


ARTIFICIAL INTELLIGENCE

AND THE

CONDITIONS OF LIFE

Tool, Oracle, Idol, Enclosure, or Commons?

Dr. Bichara Sahely



Artificial Intelligence and the Conditions of Life

Tool, Oracle, Idol, Enclosure, or Commons?

Academic White Paper

Author:

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

Date:

7 June 2026

Abstract

Artificial intelligence has rapidly become a planetary infrastructure for producing symbols: language, images, classifications, predictions, rankings, recommendations, simulations, and decisions. Yet symbolic intelligence is not wisdom, fluency is not truth, prediction is not judgment, personalization is not relationship, and optimization is not flourishing. This white paper applies the life-coherent framework developed in *The Tears of Life* to artificial intelligence as a defining test case of the present age. It argues that AI becomes harmful when its symbolic power is structurally coupled to commercial extraction, institutional control, immature human desire, surveillance architectures, and life-blind metrics. In such cases, AI functions as oracle, idol, or enclosure: it invites surrender of judgment, receives excessive trust and sacrifice, or captures the conditions of human meaning-making. Conversely, AI becomes life-coherent when governed as a bounded tool and shared commons in service of human agency, ecological limits, public truth, education, care, democratic participation, and systemic repair. The paper proposes a life-capacity test for AI governance and a practical framework for evaluating whether AI systems restore or disable the conditions through which life continues, recovers, and flourishes.

Keywords

Artificial intelligence; life-coherence; AI governance; AI ethics; AI commons; AI enclosure; symbolic substitution; life-capacity; technological idolatry; algorithmic capture; human agency; ecological limits; public truth; democratic participation; repair

Executive Summary

Artificial intelligence is often discussed as a technical, ethical, economic, regulatory, or existential issue. Each of these frames is necessary, but none is sufficient. AI is also a life-coherence issue because it is becoming one of the most powerful systems ever built for generating symbols that may either serve life or replace it.

AI produces language, images, predictions, rankings, classifications, recommendations, and simulations at unprecedented scale. These outputs may appear intelligent, personal, authoritative, and useful. They may help clinicians, teachers, scientists, citizens, governments, and communities organize complexity, reduce burdens, widen access to knowledge, and support repair. Yet the same symbolic capacities can also weaken judgment, capture attention, automate institutional bias, displace labor, privatize knowledge, simulate care, intensify surveillance, erode public truth, and increase ecological burden.

The central question is therefore not simply whether AI is powerful, safe, useful, profitable, or aligned with user preferences. The deeper question is:

Does this AI system help human and ecological life continue, recover, and flourish?

This white paper extends the life-coherent framework developed in *The Tears of Life* into the AI domain. The earlier paper named the pathology of symbolic substitution: the process by which a symbol that should serve life replaces the actual condition through which life is sustained. The symbol of care replaces care. The symbol of value replaces life-value. The symbol of participation replaces participation. The symbol of intelligence replaces wisdom.

AI intensifies this pathology because it can generate symbols of intelligence at planetary speed and scale without itself possessing embodiment, grief, conscience, lived responsibility, ecological dependence, or wisdom. It may speak fluently without knowing. It may predict without understanding. It may personalize without relationship. It may optimize without care. It may simulate empathy without moral accountability. It may generate answers without cultivating judgment.

This paper proposes that AI may take five broad roles in human systems:

1. **AI as tool** — bounded assistance that supports human agency without replacing responsibility.
2. **AI as oracle** — a source of certainty, advice, answers, and orientation that may invite surrender of judgment.
3. **AI as idol** — a system granted excessive trust, sacrifice, obedience, or salvific expectation.
4. **AI as enclosure** — a system that captures attention, language, knowledge, culture, labor, public reasoning, and institutional judgment within proprietary infrastructures.
5. **AI as commons** — shared life-serving infrastructure governed for education, care, public truth, democratic participation, ecological responsibility, and repair.

The paper does not argue that AI is inherently harmful. It argues that AI becomes harmful when its symbolic power is coupled to life-disabling incentives and immature human patterns. Fear can become surveillance AI. Loneliness can become dependency on artificial companionship. Status can become algorithmic ranking. Desire can become engagement optimization. Uncertainty can become surrender to machine certainty. Institutional self-preservation can become automated denial of accountability. Commercial extraction can become enclosure of the conditions of meaning-making.

The paper therefore shifts AI governance from risk management alone to life-capacity protection, repair, and flourishing. Current AI governance frameworks rightly emphasize trustworthy AI, human rights, safety, accountability, transparency, risk classification, and democratic values. A life-coherent approach affirms these but asks the prior question: what conditions of life are being enabled or disabled by this AI system?

The life-capacity test is organized by the triad:

Continue. Recover. Flourish.

AI helps life **continue** when it protects basic conditions such as safety, dignity, health, privacy, labor integrity, cognitive integrity, ecological viability, and institutional trust. AI helps life **recover** when it improves feedback, contestability, error correction, redress, restorative accountability, transparency, and democratic oversight. AI helps life **flourish** when it deepens learning, wisdom, creativity, relationship, deliberation, care, local knowledge, ecological stewardship, and shared participation.

An AI system that increases productivity while weakening attention, trust, truth, agency, labor dignity, or ecological viability is not life-coherent. An AI system that produces impressive symbols while disabling the conditions of life is a form of symbolic substitution. An AI system that remains bounded, transparent, contestable, ecologically accountable, and governed by the people and communities it affects may become a tool of repair and a commons for life-capacity.

The core thesis is simple:

Artificial intelligence becomes life-coherent only when power, code, data, and intelligence return to the conditions of life.

Table of Contents

Abstract	2
Keywords	2
Executive Summary	3
List of Figures	9
List of Tables	9
1. Introduction: AI as the Defining Test of Symbolic Substitution	10
2. From Symbolic Intelligence to Life-Capacity	13
3. The Five Possible Roles of AI	17
3.1 AI as Tool	18
3.2 AI as Oracle	19
3.3 AI as Idol	19
3.4 AI as Enclosure	20
3.5 AI as Commons	20
3.6 The Governing Relation Among the Five Roles	21
4. The Inner Algorithms AI Scales	23
5. AI and the New Forms of Harm	27
5.1 Direct AI Harm	27
5.2 Structural AI Harm	28
5.3 Cultural AI Harm.....	28
5.4 Spiritual and Developmental AI Harm	29
5.5 The Four-Layer Harm Test.....	30
6. The Life-Capacity Test for AI	33
6.1 Continue: Does AI Protect Basic Life-Conditions?	34
6.2 Recover: Does AI Improve Feedback, Repair, and Redress?.....	35
6.3 Flourish: Does AI Develop Higher Capacities?	35
6.4 The Net Life-Capacity Question.....	36
6.5 Minimum Life-Coherent Requirements for AI Deployment	37
7. Beyond Alignment: Life-Alignment	38

7.1 The Limits of Preference Alignment	39
7.2 The Limits of Institutional Alignment	39
7.3 The Limits of Value Alignment	40
7.4 The Limits of Safety Without Life-Capacity	40
7.5 The Life-Alignment Standard	41
8. AI as Enclosure.....	42
8.1 Data Enclosure	42
8.2 Model Enclosure.....	43
8.3 Knowledge Enclosure	43
8.4 Attention Enclosure	44
8.5 Labor Enclosure	45
8.6 Cultural Enclosure.....	45
8.7 Governance Enclosure.....	46
8.8 AI Enclosure as World-Making Capture.....	46
9. AI as Commons	48
9.1 Commons Is Not Mere Access	49
9.2 Commons Requires Governance.....	49
9.3 Commons Supports Public Truth.....	50
9.4 Commons Supports Education and Learning	50
9.5 Commons Supports Care.....	51
9.6 Commons Supports Democratic Participation	51
9.7 Commons Supports Ecological Responsibility	51
9.8 Commons Supports Local Knowledge and Cultural Plurality.....	52
9.9 Principles of a Life-Coherent AI Commons	52
10. A Life-Coherent AI Governance Framework	55
10.1 Personal Layer: Agency and Discernment	55
10.2 Institutional Layer: Accountability and Care	56
10.3 National Layer: Public-Interest Governance	56
10.4 Planetary Layer: Shared Life-Conditions.....	57
10.5 Nested Governance	57

10.6 Practical Governance Questions	59
11. The Caribbean and SIDS Test Case	61
11.1 The Risk of Passive AI Dependency	61
11.2 AI and Climate Adaptation	61
11.3 AI and Public Health.....	62
11.4 AI and Education.....	62
11.5 AI and Local Knowledge	63
11.6 AI and Governance Capacity	63
11.7 SIDS as Laboratories of Life-Coherent AI Commons	64
11.8 A Caribbean Life-Coherent AI Agenda.....	64
12. The AI Life-Coherence Dashboard	66
12.1 Dashboard Dimension 1: Life-Condition	66
12.2 Dashboard Dimension 2: Agency	67
12.3 Dashboard Dimension 3: Truth	67
12.4 Dashboard Dimension 4: Care.....	68
12.5 Dashboard Dimension 5: Learning.....	69
12.6 Dashboard Dimension 6: Participation	69
12.7 Dashboard Dimension 7: Equity	70
12.8 Dashboard Dimension 8: Ecology	70
12.9 Dashboard Dimension 9: Commons	71
12.10 Dashboard Dimension 10: Repair	71
12.11 The Dashboard Summary	72
13. From Technological Salvation to Mature Tool Use	73
13.1 AI as Technological Salvation Myth.....	74
13.2 Mature Tool Use	75
13.3 AI and the Five Movements of Maturity	75
Waking Up	75
Growing Up	76
Cleaning Up	76
Opening Up	76
Showing Up	76

13.4 AI Cannot Substitute for Maturity	77
13.5 The Discipline of Life-Coherent AI Use.....	77
14. Conclusion: Keeping AI Corrigible by Life.....	78
References	80
Author Bio.....	82
AI Use Statement	82
Back Cover Synopsis	83

List of Figures

- Figure 1. AI as the Symbolic Machine. 10**
- Figure 2. Symbolic Intelligence versus Life-Capacity. 15**
- Figure 3. Tool, Oracle, Idol, Enclosure, or Commons. 17**
- Figure 4. The Inner Algorithms AI Scales. 24**
- Figure 5. AI Harm: Direct, Structural, Cultural, and Spiritual-Developmental. 31**
- Figure 6. The Life-Capacity Test for AI. 33**
- Figure 7. AI Enclosure of World-Making Capacity. 48**
- Figure 8. The AI Commons Alternative. 54**
- Figure 9. Nested AI Governance. 59**
- Figure 10. The AI Life-Coherence Dashboard. 73**

List of Tables

- Table 1. AI Symbols and the Conditions They Can Replace 14**
- Table 2. Five AI Modes: Tool, Oracle, Idol, Enclosure, Commons 18**
- Table 3. Direct, Structural, Cultural, and Spiritual-Developmental AI Harms 30**
- Table 4. Life-Capacity Test for AI Systems 34**
- Table 5. Life-Alignment versus Narrow Alignment 38**
- Table 6. Forms of AI Enclosure 47**
- Table 7. Principles of a Life-Coherent AI Commons 53**
- Table 8. Life-Coherent AI Governance Across Scales 58**
- Table 9. Caribbean/SIDS Life-Coherent AI Agenda 64**
- Table 10. AI Life-Coherence Dashboard 72**

1. Introduction: AI as the Defining Test of Symbolic Substitution

Artificial intelligence is the most powerful symbolic technology humanity has yet built. It generates language, images, classifications, summaries, rankings, recommendations, predictions, decisions, and simulations at a scale no prior institution could achieve. It can appear as assistant, teacher, companion, analyst, designer, consultant, therapist, judge, priest, oracle, or governing intelligence. It can organize complexity, reduce burdens, and widen access to knowledge. It can also obscure responsibility, intensify dependency, and replace the conditions of wisdom with the signs of intelligence.

