does not confuse intelligence with the absence of consequence.

The hidden life-ground of AI teaches a simple but profound lesson: there is no intelligence outside life. Every system that thinks, calculates, predicts, or generates depends on a world that must remain livable. The measure of artificial intelligence is therefore not only what it can produce, but what it helps preserve, repair, and bring forth.

If AI is governed by the logic of unbounded symbolic expansion, it will convert increasing portions of the life-ground into output, dependency, and enclosure. If AI is governed by life-coherence, it can become a bounded tool within a civil commons: useful, accountable, proportionate, transparent, and directed toward shared flourishing.

The future of AI will not be decided only in laboratories, markets, or data centers. It will be decided in the moral grammar by which societies answer one question: Does this intelligence serve life?

If the answer is no, restraint is wisdom. If the answer is yes, governance is responsibility. And if the answer is uncertain, the life-ground must be allowed to speak before the machine continues.

References

- Aczel, M., Chamanara, S., Matin, M., Farsi, A., Marwala, T., & Madani, K. (2026). *Environmental cost of AI's energy use: Carbon, water and land footprints*. United Nations University Institute for Water, Environment and Health. <https://doi.org/10.53328/INR26RMA002>
- Baldé, C. P., Kuehr, R., Yamamoto, T., McDonald, R., D'Angelo, E., Althaf, S., Bel, G., Deubzer, O., Fernandez-Cubillo, E., Forti, V., Gray, V., Herat, S., Honda, S., Iattoni, G., Khetriwal, D. S., Luda di Cortemiglia, V., Lobuntsova, Y., Nnorom, I., Pralat, N., & Wagner, M. (2024). *The global e-waste monitor 2024*. International Telecommunication Union and United Nations Institute for Training and Research. <https://ewastemonitor.info/the-global-e-waste-monitor-2024/>
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Crawford, K., & Joler, V. (2018). *Anatomy of an AI system: The Amazon Echo as an anatomical map of human labor, data and planetary resources*. AI Now Institute and Share Lab. <https://anatomyof.ai/>
- Dodge, J., Prewitt, T., Tachet des Combes, R., Odmark, E., Schwartz, R., Strubell, E., Luccioni, A. S., Smith, N. A., DeCario, N., & Buchanan, W. (2022). Measuring the carbon intensity of AI in cloud instances. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1877–1894. <https://doi.org/10.1145/3531146.3533234>
- Floridi, L., & Cows, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>
- Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D., & Pineau, J. (2020). Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research*, 21(248), 1–43. <https://www.jmlr.org/papers/v21/20-312.html>
- International Energy Agency. (2025). *Energy and AI*. IEA. <https://www.iea.org/reports/energy-and-ai>
- Intergovernmental Panel on Climate Change. (2022). *Climate change 2022: Impacts, adaptation and vulnerability. Working Group II contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press. <https://doi.org/10.1017/9781009325844>
- Lacoste, A., Luccioni, A. S., Schmidt, V., & Dandres, T. (2019). *Quantifying the carbon emissions of machine learning*. arXiv. <https://arxiv.org/abs/1910.09700>
- Li, P., Yang, J., Islam, M. A., & Ren, S. (2023). *Making AI less “thirsty”: Uncovering and addressing the secret water footprint of AI models*. arXiv. <https://arxiv.org/abs/2304.03271>
- Luccioni, A. S., Viguier, S., & Ligozat, A.-L. (2023). Estimating the carbon footprint of BLOOM, a 176B parameter language model. *Journal of Machine Learning Research*, 24(253), 1–15. <https://www.jmlr.org/papers/v24/23-0069.html>
- Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986. <https://doi.org/10.1126/science.aba3758>
- Maturana, H. R., & Varela, F. J. (1992). *The tree of knowledge: The biological roots of human understanding* (Rev. ed.). Shambhala.
- McMurtry, J. (2002). *Value wars: The global market versus the life economy*. Pluto Press.
- McMurtry, J. (2009–2011). What is good? What is bad? The value of all values across time, place and theories. In *Philosophy and world problems*. UNESCO–Encyclopedia of Life Support Systems.
- McMurtry, J. (2013). *The cancer stage of capitalism: From crisis to cure* (2nd ed.). Pluto Press.

- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511807763>
- Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). *Carbon emissions and large neural network training*. arXiv. <https://arxiv.org/abs/2104.10350>
- Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650. <https://doi.org/10.18653/v1/P19-1355>
- United Nations Development Programme. (2024). *Small island digital states: How digital can catalyse SIDS development*. UNDP. <https://www.undp.org/publications/small-island-digital-states-how-digital-can-catalyse-sids-development>
- United Nations Office of the High Representative for the Least Developed Countries, Landlocked Developing Countries and Small Island Developing States. (n.d.). *About Small Island Developing States*. United Nations. <https://www.un.org/ohrlls/content/about-small-island-developing-states>

Appendix A. Life-Coherent AI Use Checklist

1. Is AI genuinely needed?
2. What life-capacity does the use expand?
3. Is the smallest adequate model being used?
4. Is the lightest adequate modality being used?
5. Are long text, image, or video outputs truly necessary?
6. What are the carbon, water, land, and energy implications?
7. What data center, grid, water, and land systems are involved?
8. What minerals, hardware, labor, and e-waste systems support the use?
9. Who benefits?
10. Who bears the burden?
11. Who decides?
12. Are affected communities informed and protected?
13. Does the system strengthen local capacity or deepen dependency?
14. Does it support knowledge, learning, and judgment?
15. Is there transparency, accountability, appeal, review, and repair?
16. Can the system be paused, redesigned, or withdrawn?
17. Does the use contribute to the commons of life?
18. Does it expand life-capacity within ecological limits?

Decision options: proceed; proceed with limits; redesign; refuse or withdraw.

Appendix B. Figure Source Note

Unless otherwise stated, figures in this white paper were conceptually developed by the author and created with AI-assisted image generation in 2026. The figures are interpretive diagrams based on the life-coherence framework developed in this paper and are intended as conceptual aids rather than empirical data visualizations.

Appendix C. About the Author

Dr. Bichara Sahely is a physician and independent scholar based in St. Kitts and Nevis. His work integrates medicine, public health, ecology, peace studies, political economy, systems thinking, spirituality, and technology ethics into a life-coherence framework. Through bsahely.com and the Life-Knowledge Commons, he develops essays, white papers, audiovisual resources, and practical tools for evaluating whether human systems expand or diminish the conditions of life.