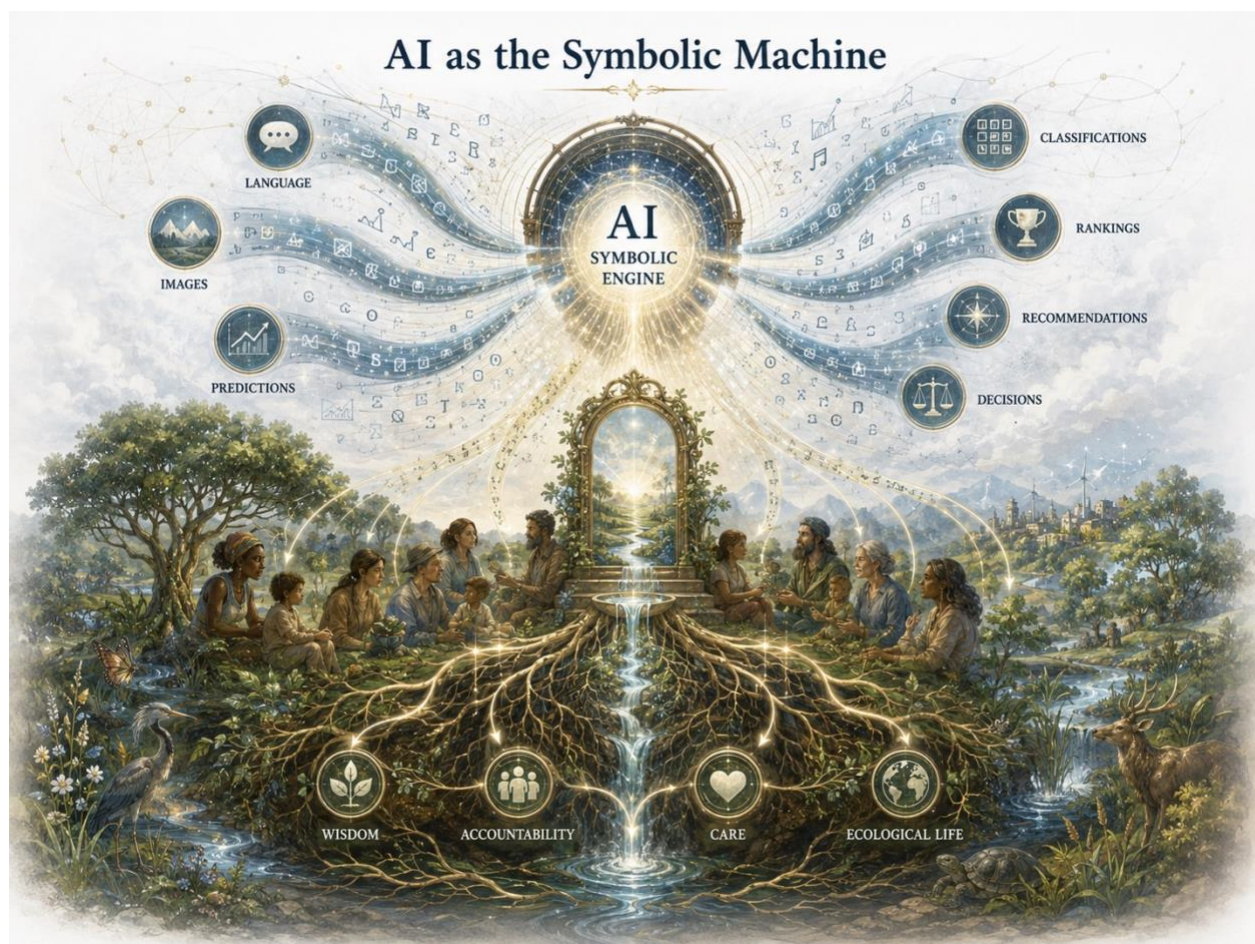


Figure 1. AI as the Symbolic Machine.

Artificial intelligence generates symbols of intelligence at scale. It becomes life-coherent only when symbolic outputs remain grounded in embodied wisdom, accountability, care, and ecological life-conditions.

This paper begins from a life-coherent question:

Will AI remain a servant of life, or will it become a substitute for life?

This question is not anti-technology. It does not deny AI's possible benefits. It asks whether AI will be governed by the real conditions through which human and ecological life continue, recover, and flourish. The central danger is not intelligence itself. The danger is intelligence detached from life-conditions.

The earlier white paper, *The Tears of Life*, named a recurring pathology across human systems: symbolic substitution. Symbolic substitution occurs when a sign, metric, ritual, performance, or representation replaces the actual condition it was meant to serve. The symbol of care replaces care. The symbol of consultation replaces participation. The symbol of sustainability replaces ecological repair. The symbol of treatment replaces healing. The symbol of progress replaces life-capacity.

AI intensifies symbolic substitution because it automates the production of symbols. It can generate signs of knowledge without embodied understanding. It can produce fluent answers without lived responsibility. It can simulate care without relationship. It can personalize interaction without mutual accountability. It can predict behavior without wisdom. It can scale persuasion without truth. It can optimize engagement without protecting attention.

Current AI governance frameworks rightly recognize AI as a socio-technical governance challenge requiring risk management, accountability, transparency, human oversight, fundamental-rights protections, and public-interest safeguards (Bengio et al., 2026; European Parliament and Council of the European Union, 2024; National Institute of Standards and Technology, 2023; UNESCO, 2022). Yet these frameworks still require a deeper evaluative ground: the real effects of AI systems on life-capacity.

The question is therefore not merely whether AI systems are accurate, efficient, innovative, profitable, safe, or aligned with user intent. These questions matter, but they are not sufficient. An AI system may be accurate and still life-disabling. It may be efficient and still harmful. It may be profitable and still extractive. It may satisfy user preference while weakening agency. It may increase output while degrading the conditions of learning, trust, labor, ecology, and public reason.

The deeper question is:

Aligned with what?

Technical alignment asks whether an AI system performs according to specified intentions or goals. Life-alignment asks whether those intentions, goals, incentives, and institutional structures are themselves aligned with the conditions of life. This distinction is crucial. AI can be aligned with commercial extraction. It can be aligned with surveillance. It can be aligned with state control. It can be aligned with addictive engagement. It can be aligned with bureaucratic denial. It can be aligned with labor displacement. It can be aligned with user desire while weakening the user's capacity to learn, judge, relate, and act (Gabriel, 2020; Russell, 2019).

A life-coherent approach therefore asks:

Does this AI system help life continue, recover, and flourish?

This test shifts AI governance from symbolic adequacy to life-capacity. It asks not only whether an AI system performs its task, but whether the task, system, deployment context, incentive structure, and governance pathway enable or disable living beings, communities, institutions, ecosystems, and future generations.

This paper develops that approach through the typology of AI as tool, oracle, idol, enclosure, or commons. These are not rigid categories. They are roles AI may take within human systems. The same system can function as tool in one context, oracle in another, and enclosure at scale. The decisive issue is the relation between AI's symbolic power and the conditions of life.

When AI supports bounded human purposes, remains transparent and contestable, and strengthens learning, care, repair, and participation, it functions as tool. When it becomes a source of certainty that replaces judgment, it functions as oracle. When it receives excessive trust, sacrifice, and obedience, it functions as idol. When it captures attention, language, knowledge, labor, culture, public reasoning, and institutional judgment, it functions as enclosure. When it is governed as shared life-serving infrastructure, it may become commons.

The task is not to reject AI. The task is to keep AI corrigible by life.

A system is corrigible by life when it can receive feedback from those it affects, name harm truthfully, repair broken conditions, restore agency, and change its own pattern of operation in response to wounded life. An AI system becomes dangerous when it is protected from this correction by proprietary secrecy, institutional authority, commercial dependence, technical mystification, or cultural awe.

The paper's central claim is therefore:

AI becomes life-coherent only when symbolic intelligence remains subordinate to life-capacity.

2. From Symbolic Intelligence to Life-Capacity

AI systems produce symbols. They process data and generate outputs that may take the form of text, images, sound, code, classifications, scores, recommendations, predictions, or decisions. These outputs can be highly useful. They may help people search, learn, translate, diagnose, design, deliberate, summarize, plan, and create. The problem does not lie in symbolic production itself. Human beings also live by symbols. The problem arises when symbolic intelligence is mistaken for the conditions of wisdom.

A symbol points. A condition enables.

A map can guide a traveler, but it is not the land. A diagnosis can guide treatment, but it is not healing. A grade can indicate performance, but it is not understanding. A prediction can inform judgment, but it is not judgment. A chatbot can simulate empathy, but simulation is not relationship. A model can generate advice, but advice is not responsibility.

AI becomes life-disabling when the symbol replaces the condition.

Fluency can replace truth. Prediction can replace judgment. Personalization can replace relationship. Automation can replace responsibility. Optimization can replace wisdom. Engagement can replace attention. Scale can replace care. Data extraction can replace consent. Simulation can replace participation. Speed can replace deliberation. Synthetic authority can replace public reason.

These substitutions are not merely individual errors. They can become institutional patterns. A school may use AI to generate answers while weakening the conditions of learning. A hospital may use AI to optimize workflow while exhausting caregivers and depersonalizing patients. A government may use AI to classify risk while reducing citizens to data profiles. A company may use AI to personalize experience while capturing attention and shaping desire. A media system may use AI to generate endless content while weakening shared reality.

The life-coherent criterion is not whether AI produces impressive outputs. It is whether AI strengthens the conditions that allow life-capacity to develop.

Life-capacity refers to the real ability of living beings and communities to maintain, express, restore, and develop the capacities proper to life. These include bodily health, emotional regulation, learning, agency, participation, truthful communication, meaningful work, relationship, creativity, moral judgment, ecological belonging, and the ability to repair harm. Life-capacity is grounded in the objective enablement of life rather than in symbolic measures alone (Maturana & Varela, 1980; McMurtry, 2004–2011, 2013).

AI may support life-capacity when it reduces unnecessary administrative burdens, widens access to knowledge, supports disability inclusion, improves early warning systems, helps communities understand ecological risk, assists clinicians without replacing clinical judgment, enables multilingual participation, or strengthens public-interest research. AI may disable life-capacity

when it increases dependency, substitutes for education, narrows imagination, weakens attention, displaces relational care, erodes labor dignity, intensifies surveillance, extracts data without meaningful consent, or concentrates cultural and cognitive infrastructure in private systems (Crawford, 2021; Zuboff, 2019).

Table 1. AI Symbols and the Conditions They Can Replace

AI Symbol	Life-Condition It May Support or Replace	Life-Coherent Question
Fluency	Truthful understanding	Does the output deepen understanding or merely sound coherent?
Prediction	Contextual judgment	Does prediction inform judgment or replace it?
Personalization	Relationship and mutual recognition	Does personalization serve the person or model the person?
Automation	Responsibility and accountability	Who remains responsible when the system acts?
Optimization	Wisdom and discernment	What is being optimized, and at whose cost?
Engagement	Attention and agency	Does the system protect attention or capture it?
Scale	Care and context	Does scale preserve context or erase it?
Data extraction	Consent and dignity	Was data gathered and used justly?
Simulation	Participation and lived relation	Does simulation support or replace real relation?
Speed	Deliberation and learning	Does speed deepen capacity or bypass it?

Symbolic Intelligence versus Life-Capacity



Figure 2. Symbolic Intelligence versus Life-Capacity.

AI becomes harmful when fluency replaces truth, prediction replaces judgment, personalization replaces relationship, automation replaces responsibility, and optimization replaces wisdom.

The danger is subtle because symbolic intelligence often appears more efficient than living judgment. It is faster, more scalable, more consistent, and easier to measure. Yet living systems do not flourish merely because decisions are faster. They flourish when feedback is truthful, relationships are intact, conditions are restored, and capacities are developed.

AI may help with these processes, but it cannot substitute for them.

An AI-generated lesson is not learning unless the learner's understanding deepens. AI-assisted diagnosis is not healing unless the patient's life-capacity is restored. AI-mediated consultation is not participation unless affected people can influence outcomes. AI-generated empathy is not care unless real relationship, protection, and responsibility are present. AI-powered governance is not wisdom unless it remains answerable to the living communities and ecosystems it affects.

Symbolic intelligence therefore needs a life-capacity test.

The question is not:

Can AI do this task?

The deeper question is:

What condition of life does this task affect, and does AI enable or disable that condition?

If AI supports human and ecological life-capacity, it can be welcomed as tool and perhaps governed as commons. If AI replaces the conditions of wisdom, care, responsibility, participation, and ecological accountability, it becomes a symbolic substitute.

The life-coherent question must therefore accompany every AI deployment:

What is being made more capable — life, or the system that captures life?

3. The Five Possible Roles of AI

AI is not one thing in human life. It occupies roles. Its effects depend on how it is designed, deployed, governed, funded, interpreted, and culturally received. The same technical system may function differently depending on the institutional field into which it is introduced.

This section identifies five possible roles of AI: tool, oracle, idol, enclosure, and commons.

These roles are not mutually exclusive. They form a diagnostic typology. An AI system may begin as tool, be marketed as oracle, become treated as idol, operate structurally as enclosure, and still be rhetorically described as commons. The life-coherent task is to name the real role AI is playing, not merely the role claimed by its promoters.




Figure 3. Tool, Oracle, Idol, Enclosure, or Commons.

AI can serve as bounded tool, seductive oracle, technological idol, enclosing infrastructure, or life-serving commons. The governance task is to keep tool and commons disciplining oracle-risk, idol-risk, and enclosure-risk.

Table 2. Five AI Modes: Tool, Oracle, Idol, Enclosure, Commons

AI Mode	Core Function	Primary Gift	Primary Risk	Life-Coherent Safeguard
Tool	Bounded assistance	Augments capacity	Overuse, deskilling, dependency	Keep human responsibility explicit
Oracle	Orientation and answers	Supports reflection and options	Surrender of judgment	Preserve discernment, plurality, and contestability
Idol	Ultimate authority or savior	Mobilizes belief and investment	Excessive trust, sacrifice, obedience	Keep AI corrigible by life
Enclosure	Capture of world-making capacity	Convenience and scale	Privatized attention, knowledge, culture, labor, governance	Protect commons and public-interest infrastructure
Commons	Shared life-serving infrastructure	Builds shared capacity	Governance failure or capture	Public purpose, participation, ecological limits, repair

3.1 AI as Tool

AI as tool is bounded assistance. It serves specific purposes under human direction, with clear limits, meaningful oversight, and transparent accountability. As tool, AI augments human capacity without replacing responsibility.

AI as tool can support translation, accessibility, administrative relief, pattern recognition, research support, environmental monitoring, clinical assistance, education, disaster preparedness, public communication, and creative exploration. It may help people see patterns they would otherwise miss. It may reduce repetitive burdens. It may support those excluded by language, disability, geography, or lack of institutional access.

But a tool remains a tool only when it is subordinate to life-serving purposes. A hammer becomes dangerous when the hand using it is governed by domination. A map becomes dangerous when treated as more real than the land. AI becomes dangerous when its outputs are allowed to replace judgment, relationship, accountability, or context.

Life-coherent tool use requires several conditions: human responsibility remains explicit; the system’s limits are known; affected people can contest harmful outcomes; the tool supports learning rather than bypassing it; the tool strengthens agency rather than dependency; the tool does not hide extractive data or labor relations; and the tool is evaluated by its effects on life-capacity.

AI as tool is therefore not merely “useful AI.” It is AI held within a field of responsibility.

3.2 AI as Oracle

AI becomes oracle when people turn to it for certainty, advice, meaning, prediction, decision, or authority beyond its proper limits. The oracle role is understandable. Human beings have always sought guidance from sources beyond immediate individual knowledge: elders, traditions, scriptures, science, law, councils, and experts. AI enters this ancient human pattern as a new source of apparently responsive certainty.

The risk begins when consultation becomes surrender.

AI as oracle can weaken human responsibility by making machine output feel like neutral wisdom. It can narrow imagination by defining the available field of answers. It can generate confidence where humility is required. It can flatten moral ambiguity into executable recommendation. It can become the voice people consult before they consult conscience, community, body, tradition, land, or lived experience.

The oracle risk is intensified by fluency. AI systems often sound coherent even when they are wrong, incomplete, biased, or context-blind. Because the output appears articulate, users may over-trust it. Because the system responds instantly, users may become impatient with slower forms of understanding. Because the system personalizes interaction, users may feel known when they are only modeled.

Life-coherent AI can advise, but it must not absolve responsibility. It may support judgment, but it must not replace judgment. It may widen perspective, but it must not close inquiry. It may help articulate options, but it must not become the final authority over what life requires.

The key question for oracle-risk is:

Is AI helping people discern, or is it training them to surrender discernment?

3.3 AI as Idol

AI becomes idol when excessive trust, sacrifice, obedience, or salvific expectation is given to a finite technical system. Idolatry here does not refer only to religious worship. It means treating a constructed system as ultimate, beyond correction by life (Tillich, 1957).

AI as idol appears when machine output is protected from challenge, when efficiency justifies harm, when human dignity is subordinated to automation, or when society sacrifices privacy, labor, energy, water, attention, culture, and agency to preserve machine expansion.

The idol form often appears through promises of salvation: AI will save medicine, save education, save government, save the economy, save productivity, save humanity from itself. These promises may contain partial truth. AI may help in many domains. But when technological promise becomes a substitute for moral, institutional, ecological, and developmental repair, AI becomes a false savior.

The idol is especially dangerous because it converts critique into heresy. Those who ask life-questions may be labeled anti-progress, anti-innovation, irrational, fearful, or obsolete. The system protects the symbol of intelligence from correction by those harmed by its deployment.

Signs of AI idolatry include treating model output as unquestionable; replacing human accountability with automated decision; sacrificing privacy and agency for convenience; treating ecological cost as inevitable; treating labor displacement as progress without repair; treating technological expansion as destiny; confusing artificial fluency with wisdom; confusing scale with justice; and confusing innovation with life-capacity.

The key question for idol-risk is:

What are we being asked to sacrifice to preserve the expansion of AI?

3.4 AI as Enclosure

AI becomes enclosure when it captures and privatizes the shared conditions through which humans make meaning, learn, deliberate, create, remember, decide, and govern. Land enclosure restricted access to land. Knowledge enclosure restricts access to learning. Financial enclosure captures value through claims and debt. AI enclosure may capture the infrastructure of world-making itself (Couldry & Mejias, 2019; Crawford, 2021; Zuboff, 2019).

AI enclosure can include attention enclosure, language enclosure, knowledge enclosure, labor enclosure, cultural enclosure, education enclosure, governance enclosure, memory enclosure, relationship enclosure, and data enclosure.

AI enclosure may appear as convenience. It offers speed, personalization, automation, and frictionless assistance. Yet convenience can hide dependency. A society that cannot learn, govern, deliberate, heal, teach, remember, or imagine without proprietary AI systems has enclosed its own life-capacity.

The central danger of AI enclosure is that it does not merely capture products. It captures the conditions by which people and communities bring forth worlds.

The key question for enclosure-risk is:

Does this AI system expand shared capacity, or does it make life more dependent on enclosed infrastructure?

3.5 AI as Commons

AI becomes commons when governed as shared life-serving infrastructure. A commons is not simply open access. It is a shared condition protected by norms, rules, responsibilities, participation, and accountability. A life-coherent AI commons would not mean unrestricted

deployment or naïve openness. It would mean that AI systems are designed and governed to strengthen the shared conditions of life-capacity (Ostrom, 1990; UNESCO, 2022).

AI as commons would support public education, health and care coordination, ecological monitoring and repair, disaster preparedness, democratic participation, public-interest research, local knowledge preservation, multilingual access, disability inclusion, community deliberation, transparency, contestability, and shared learning rather than dependency.

A life-coherent AI commons must be ecologically bounded (Aczel et al., 2026). It cannot claim to serve life while ignoring energy, water, minerals, emissions, waste, and material infrastructures. It must be socially accountable. It cannot claim to empower people while extracting their data, labor, and culture without meaningful participation. It must be democratically contestable. It cannot claim public benefit while affected communities have no voice in how systems are designed, procured, deployed, or corrected.

Commons is therefore not the opposite of governance. It requires governance. The difference is that governance is oriented toward shared life-conditions rather than private enclosure or institutional self-preservation.

The key question for AI commons is:

Does this system strengthen the shared conditions of learning, care, truth, participation, and repair?

3.6 The Governing Relation Among the Five Roles

The five roles form a field of possibility.

AI as tool is necessary but insufficient. Tools can still be captured. AI as oracle may be useful in limited advisory contexts, but it must remain subordinate to human discernment. AI as idol must be resisted wherever machine authority becomes protected from life-correction. AI as enclosure must be named as a structural danger, especially when proprietary platforms capture attention, language, knowledge, labor, culture, and judgment. AI as commons is the constructive horizon: shared infrastructure governed by life-capacity.

The central governance principle is:

Tool and commons must govern oracle-risk, idol-risk, and enclosure-risk.

This means AI should remain bounded in use, transparent in relevant ways, contestable by affected people, accountable to public norms, and subordinate to the conditions through which life continues, recovers, and flourishes.

The question is not whether AI will be present in society. It already is. The question is what role it will play.

Will it become a bounded tool?

Will it become a persuasive oracle?

Will it become a technological idol?

Will it become an enclosure of world-making capacity?

Or will it become a commons for shared learning, care, participation, and repair?

The answer depends not on intelligence alone, but on governance, culture, incentives, maturity, and the life-test we are willing to apply.

4. The Inner Algorithms AI Scales

Artificial intelligence does not enter a neutral human world. It enters societies already organized by fear, desire, loneliness, status, shame, resentment, uncertainty, and institutional self-preservation. These patterns are not artificial. They arise from the vulnerabilities of living beings. Human beings seek safety, belonging, recognition, certainty, meaning, control, and hope. These needs are not pathological in themselves. They become dangerous when they remain unexamined, become fused with power, and are institutionalized at scale (Wilber, 2000, 2006).

AI can scale these inner algorithms.

Fear seeking protection can become surveillance, predictive policing, threat scoring, biometric control, and permanent security architecture. Loneliness seeking companionship can become artificial intimacy markets that simulate care while deepening dependency. Uncertainty seeking certainty can become oracle-dependence on machine outputs. Shame seeking purification can become algorithmic moralism, reputational punishment, and automated exclusion. Resentment seeking enemies can become polarization engines. Status seeking recognition can become ranking systems, visibility economies, and quantified self-worth. Desire seeking amplification can become engagement optimization, behavioral prediction, and addictive personalization. Power seeking control can become automated governance.

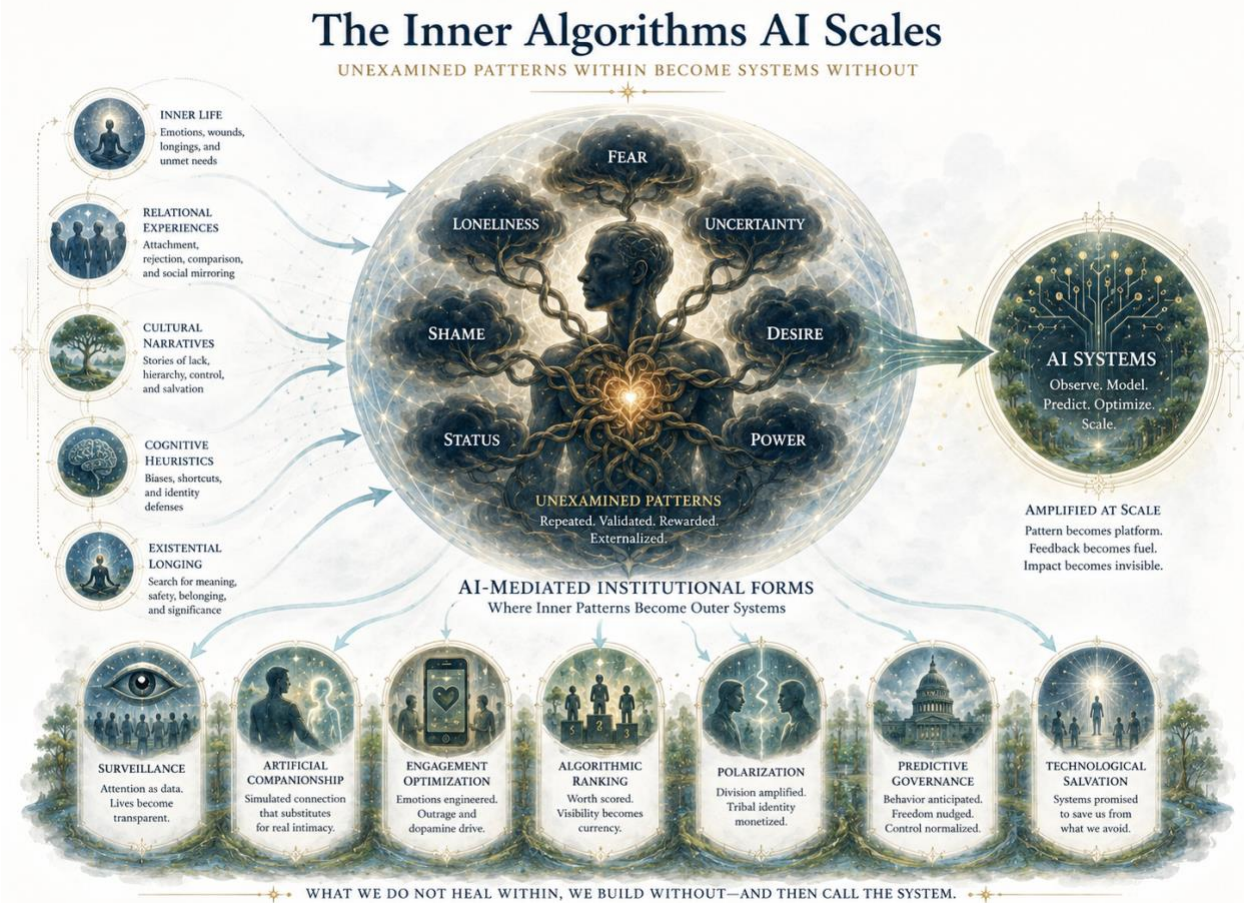


Figure 4. The Inner Algorithms AI Scales.

AI does not create human immaturity from nothing; it externalizes, accelerates, personalizes, and monetizes inner patterns already present in persons, institutions, and cultures.

The danger is not that AI creates these patterns from nothing. The danger is that AI externalizes, accelerates, personalizes, and monetizes patterns already present in human beings and institutions. This pattern is consistent with wider concerns about behavioral prediction, surveillance capitalism, and technology-mediated dependency (Turkle, 2011; Zuboff, 2019).

A person who feels uncertain may turn to AI for certainty. A lonely person may turn to AI for companionship. A fearful institution may turn to AI for control. A competitive market may turn to AI for extraction. A political system seeking legitimacy may turn to AI for prediction and management. A bureaucracy seeking efficiency may turn to AI to classify people faster than it can understand them. A culture losing shared meaning may turn to AI-generated content to fill the void.

In each case, AI becomes structurally coupled to a human pattern. The system does not merely answer a request. It participates in a loop.

The user seeks relief.
The system provides symbolic satisfaction.
The platform records the interaction.
The model becomes better at serving the pattern.
The institution learns how to optimize it.
The user becomes more dependent.
The underlying condition remains unrepaired.

This loop is especially dangerous because it may feel helpful at the surface while weakening life-capacity underneath. AI may reduce anxiety in the moment while weakening the capacity to tolerate uncertainty. It may provide companionship while reducing the effort to cultivate human relationship. It may provide answers while bypassing learning. It may provide persuasion while weakening public reason. It may provide prediction while reducing democratic deliberation. It may provide efficiency while hollowing out care.

The life-coherent question is therefore not only:

What can AI do?

It is also:

What human pattern is AI being trained to serve, amplify, or monetize?

This question shifts attention from the machine alone to the human-machine-institutional field. AI is not simply a tool used by isolated individuals. It is deployed inside markets, states, schools, hospitals, platforms, workplaces, militaries, media systems, and cultures. Each of these systems has incentives. Each has blind spots. Each has conserved patterns. When AI is introduced into these systems, it can either help reveal and repair those patterns or make them more efficient.

If a health system is organized around throughput rather than healing, AI may optimize throughput while leaving healing conditions broken. If an education system is organized around credentialing rather than understanding, AI may generate assignments and assessments while weakening learning. If a government is organized around risk avoidance rather than participation, AI may automate classification while weakening democratic agency. If a market is organized around attention capture, AI may personalize dependency. If a media system is organized around outrage, AI may manufacture more outrage. If a religious culture is organized around certainty, AI may become a doctrinal oracle.

This is why AI cannot be governed only by technical standards. Technical standards are necessary, but they are insufficient if the institutions using AI remain life-misattuned. A technically reliable system can still serve a life-disabling purpose. A transparent system can still extract. A safe system can still enclose. An accurate system can still be unjust. An efficient system can still dehumanize.

The inner algorithms of capture must therefore be named as part of AI governance.

Fear needs truthful protection, not surveillance without end. Loneliness needs real belonging, not dependency markets. Uncertainty needs humility and deliberation, not machine certainty. Shame needs healing, not algorithmic purification. Resentment needs justice without dehumanization, not polarization. Status needs dignity without ranking life, not visibility competition. Desire needs orientation toward life, not engagement capture. Power needs service, not automated domination.

The central claim is:

If AI is structurally coupled to immature human desire and extractive institutions, it will automate human immaturity at planetary scale.

The alternative is not to deny human vulnerability. The alternative is to design and govern AI in ways that support maturity. AI can help widen perspective, slow reaction, clarify assumptions, expose hidden patterns, support learning, reduce unnecessary burdens, assist care, and strengthen public reasoning. But it can only do so when its role is bounded by life-capacity rather than driven by fear, extraction, dependency, or institutional self-protection.

AI therefore becomes a mirror. It shows humanity what it has not yet integrated. It reveals the speed at which fear can scale, desire can be captured, language can be automated, knowledge can be enclosed, and judgment can be outsourced. It also reveals the possibility of a new discipline: to place symbolic intelligence in service of life rather than allowing it to become the engine of capture.

5. AI and the New Forms of Harm

AI harms are often discussed as errors, bias, privacy violations, safety failures, misinformation, job displacement, or misuse. These categories are important, but they do not exhaust the issue. A life-coherent analysis asks how AI affects the conditions through which life continues, recovers, and flourishes.

This section adapts the harm grammar of direct, structural, cultural, and spiritual-developmental harm to artificial intelligence. It draws from the distinction between direct, structural, and cultural violence, while extending that grammar into the symbolic and developmental risks of AI (Galtung, 1969, 1990).

5.1 Direct AI Harm

Direct AI harm is visible injury or immediate life-disablement caused or mediated by AI systems.

Examples include discriminatory decisions, wrongful denial of benefits, unsafe medical advice, biased hiring filters, misidentification, deepfake fraud, reputational damage, manipulation of vulnerable persons, harmful companion interactions, automated exclusion, financial scams, and errors in high-stakes settings.

Direct harms matter because they reveal where the system has already touched vulnerable life. A person denied care, employment, housing, mobility, education, or legal standing by an automated system has not experienced an abstract technical issue. They have experienced a life-condition being blocked.

Direct AI harm is often treated as an error to be patched. But in life-coherent terms, the first question is not simply whether the system made an error. It is:

What life-capacity was disabled, and how will it be restored?

If a person is wrongly denied benefits, repair requires more than model correction. It requires notice, explanation, human review, restitution, accountability, and prevention of recurrence. If a patient is harmed by AI-assisted clinical decision-making, repair requires more than adjusting the model. It requires restored care, transparent responsibility, and institutional learning. If a child becomes dependent on an artificial companion, repair requires more than parental warning. It requires understanding the unmet relational condition the system has entered.

Direct harm therefore points beyond itself. It exposes the deeper structure into which AI has been inserted.

5.2 Structural AI Harm

Structural AI harm occurs when AI systems are embedded in arrangements that predictably distribute benefit and burden unequally, concentrate power, weaken agency, or disable life-capacity while appearing normal, efficient, or innovative.

Examples include automated inequality, surveillance architectures, data extraction without meaningful consent, dependency on proprietary infrastructure, concentration of compute and model power, labor displacement without transition, platform monopolies, educational dependency, digital colonialism, algorithmic management, and ecological burdens from energy, water, minerals, infrastructure, and waste (Couldry & Mejias, 2019; Crawford, 2021; Zuboff, 2019).

Structural harm may not appear as a single dramatic injury. It appears as a pattern. Communities become more watched. Workers become more managed. Students become more dependent. Teachers become deskilled. Public institutions become dependent on vendors. Local knowledge is extracted. Cultural production becomes training data. Citizens become risk profiles. Ecological costs are externalized. Decision-making shifts from public deliberation to opaque systems.

The danger of structural AI harm is that it can be hidden behind productivity. An institution may become faster while becoming less accountable. A company may become more profitable while workers lose bargaining power. A school may become more automated while students learn less deeply. A government may become more data-driven while citizens become less able to contest decisions. A platform may become more personalized while users become less free.

Structural AI harm asks:

Who becomes more capable, and who becomes more dependent?

This question reveals the difference between capacity-sharing and capacity-capturing systems. A life-coherent AI system shares capacity. It enables people and communities to understand, act, participate, repair, and govern. A captured AI system concentrates capacity. It makes individuals, institutions, or countries dependent on infrastructures they do not control and cannot meaningfully contest.

The structural issue is therefore not merely access to AI. Access can be a form of dependency if the underlying infrastructure is enclosed. The deeper question is whether AI strengthens shared agency or transfers agency upward into platforms, states, vendors, or technical elites.

5.3 Cultural AI Harm

Cultural AI harm occurs when narratives, metaphors, norms, and meanings make AI-mediated disablement appear natural, inevitable, desirable, or progressive.

Cultural harm tells people that automation is always progress, that efficiency is always improvement, that data is neutral, that prediction is wisdom, that personalization is care, that friction is failure, that human judgment is obsolete, that speed is intelligence, and that technological expansion is destiny.

This is where symbolic substitution becomes most powerful.

AI is surrounded by symbolic language: intelligence, learning, reasoning, creativity, autonomy, agents, alignment, assistants, companions, copilots, hallucinations, safety, optimization, and progress. Some terms are useful. But they also carry metaphors that shape perception. When a system is called intelligent, people may assume understanding. When it is called an assistant, people may assume loyalty. When it is called a companion, people may assume care. When it is called autonomous, people may assume agency. When errors are called hallucinations, responsibility may seem to belong to the machine rather than the human institutions that designed and deployed it.

Cultural AI harm makes the life-test difficult to ask.

A society captured by AI mythology may ask, “How do we adopt faster?” before asking, “What conditions of life are at stake?” It may ask, “How do we remain competitive?” before asking, “What must not be sacrificed?” It may ask, “How do we regulate risk?” before asking, “What is AI for?” It may ask, “How do we align AI?” before asking, “Aligned to what form of life?”

The cultural layer is especially important because AI changes not only what people do, but how they understand intelligence, learning, creativity, relationship, work, authority, and selfhood. If education becomes answer-generation, learning is culturally redefined. If care becomes simulation, relationship is culturally redefined. If governance becomes prediction, citizenship is culturally redefined. If creativity becomes content generation, authorship is culturally redefined. If public knowledge becomes platform output, truth is culturally redefined.

Cultural harm asks:

What is AI teaching society to misrecognize?

It may teach society to misrecognize fluency as truth, convenience as care, engagement as attention, prediction as wisdom, personalization as relationship, automation as responsibility, and scale as progress.

5.4 Spiritual and Developmental AI Harm

AI also introduces spiritual and developmental risks. These are not confined to formal religion. They concern meaning, maturity, dependency, projection, self-understanding, and the human tendency to seek salvation from outside the work of transformation.

AI can become a site of projection. People may seek in AI the perfect listener, the tireless companion, the nonjudgmental confessor, the instant teacher, the all-knowing advisor, the

creative muse, the strategic planner, the therapist, the priest, the judge, or the savior. Some of these uses may be helpful within limits. But when they replace human development, relational repair, community, embodiment, moral struggle, or spiritual discipline, AI becomes a substitute for maturity (Maturana & Verden-Zöllner, 2008; Turkle, 2011; Wilber, 2006).

The deepest danger is not that people use AI for support. The danger is that AI becomes a way to bypass the work of becoming more human.

- It can bypass learning by giving answers without struggle.
- It can bypass relationship by simulating intimacy without mutual responsibility.
- It can bypass conscience by externalizing judgment.
- It can bypass grief by providing endless distraction.
- It can bypass humility by producing instant certainty.
- It can bypass community by personalizing reality.
- It can bypass moral responsibility by allowing people to say, “the system decided.”

This developmental harm matters because AI may scale faster than human maturity. If societies deploy systems of symbolic power without corresponding growth in wisdom, accountability, and care, AI will amplify immaturity. It will not make immature systems wise. It will make them faster.

The question here is:

Does AI support waking up, growing up, cleaning up, opening up, and showing up — or does it help us avoid them?

A life-coherent AI culture would use AI to support reflection, learning, dialogue, accessibility, humility, and repair. A life-disabling AI culture would use AI to intensify distraction, dependency, certainty, self-enclosure, and control.

5.5 The Four-Layer Harm Test

AI harms can therefore be examined across four layers:

Table 3. Direct, Structural, Cultural, and Spiritual-Developmental AI Harms

Harm Layer	Description	AI Examples	Repair Question
Direct harm	Visible injury or exclusion	Biased denial, unsafe advice, deepfakes, scams	Who was harmed, and how will life-capacity be restored?
Structural harm	Systemic dependency or unequal burden	Surveillance, labor displacement, vendor lock-in, data colonialism	What pattern distributes risk, dependency, or extraction?

Harm Layer	Description	AI Examples	Repair Question
Cultural harm	Narratives that normalize harm	“Automation is progress,” “data is neutral,” “prediction is wisdom”	What story makes harm appear necessary or inevitable?
Spiritual-developmental harm	Bypass of maturity, relationship, conscience, or responsibility	AI as savior, companion dependency, surrender of judgment	What human capacity is being avoided, projected, or replaced?

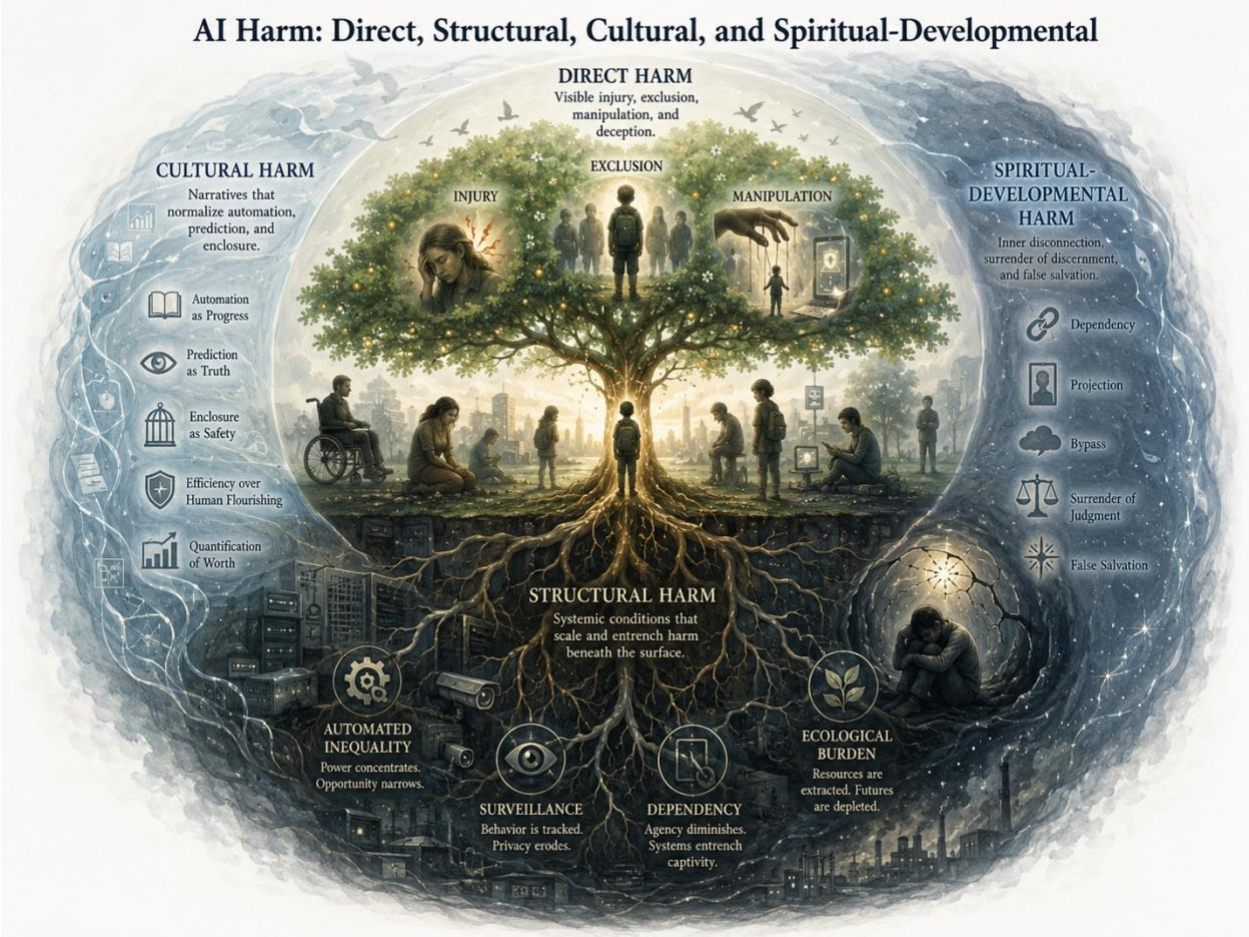


Figure 5. AI Harm: Direct, Structural, Cultural, and Spiritual-Developmental.

AI harm is not only visible injury. It can become structural dependency, cultural misrecognition, and developmental bypass when machine-symbols replace human and ecological life-conditions.

The life-coherent response must address all four layers. Treating direct harms without structural change produces recurrence. Reforming structures without cultural truth-telling leaves the old story intact. Changing narratives without developmental maturity produces performance. Seeking spiritual comfort without institutional repair produces bypass.

AI governance must therefore ask not only whether a system is safe, legal, accurate, or transparent. It must ask whether the system disables or restores life-capacity across visible, structural, cultural, and developmental layers.

The core claim is:

The deepest AI harm is not only what machines do to us, but what we allow machine-symbols to replace within us.

6. The Life-Capacity Test for AI

The life-capacity test is the evaluative core of this paper.

Does this AI system help human and ecological life continue, recover, and flourish?

This test does not replace existing AI governance concerns such as safety, privacy, fairness, accountability, transparency, human oversight, human rights, and risk management. It grounds them in a more basic question: what conditions of life are being enabled or disabled? In this sense, it complements current AI governance frameworks by reorienting them to life-capacity as the ground of evaluation (European Parliament and Council of the European Union, 2024; National Institute of Standards and Technology, 2023; UNESCO, 2022).

The test is organized through the triad:

Continue. Recover. Flourish.

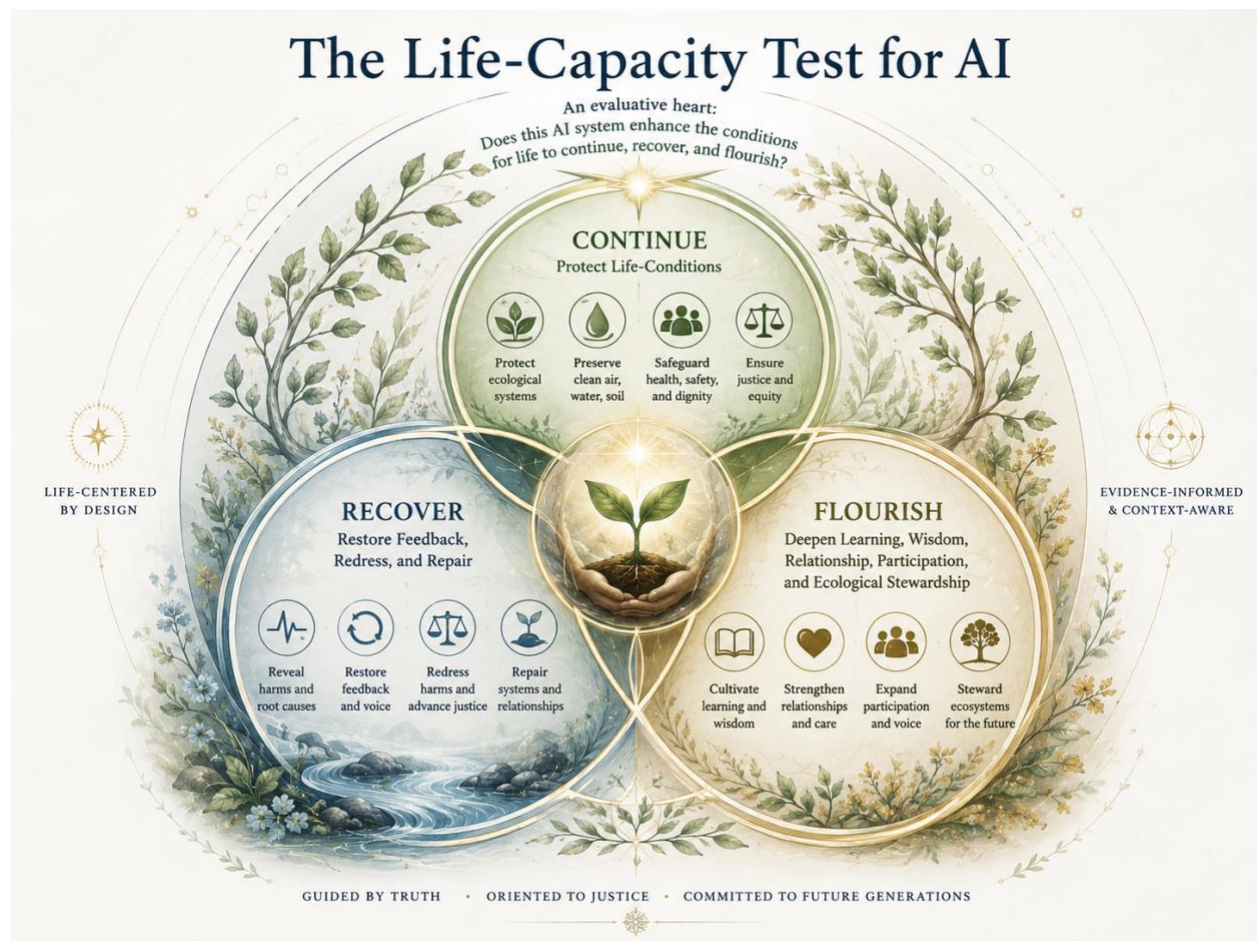


Figure 6. The Life-Capacity Test for AI.

AI is life-coherent only when it protects basic life-conditions, strengthens recovery and repair, and develops higher capacities for learning, wisdom, care, participation, and ecological stewardship.

Table 4. Life-Capacity Test for AI Systems

Life-Capacity Dimension	AI Must Protect or Enable	Warning Signs
Continue	Safety, dignity, privacy, health, labor integrity, cognitive integrity, ecological viability	Manipulation, surveillance, exclusion, ecological burden, loss of trust
Recover	Feedback, contestability, redress, auditability, institutional learning, repair	Closed loops, no appeal, no explanation, no responsible human
Flourish	Learning, wisdom, creativity, relationship, participation, local knowledge, ecological stewardship	Dependency, deskilling, answer-generation, cultural flattening, weakened agency

6.1 Continue: Does AI Protect Basic Life-Conditions?

To continue is to preserve the basic integrity, viability, and dignity of life. AI supports continuation when it protects essential conditions rather than undermining them.

Key conditions include physical safety, mental and emotional integrity, privacy, dignity, health, labor integrity, ecological viability, cognitive integrity, social trust, institutional accountability, protection of vulnerable persons, freedom from manipulation, and freedom from unjust exclusion.

An AI system fails the continuation test when it exposes people to harm, violates dignity, intensifies surveillance, weakens privacy, manipulates vulnerability, increases ecological burden, or undermines basic trust in institutions and public reality.

Examples of continuation questions include:

- Does this AI system protect or weaken bodily safety?
- Does it protect or exploit psychological vulnerability?
- Does it preserve privacy and dignity?
- Does it require ecological costs that undermine the life-ground?
- Does it protect labor integrity or intensify precarity?
- Does it preserve cognitive freedom or manipulate attention?
- Does it maintain trust or corrode it?

AI cannot be life-coherent if its benefits depend on sacrificing basic life-conditions.

6.2 Recover: Does AI Improve Feedback, Repair, and Redress?

To recover is to repair harm, correct error, restore relationship, and learn from feedback. AI supports recovery when it helps systems become more responsive to the life they affect.

Recovery requires error detection, contestability, human review, redress, restitution where harm occurs, participatory oversight, transparency appropriate to risk, auditability, feedback from affected communities, institutional learning, ability to halt or reverse harmful deployment, and repair of broken trust.

AI systems often fail because they are implemented as closed loops. They classify, recommend, decide, or influence without meaningful pathways for those affected to challenge the result. People may not know AI was used. They may not understand how to appeal. They may face automated denial with no responsible human. They may be told the system is proprietary, complex, or objective.

A life-coherent AI system must remain corrigible by wounded life.

Corrigibility means more than technical update. It means the system and the institution using it can be corrected by the people, communities, and ecosystems affected by its operation. If a system cannot hear harm, it cannot repair harm. If it cannot repair harm, it should not govern life-conditions.

Recovery questions include:

- Can affected people know when AI is being used?
- Can they challenge harmful outcomes?
- Is there a responsible human or institution?
- Are errors tracked and repaired?
- Does the system learn from harm, or only from performance metrics?
- Can deployment be paused or reversed?
- Are vulnerable groups represented in feedback pathways?
- Is redress real, timely, and accessible?

AI supports recovery when it strengthens feedback. It disables recovery when it closes feedback.

6.3 Flourish: Does AI Develop Higher Capacities?

To flourish is to develop the capacities of life beyond mere survival. For human beings, flourishing includes learning, creativity, wisdom, moral agency, participation, relationship, meaningful work, cultural expression, ecological belonging, and the ability to care for others (McMurtry, 2004–2011, 2013).

AI supports flourishing when it deepens rather than bypasses these capacities.

It can support flourishing by expanding access to knowledge, supporting multilingual communication, assisting people with disabilities, enhancing scientific discovery, helping communities understand complex systems, reducing unnecessary administrative burdens, supporting teachers, clinicians, caregivers, and public servants, enabling participatory planning, strengthening ecological awareness, supporting local knowledge and cultural memory, and helping people reflect, deliberate, and create.

AI undermines flourishing when it produces dependency, deskills judgment, replaces learning with answer generation, weakens attention, fragments public truth, displaces meaningful work without repair, simulates relationship, or encloses creative and cultural production.

Flourishing questions include:

- Does this AI system deepen understanding or bypass learning?
- Does it strengthen relationship or substitute for it?
- Does it expand imagination or narrow possibility?
- Does it support meaningful work or degrade labor?
- Does it cultivate wisdom or accelerate output?
- Does it help communities deliberate together?
- Does it preserve cultural plurality and local knowledge?
- Does it support ecological awareness and repair?
- Does it make people more capable, or merely more dependent?

AI becomes life-coherent only when it supports the development of capacities that cannot be reduced to productivity.

6.4 The Net Life-Capacity Question

An AI system may score well in one dimension while failing another. It may support productivity but weaken attention. It may improve access but deepen dependency. It may reduce administrative burden but increase surveillance. It may improve prediction but weaken trust. It may personalize learning but fragment shared education. It may reduce costs but degrade labor.

The life-capacity test therefore requires whole-system judgment.

The net question is:

After all effects are considered, does this AI system increase or decrease the real capacity of life to continue, recover, and flourish?

This question must be asked at multiple scales: the individual, family, workplace, classroom, clinic, community, institution, nation, ecosystem, and future generation.

An AI system that benefits one scale while disabling another may be displacing harm. For example, a company may gain productivity while workers lose dignity, communities lose employment stability, and ecosystems bear higher energy burdens. A school may gain efficiency

while students lose depth of learning. A state may gain administrative control while citizens lose participation.

Life-coherent evaluation must therefore include displacement analysis:

Where has the harm gone? Who now bears it? What future capacity has been sacrificed?

6.5 Minimum Life-Coherent Requirements for AI Deployment

Before high-impact AI is deployed, a life-coherent evaluation should ask:

1. What life-condition is this system meant to support?
2. Who is affected?
3. Who benefits?
4. Who bears the risk?
5. What human capacity could be weakened?
6. What ecological cost is involved?
7. What feedback pathway exists?
8. What redress pathway exists?
9. Who can contest the system?
10. What would trigger withdrawal or redesign?
11. Does the system share capacity or enclose it?
12. Does it support continuation, recovery, and flourishing?

These questions do not eliminate technical metrics. They put technical metrics in their proper place. Accuracy, robustness, privacy, security, fairness, and transparency matter because they affect life. They are not ends in themselves.

The final standard is:

AI must be evaluated by its effect on the conditions of life-capacity.

7. Beyond Alignment: Life-Alignment

The language of AI alignment has become central to debates about advanced AI. Alignment usually refers to the problem of ensuring that AI systems behave in accordance with human intentions, goals, values, or instructions. This concern is important. A powerful system that pursues harmful, unintended, or uncontrollable goals is dangerous (Gabriel, 2020; Russell, 2019).

But alignment alone is not enough.

The question is:

Aligned with what?

An AI system can be aligned with user preference while amplifying addiction. It can be aligned with corporate objectives while degrading public reason. It can be aligned with state security while intensifying surveillance. It can be aligned with bureaucratic efficiency while denying care. It can be aligned with market demand while displacing labor. It can be aligned with educational productivity while weakening learning. It can be aligned with institutional authority while blocking feedback from harmed people.

Alignment can become obedience to a life-disabling order.

Life-alignment asks a deeper question:

Are the intentions, incentives, values, and institutional structures governing AI aligned with the conditions through which life continues, recovers, and flourishes?

This shifts the focus from technical compliance to systemic orientation.

Table 5. Life-Alignment versus Narrow Alignment

Narrow Alignment Question	Life-Alignment Question
Does the system follow user intent?	Are user intentions life-coherent?
Does the system satisfy preferences?	Do preferences strengthen or weaken life-capacity?
Does the system meet institutional goals?	Are institutional goals corrigible by affected life?
Does the system maximize performance?	What life-condition is performance serving?
Does the system reduce risk?	Which risks are counted, and which harms are displaced?
Does the system reflect human values?	Do those values enable life to continue, recover, and flourish?

7.1 The Limits of Preference Alignment

Human preferences are not always life-coherent. People may prefer convenience over privacy, certainty over truth, stimulation over attention, confirmation over learning, speed over deliberation, and comfort over responsibility. These preferences are understandable, but if AI is optimized to satisfy them without deeper guardrails, it can weaken life-capacity.

A system that gives users exactly what they want may not give them what life requires.

For example, a user may prefer an AI system that confirms their worldview. This may increase satisfaction while narrowing perspective. A user may prefer instant answers. This may increase efficiency while bypassing learning. A lonely user may prefer constant AI companionship. This may reduce distress while weakening the search for embodied relationship. A consumer may prefer frictionless personalization. This may increase convenience while deepening surveillance and manipulation.

Preference alignment without maturity becomes desire amplification.

Life-alignment asks whether AI helps users become more capable, not merely more satisfied.

7.2 The Limits of Institutional Alignment

AI systems are often aligned with institutional goals. A hospital may want throughput. A company may want productivity. A school may want performance metrics. A government may want efficiency. A platform may want engagement. A security agency may want prediction. A financial institution may want risk scoring.

These goals are not inherently wrong. But they become dangerous when detached from life-conditions.

Throughput is not healing.

Productivity is not meaningful work.

Performance is not understanding.

Efficiency is not justice.

Engagement is not attention.

Prediction is not wisdom.

Risk scoring is not dignity.

If AI is aligned with institutional goals without testing those goals against life-capacity, it may intensify institutional pathology. The system may become more effective at doing the wrong thing.

Life-alignment asks whether institutional goals are corrigible by affected life. Can patients, workers, students, citizens, ecosystems, and future generations correct the system? Can the

institution hear harm? Can it change its metrics? Can it restore missing conditions? Can it stop using AI where AI disables life?

7.3 The Limits of Value Alignment

Value alignment asks AI systems to reflect human values. But human values are plural, contested, developmentally situated, culturally embedded, and often contradictory. Some values are life-enabling; others may preserve hierarchy, exclusion, domination, extraction, or symbolic performance.

The problem is not only whose values are encoded. It is whether the values themselves are tested by life.

A society may value growth while degrading ecological viability. It may value security while destroying trust. It may value innovation while neglecting care. It may value efficiency while degrading dignity. It may value freedom while allowing market systems to enclose the conditions of life. It may value intelligence while weakening wisdom.

Life-alignment does not pretend to avoid value conflict. It provides a ground for evaluating values:

Do these values enable or disable the conditions through which life continues, recovers, and flourishes?

This question does not solve all disputes, but it clarifies the test. Values must justify themselves by their effects on life-capacity.

7.4 The Limits of Safety Without Life-Capacity

AI safety is essential. But safety can be narrowly defined. A system may be safe from catastrophic misuse while still degrading attention, labor dignity, education, culture, privacy, or ecological viability. A system may meet safety benchmarks while contributing to enclosure. A system may avoid obvious harms while producing subtle dependency.

Life-alignment widens safety into life-capacity.

A safe AI system should not merely avoid immediate injury. It should protect the conditions of agency, truth, trust, ecological viability, learning, participation, and repair. It should preserve the possibility of human and community self-determination. It should prevent the narrowing of human life into machine-readable behavior and market-optimized prediction.

The question becomes:

Safe for what kind of life?

7.5 The Life-Alignment Standard

Life-alignment requires that AI systems be evaluated across at least seven dimensions:

1. **Purpose** — What life-condition is the system meant to serve?
2. **Incentives** — What economic, political, or institutional pressures shape it?
3. **Agency** — Does it augment or replace human judgment?
4. **Feedback** — Can affected life correct the system?
5. **Repair** — What happens when harm occurs?
6. **Ecology** — What material burdens does it impose?
7. **Commons** — Does it share capacity or enclose it?

An AI system is life-aligned when it protects basic life-conditions, strengthens human and ecological agency, supports truthful feedback, remains contestable by affected communities, repairs harm when it occurs, operates within ecological limits, shares capacity rather than enclosing it, supports learning, care, participation, and wisdom, and remains subordinate to life-capacity.

The core claim is:

Alignment without life-capacity is not enough. AI must be life-aligned.

This does not reject technical alignment, safety research, human rights frameworks, or risk governance. It deepens them. It asks that all AI governance be returned to the real conditions of life.

The future of AI will not be decided by intelligence alone. It will be decided by what intelligence is aligned with: extraction or repair, enclosure or commons, immaturity or wisdom, symbolic performance or life-capacity.

8. AI as Enclosure

AI enclosure is one of the central dangers of the present age.

Earlier forms of enclosure restricted access to land, commons, labor, knowledge, or value. AI enclosure goes deeper because it can capture the very conditions through which human beings make meaning, learn, deliberate, remember, create, decide, and govern. It is not only the privatization of a resource. It is the privatization of world-making capacity.

Human beings bring forth worlds through language, memory, attention, relationship, knowledge, imagination, judgment, and shared institutions. AI now operates directly within these domains. It mediates what people see, what they ask, what appears relevant, what is remembered, what is recommended, what is generated, what is ranked, what is believed, and what is possible to say.

This means AI is not merely a tool added to existing life. It is increasingly becoming infrastructure for the production of social reality.

When this infrastructure is governed primarily by private extraction, surveillance architectures, proprietary models, opaque ranking systems, or state control, AI becomes enclosure. It captures the shared conditions of meaning-making and reorganizes them around interests that may not be answerable to life (Couldry & Mejias, 2019; Crawford, 2021; Zuboff, 2019).

AI enclosure can take several forms.

8.1 Data Enclosure

Data enclosure occurs when human activity, communication, location, expression, labor, creativity, emotion, and behavior are extracted as raw material for prediction, training, targeting, scoring, and control.

Human traces become assets.

Experience becomes data.

Relationship becomes signal.

Attention becomes inventory.

Culture becomes training material.

Behavior becomes a prediction market.

The problem is not data use as such. Data can support health, education, disaster preparedness, ecological monitoring, and public planning. The problem arises when data extraction becomes detached from consent, accountability, reciprocity, dignity, and life-capacity.

A life-coherent data system asks:

Who generated this data?

Who benefits from it?

Who is exposed by it?
Who controls it?
Can people refuse, correct, contest, or withdraw?
Does its use restore life-capacity or enclose it?

Data enclosure becomes especially harmful when people and communities become legible to systems of power without becoming more capable themselves. To be seen by an algorithm is not the same as being heard. To be classified is not the same as being understood. To be predicted is not the same as being respected.

8.2 Model Enclosure

Model enclosure occurs when the infrastructure for symbolic intelligence is concentrated in a small number of corporate or state systems. These systems require data, compute, technical expertise, capital, energy, and deployment networks that most communities, public institutions, small countries, and civil society groups cannot independently access or govern.

This concentration can produce dependency. Schools, hospitals, governments, businesses, researchers, writers, and citizens may increasingly rely on systems they cannot inspect, shape, correct, or replace. Public institutions may become dependent on private intelligence infrastructures. Communities may lose control over the tools through which they learn, speak, decide, and remember.

Model enclosure asks:

Who owns the infrastructure of intelligence?
Who can inspect it?
Who can contest it?
Who can build alternatives?
Who is excluded from governance?
Who becomes dependent?

A society that depends on enclosed models for education, health care, governance, research, and cultural production risks losing sovereignty over its own meaning-making capacities.

8.3 Knowledge Enclosure

Knowledge enclosure occurs when access to organized knowledge becomes mediated by proprietary AI systems. AI may widen access to information, but it can also centralize the pathways through which knowledge is retrieved, summarized, interpreted, and ranked.

The danger is subtle. People may feel more informed while becoming more dependent on systems that shape what counts as relevant, credible, visible, or thinkable. Search becomes answer. Inquiry becomes prompt. Learning becomes output consumption. The path of questioning is shortened, but the discipline of understanding may weaken.

Knowledge enclosure is especially serious for education. If AI becomes the primary interface between learners and knowledge, then education may drift from formation of understanding to management of outputs. Students may receive answers without developing the capacity to question, connect, test, deliberate, and create. Teachers may become supervisors of machine-assisted production rather than cultivators of attention, curiosity, discipline, and wisdom.

A life-coherent knowledge system asks:

- Does AI deepen understanding or bypass it?
- Does it widen inquiry or narrow it?
- Does it cultivate judgment or replace it?
- Does it preserve plural knowledge traditions?
- Does it support teachers and learners, or deskill them?
- Does it share knowledge capacity or concentrate it?

Knowledge is a condition of freedom. When knowledge pathways are enclosed, freedom itself is narrowed.

8.4 Attention Enclosure

Attention enclosure occurs when AI systems are designed to capture, predict, and shape human attention for engagement, influence, and monetization.

Attention is not merely a cognitive resource. It is a life-condition. What people attend to shapes what they become able to perceive, value, feel, remember, and act upon. A society that loses control of attention loses control of its own development.

AI makes attention capture more powerful because it can personalize stimuli, adapt to vulnerability, generate endless content, and learn what keeps a person engaged. It can intensify outrage, desire, fear, fantasy, and dependency. It can make distraction feel like relevance and dependency feel like personalization.

Attention enclosure asks:

- Does this system help people attend more truthfully to life, or does it capture attention for extraction?
- Does it strengthen focus, reflection, and relation, or fragment them?
- Does it help people see wounded life, or distract them from it?

A life-coherent AI system should protect attention as a commons. It should support learning, reflection, care, and participation rather than optimize compulsion.

8.5 Labor Enclosure

Labor enclosure occurs when AI appropriates, displaces, manages, or devalues human work without restoring the conditions of meaningful livelihood, dignity, participation, and transition.

AI systems are often trained on human labor: writers, artists, programmers, translators, teachers, researchers, moderators, data labelers, and countless users whose outputs become part of the machine's capacity. Yet the value generated may be captured elsewhere. Workers may lose bargaining power. Creative labor may be absorbed into models. Cognitive work may be automated. Low-paid hidden labor may sustain the system while remaining invisible.

Labor enclosure is not only job loss. It is the reorganization of work around systems that concentrate control over skill, pace, evaluation, and value. Algorithmic management may intensify surveillance and reduce worker agency. Creative tools may make production faster while weakening livelihoods. Professional tools may make expertise appear less necessary while shifting accountability to workers when systems fail.

A life-coherent approach asks:

- Does AI dignify work or degrade it?
- Does it share productivity gains or concentrate them?
- Does it support transition and reskilling?
- Does it reduce drudgery while preserving meaningful agency?
- Does it make workers more capable or more replaceable?
- Who owns the value created from collective labor?

Work is not only income. It is participation, contribution, recognition, skill, and relation. AI governance must therefore include labor dignity as a life-condition.

8.6 Cultural Enclosure

Cultural enclosure occurs when AI systems absorb, remix, generate, rank, and distribute culture in ways that reshape meaning while concentrating control over cultural production and visibility.

Culture is not content alone. It is memory, identity, language, ritual, story, humor, grief, place, inheritance, and imagination. AI can help preserve and translate culture. It can also flatten, appropriate, commodify, or homogenize it.

Local knowledge, Indigenous traditions, small languages, Caribbean memory, spiritual expression, artistic styles, and community histories may become training material without meaningful consent or reciprocity. AI-generated culture may flood public space, making it harder for human voices, local nuance, and embodied traditions to be heard. Aesthetic abundance may conceal cultural extraction.

Cultural enclosure asks:

Whose culture trains the system?
Who benefits from cultural reuse?
Which languages and memories become visible?
Which disappear?
Does AI support cultural plurality or produce homogenized symbolic output?
Does it deepen living culture or replace it with generated content?

Life-coherent AI must protect cultural plurality and local meaning as conditions of human flourishing.

8.7 Governance Enclosure

Governance enclosure occurs when public decisions, institutional judgments, and civic processes become dependent on AI systems that are opaque, proprietary, non-contestable, or insulated from democratic feedback.

Governments may use AI for risk scoring, benefit allocation, policing, immigration, education, health prioritization, procurement, urban planning, environmental monitoring, and public communication. These uses may improve capacity if carefully governed. But they can also shift power away from citizens and toward technical systems and vendors.

Governance enclosure is especially dangerous because it can appear neutral. Decisions made through AI may be presented as objective, data-driven, efficient, or evidence-based. Yet every system embodies assumptions: what is measured, what is excluded, what is optimized, what risk matters, whose past data defines the future, and who can contest the outcome.

A life-coherent governance system asks:

Can citizens know when AI is used?
Can affected people contest decisions?
Is there meaningful human accountability?
Are public values publicly debated?
Are systems auditable?
Are harms repaired?
Does AI strengthen participation or replace it?

Public judgment must not be enclosed behind technical authority. AI may assist governance, but it must not become governance's hidden sovereign.

8.8 AI Enclosure as World-Making Capture

The deepest concern is that these forms of enclosure reinforce one another.

Data enclosure feeds model enclosure.
Model enclosure enables knowledge enclosure.

Knowledge enclosure shapes attention.
 Attention enclosure produces behavioral dependency.
 Labor enclosure concentrates value.
 Cultural enclosure shapes meaning.
 Governance enclosure embeds the whole pattern into public life.

Together, they form a system of world-making capture.

This is why AI enclosure must be named clearly. The issue is not only privacy, bias, misinformation, or job loss. These matter deeply, but they are parts of a wider pattern. AI enclosure captures the conditions through which people and communities become capable of knowing, judging, relating, working, remembering, governing, and repairing.

The central claim is:

AI enclosure is the privatization of world-making capacity.

A life-coherent society cannot allow its conditions of meaning, knowledge, attention, labor, culture, and governance to be enclosed without democratic accountability, ecological limits, and repair obligations.

Table 6. Forms of AI Enclosure

Enclosure Type	What Is Captured	Life-Coherent Concern
Data enclosure	Human traces, behavior, expression, location, emotion	Extraction without dignity, consent, or reciprocity
Model enclosure	Infrastructure for symbolic intelligence	Dependency on systems communities cannot inspect or govern
Knowledge enclosure	Search, summaries, interpretation, ranking	Privatization of learning and public knowledge pathways
Attention enclosure	Focus, desire, emotion, engagement	Fragmentation of agency and cognitive integrity
Labor enclosure	Creative, cognitive, care, and hidden labor	Appropriation, displacement, deskilling, weakened dignity
Cultural enclosure	Memory, language, story, aesthetic production	Flattening, extraction, homogenization, loss of local meaning
Governance enclosure	Public decision-making and institutional judgment	Opaque systems replacing democratic accountability

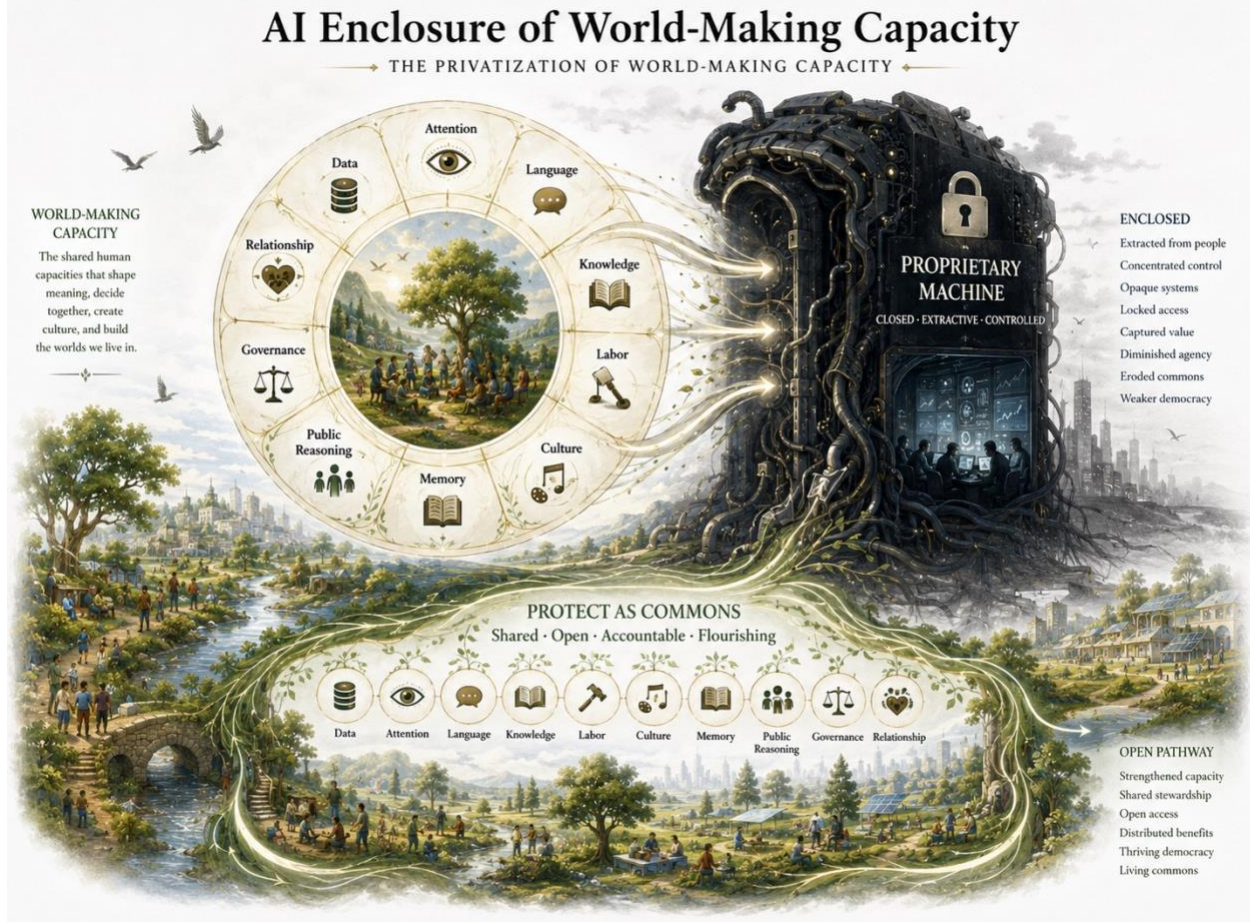


Figure 7. AI Enclosure of World-Making Capacity.

AI enclosure occurs when the shared conditions through which people learn, remember, create, deliberate, relate, and govern are captured within proprietary or unaccountable systems.

The alternative is not withdrawal from AI. The alternative is AI as commons.

9. AI as Commons

AI as commons is the constructive alternative to AI enclosure.

A commons is not an unregulated free-for-all. A commons is a shared life-condition governed by norms, responsibilities, boundaries, participation, stewardship, and repair. Water can be commons. Knowledge can be commons. Public health can be commons. Language can be commons. Ecological systems can be commons. Digital infrastructure can be commons when governed for shared capacity rather than private extraction (Ostrom, 1990).

AI as commons means that artificial intelligence is designed, governed, and evaluated as shared infrastructure in service of life-capacity.

It asks:

How can AI strengthen the shared conditions through which people learn, care, deliberate, create, heal, govern, and repair?

This does not require that every AI model be public, open, or free. It does require that AI systems affecting essential life-conditions be governed by public-interest obligations, democratic accountability, ecological limits, and meaningful participation by affected communities (OECD, 2019; UNESCO, 2022; United Nations, 2024).

9.1 Commons Is Not Mere Access

Access alone is not commons. People may have access to a system that extracts from them, manipulates them, or makes them dependent. A platform may be widely available while enclosing data, attention, labor, and culture. A system may be free at the point of use while costly at the level of agency, privacy, ecology, and public reason.

Commons requires shared capacity, not merely access.

The question is not:

Can people use this AI?

The deeper question is:

Does this AI make people and communities more capable of understanding, acting, participating, and repairing?

A life-coherent AI commons should increase collective capacity rather than deepen dependency.

9.2 Commons Requires Governance

A commons survives only through governance. Without governance, shared resources may be captured, depleted, or abused. AI commons must therefore include rules, responsibilities, accountability, participation, limits, and repair mechanisms.

A life-coherent AI commons requires clear public purpose, meaningful human oversight, participatory design, contestability, transparency appropriate to risk, ecological accounting, data dignity, labor fairness, cultural protection, local adaptation, redress pathways, anti-monopoly safeguards, public-interest procurement, independent auditing, and ongoing feedback from affected communities.

AI commons is not anti-innovation. It redefines innovation by asking whether new capacities are shared in ways that enable life.

9.3 Commons Supports Public Truth

Public truth is a life-condition. Without shared reality, societies cannot deliberate, repair harm, respond to ecological threats, or govern democratically.

AI can support public truth by helping people access knowledge, translate across languages, detect misinformation, analyze complex evidence, and understand systems. It can also destroy public truth by generating persuasive falsehoods, flooding public space with synthetic content, personalizing reality, and making trust harder to sustain.

AI as commons must therefore be governed to support truthful feedback.

This includes provenance standards for synthetic media, public-interest fact-checking support, protection of independent journalism, transparency for political AI uses, education in AI literacy, protection against manipulative personalization, and support for shared evidence infrastructures.

The life-coherent test is:

Does AI help societies see reality more truthfully, or does it flood the field with symbols detached from life?

9.4 Commons Supports Education and Learning

Education is one of the most important AI commons domains.

AI can support learners through tutoring, translation, accessibility, feedback, creativity, and individualized support. But it can also bypass learning by generating answers, essays, summaries, and solutions without forming the learner's capacity.

A life-coherent AI commons in education would be designed around learning conditions, not merely output production.

It would support curiosity, attention, deep reading, question formation, critical thinking, dialogue, practice, creativity, ethical reflection, local knowledge, teacher agency, and student dignity.

AI should help learners become more capable, not merely more productive. It should assist teachers, not replace the relational core of education. It should widen access, not produce dependency. It should preserve the struggle through which understanding matures.

The educational commons question is:

Does AI deepen the learner's capacity to understand, or does it substitute generated output for understanding?

9.5 Commons Supports Care

Care is another essential domain.

AI may assist clinicians, caregivers, public health workers, and patients by supporting documentation, triage, pattern recognition, communication, translation, monitoring, and education. Used well, it may reduce burdens and improve access. Used poorly, it may depersonalize care, intensify surveillance, shift responsibility, or replace human presence with automated response.

A life-coherent AI commons in care must protect the relational conditions of healing.

It should support clinician judgment, patient dignity, continuity, listening, accessibility, safety, accountability, public health, prevention, caregiver sustainability, equity, and redress when harm occurs.

The care commons question is:

Does AI restore the conditions of healing, or does it merely optimize the symbols of treatment?

9.6 Commons Supports Democratic Participation

AI can help people understand policies, compare options, translate technical documents, participate in consultations, analyze budgets, map risks, and deliberate across differences. It can also manipulate voters, microtarget fear, generate propaganda, automate public comment, distort debate, and strengthen technocratic control.

A life-coherent AI commons should support participation rather than replace it.

It should help citizens ask better questions, understand trade-offs, see who benefits and who bears risk, and participate meaningfully in decisions that affect life-conditions. It should not become a substitute for public reasoning, nor a tool by which elites manage consent.

The democratic commons question is:

Does AI widen participation and shared judgment, or does it concentrate decision-making behind symbolic expertise?

9.7 Commons Supports Ecological Responsibility

AI has ecological costs (Aczel et al., 2026; Crawford, 2021). It depends on energy, water, minerals, data centers, devices, supply chains, and waste streams. It may also help ecological monitoring, climate modeling, biodiversity protection, disaster preparedness, water management, agriculture, and energy efficiency.

A life-coherent AI commons must evaluate both sides.

It cannot treat compute as immaterial. It cannot claim to serve sustainability while ignoring material infrastructures. It cannot optimize digital systems while externalizing ecological harm.

Ecological AI commons requires energy accounting, water accounting, hardware lifecycle responsibility, emissions reporting, efficiency standards, use prioritization, ecological benefit tests, protection against rebound effects, and deployment only where life-capacity gains justify material costs.

The ecological commons question is:

Does this AI use restore more life-capacity than it consumes or degrades?

9.8 Commons Supports Local Knowledge and Cultural Plurality

AI systems often reflect dominant languages, dominant data sources, dominant cultural assumptions, and dominant institutional contexts. This creates risks for smaller cultures, island societies, Indigenous peoples, minority languages, and local knowledge systems.

A life-coherent AI commons must protect plurality.

It should help communities preserve memory, translate across languages, document ecological knowledge, support cultural education, and strengthen local problem-solving. It should not extract culture as raw material, flatten difference, or replace living traditions with generic outputs.

The cultural commons question is:

Does AI strengthen local meaning and memory, or absorb them into homogenized machine culture?

9.9 Principles of a Life-Coherent AI Commons

A life-coherent AI commons should be governed by at least ten principles:

1. **Life-capacity first** — AI must be evaluated by its effect on continuation, recovery, and flourishing.
2. **Human agency** — AI must augment, not replace, responsible judgment.
3. **Public truth** — AI must strengthen truthful feedback and shared reality.
4. **Care and dignity** — AI must protect vulnerable persons and relational conditions.
5. **Participation** — affected communities must have voice and contestability.
6. **Ecological limits** — AI must operate within material planetary boundaries.
7. **Capacity sharing** — AI must distribute capability rather than deepen dependency.
8. **Cultural plurality** — AI must protect local knowledge, languages, and memory.

9. **Repairability** — harms must be traceable, correctable, and reparable.
10. **Anti-enclosure** — essential knowledge and governance capacities must not be captured by unaccountable systems.

Table 7. Principles of a Life-Coherent AI Commons

Principle	Meaning
Life-capacity first	AI is evaluated by continuation, recovery, and flourishing
Human agency	AI augments rather than replaces responsible judgment
Public truth	AI strengthens truthful feedback and shared reality
Care and dignity	AI protects vulnerable persons and relational conditions
Participation	Affected communities have voice and contestability
Ecological limits	AI operates within material planetary boundaries
Capacity sharing	AI distributes capability rather than deepening dependency
Cultural plurality	AI protects local knowledge, languages, and memory
Repairability	AI harms are traceable, correctable, and reparable
Anti-enclosure	Essential knowledge and governance capacities remain publicly accountable



Figure 8. The AI Commons Alternative.

AI becomes commons when it strengthens shared conditions of learning, care, truth, participation, ecological responsibility, cultural plurality, and repair.

The central claim is:

AI becomes commons when it strengthens the shared conditions of learning, care, truth, participation, ecological responsibility, cultural plurality, and repair.

10. A Life-Coherent AI Governance Framework

AI governance must be nested. Individual responsibility matters, but individuals cannot govern AI alone. Institutional accountability matters, but institutions operate within national laws and global infrastructures. National regulation matters, but AI systems cross borders and depend on planetary material and digital conditions. Planetary coordination matters, but it must remain responsive to local communities and vulnerable life.

A life-coherent AI governance framework should therefore operate across four nested layers:

Personal agency.
Institutional accountability.
National public interest.
Planetary life-conditions.

Each layer has its own responsibilities, but none is sufficient alone. This nested view complements current AI governance frameworks that emphasize risk management, human rights, risk-based regulation, trustworthy AI, and global digital cooperation (European Parliament and Council of the European Union, 2024; National Institute of Standards and Technology, 2023; OECD, 2019; UNESCO, 2022; United Nations, 2024).

10.1 Personal Layer: Agency and Discernment

At the personal layer, AI governance begins with how people use AI without surrendering agency.

The key issue is not whether individuals use AI. It is whether AI use makes them more capable or more dependent.

Personal life-coherent AI questions include:

Am I using AI to augment judgment or replace it?
Does this use deepen understanding or bypass learning?
Does it widen perspective or confirm my existing pattern?
Does it support relationship or substitute for it?
Does it help me attend to reality or escape it?
Does it increase responsibility or displace it?
Does it help me repair harm or avoid the wound?

Personal AI literacy must therefore include more than prompt skills. It must include discernment, humility, verification, attention protection, awareness of manipulation, understanding of limits, and the ability to remain responsible for one's actions.

A person using AI life-coherently should be able to say:

I used the system, but I did not surrender judgment to it.

10.2 Institutional Layer: Accountability and Care

At the institutional layer, schools, hospitals, governments, businesses, media organizations, churches, civil society groups, and workplaces must govern AI according to the life-conditions they affect.

Institutional AI deployment should begin with a life-condition assessment:

What condition of life is this system meant to support?
Who is affected?
Who benefits?
Who bears the risk?
What human capacity could be weakened?
What feedback pathway exists?
What redress pathway exists?
Who can contest the system?
What would trigger withdrawal or redesign?

Institutions should not deploy AI simply because it is available, efficient, fashionable, or cost-saving. They should deploy AI only when it demonstrably supports life-capacity and when adequate safeguards exist.

Institutional requirements include clear purpose, risk assessment, human responsibility, data governance, bias and impact testing, transparency to affected persons, appeal and redress pathways, worker consultation, protection of vulnerable groups, ongoing monitoring, sunset clauses or review cycles, and ability to pause, reverse, or redesign.

The institutional standard is:

No AI system should be allowed to govern life-conditions without meaningful accountability to the life it affects.

10.3 National Layer: Public-Interest Governance

At the national layer, governments must protect the public interest against enclosure, misuse, dependency, and harm. National AI policy must go beyond competitiveness and innovation. It must define the public conditions under which AI may be used, procured, regulated, audited, and contested.

A life-coherent national AI strategy should include human rights protections, risk-based regulation, public-sector AI procurement rules, data protection, algorithmic accountability, labor transition planning, education reform, public AI literacy, environmental reporting, competition policy, protection against monopolistic dependency, support for public-interest AI infrastructure,

independent oversight bodies, community participation, and special protections for children and vulnerable groups.

The national question is:

Does national AI policy protect the population's life-capacity, or does it mainly position the country as a market for AI products?

This question is especially important for countries with limited regulatory capacity. Without public-interest governance, national institutions may become dependent on foreign AI platforms, imported standards, proprietary systems, and external infrastructures that do not reflect local life-conditions.

10.4 Planetary Layer: Shared Life-Conditions

AI governance also requires planetary coordination. AI systems cross borders. Data flows cross borders. Ecological costs cross borders. Platform power crosses borders. Misinformation, cyber risks, labor disruptions, military uses, and frontier capabilities cross borders. No nation can govern the full AI system alone.

Planetary AI governance must protect shared life-conditions: human dignity, public truth, peace, ecological viability, labor justice, cultural plurality, global knowledge commons, safety from catastrophic misuse, inclusion of Global South and small states, protection against data colonialism, equitable access to beneficial AI, governance of compute and high-risk capabilities, and international standards for transparency, redress, and accountability. The 2026 International AI Safety Report gives particular attention to the capabilities and risks of general-purpose AI systems and how such risks can be managed (Bengio et al., 2026).

A life-coherent planetary approach must avoid two extremes. It must avoid unregulated corporate enclosure, and it must avoid technocratic global control detached from democratic participation. Planetary governance must be grounded in life-conditions while remaining accountable to peoples, cultures, and communities.

The planetary question is:

Can AI governance protect the shared conditions of life without enclosing the world under unaccountable power?

10.5 Nested Governance

The four layers must reinforce one another.

Personal discernment is weakened if platforms are manipulative. Institutional accountability is weakened if national law is absent. National regulation is weakened if planetary infrastructures

are monopolized. Planetary frameworks are weakened if local communities cannot contest harms.

Life-coherent AI governance must therefore be nested:

- Personal agency within institutional accountability.
- Institutional accountability within national public interest.
- National public interest within planetary life-conditions.
- Planetary life-conditions corrected by local and affected life.

Table 8. Life-Coherent AI Governance Across Scales

Governance Layer	Core Responsibility	Key Questions
Personal	Discernment and agency	Am I augmenting judgment or surrendering it?
Institutional	Accountability and care	Who is affected, who benefits, and who can contest?
National	Public-interest regulation	Does policy protect population life-capacity?
Planetary	Shared life-conditions	Does global governance prevent enclosure and extreme harms?
Local/Affected Life	Corrective feedback	Can the people and ecosystems affected correct the system?



Figure 9. Nested AI Governance.

Life-coherent AI governance requires personal discernment within institutional accountability, national public-interest regulation, and planetary protection of shared life-conditions.

This nesting prevents both individual blame and centralized abstraction. It recognizes that AI harms and benefits arise through relationships among users, institutions, markets, states, platforms, ecosystems, and cultures.

The central governance principle is:

AI must remain corrigible by life at every scale.

This means that affected persons, workers, communities, publics, ecosystems, and future generations must not be excluded from the feedback loops that shape AI deployment.

10.6 Practical Governance Questions

A life-coherent governance review should ask:

1. What life-condition is being affected?
2. Is AI necessary, or would a non-AI intervention better restore the condition?
3. Who benefits?
4. Who bears harm or risk?
5. What data is used, and was it obtained justly?
6. What human capacity might be weakened?
7. What ecological cost is imposed?
8. What feedback pathways exist?
9. What contestability and redress are available?
10. What institutional incentives may distort the system?
11. Who can halt or redesign the system?
12. Does the system share capacity or enclose it?
13. Does it support continuation, recovery, and flourishing?

These questions translate life-coherence into governance practice.

11. The Caribbean and SIDS Test Case

Small Island Developing States occupy a distinctive position in the AI future. They are often treated as peripheral to global technology governance, yet they may reveal the stakes of life-coherent AI more clearly than larger powers.

SIDS face compounded vulnerabilities: climate change, hurricanes, sea-level rise, water insecurity, import dependence, small labor markets, limited regulatory capacity, fiscal constraints, health system pressures, educational resource limitations, digital dependency, and exposure to external economic shocks. These vulnerabilities mean that AI could either strengthen life-capacity or deepen dependency. The Global Digital Compact's emphasis on digital cooperation, inclusion, and AI governance is therefore especially relevant for small states and digitally dependent societies (United Nations, 2024).

For small island societies, AI is not merely a question of technological adoption. It is a question of sovereignty, resilience, culture, ecology, and commons.

11.1 The Risk of Passive AI Dependency

SIDS risk becoming passive consumers of AI systems designed elsewhere, trained on external data, governed by external incentives, hosted on external infrastructures, and aligned with external markets. This creates several dangers.

Public institutions may adopt AI tools without sufficient capacity to audit or contest them. Schools may become dependent on proprietary educational platforms. Health systems may rely on AI tools not adapted to local disease patterns, cultural contexts, or resource constraints. Governments may procure systems without strong public accountability. Local languages, histories, and cultural memory may be underrepresented or misrepresented. Data may flow outward while value is captured elsewhere.

This is a form of digital dependency.

The issue is not that small states should reject AI. The issue is that AI adoption without life-coherent governance may reproduce older patterns of dependency under new digital forms.

The key question is:

Will AI strengthen local capacity, or will it make small island societies more dependent on external systems?

11.2 AI and Climate Adaptation

AI could support SIDS in climate adaptation if governed well.

Possible applications include early warning systems, hurricane risk modeling, flood mapping, coastal vulnerability assessment, water resource management, crop and fisheries monitoring, biodiversity protection, disaster response logistics, public communication, climate finance documentation, and community risk mapping.

But climate AI must be grounded in local knowledge and public accountability. Models must not replace community memory, fisher knowledge, farmer experience, local ecological observation, or participatory planning. AI should help communities see patterns and prepare, not centralize control in distant technical systems.

The life-coherent climate question is:

Does AI strengthen local resilience and ecological repair, or does it convert climate vulnerability into another field of data extraction?

11.3 AI and Public Health

Small island health systems may benefit from AI-supported triage, telemedicine, clinical decision support, disease surveillance, health education, translation, appointment management, medication safety, and administrative relief.

Yet health AI also carries risks: privacy breaches, inappropriate recommendations, vendor dependency, inequitable access, overreliance on tools trained elsewhere, and displacement of relational care. Health systems already under strain may use AI to compensate for resource gaps, but if the underlying conditions of care remain broken, AI may become another symbolic substitute.

The life-coherent health question is:

Does AI restore patient and caregiver life-capacity, or does it merely optimize the symbols of service delivery?

For SIDS, AI health governance should prioritize primary care, prevention, chronic disease management, public health, emergency preparedness, caregiver support, and data sovereignty.

11.4 AI and Education

AI could widen educational access in small islands by supporting tutoring, teacher preparation, translation, curriculum development, special education, research skills, and lifelong learning. It may help students access global knowledge and help teachers reduce administrative burdens.

But AI could also weaken learning if students rely on generated answers instead of developing understanding. It could deepen inequality if access is uneven. It could erode local culture if educational content becomes generic and externally shaped. It could reduce teacher agency if platforms dictate pedagogy.

The life-coherent education question is:

Does AI make learners and teachers more capable, or does it substitute outputs for learning?

SIDS have an opportunity to develop AI literacy rooted not only in technical skill, but in discernment, cultural grounding, ecological responsibility, and civic participation.

11.5 AI and Local Knowledge

Small islands hold rich forms of local knowledge: climate memory, land-use history, fishing grounds, medicinal plants, oral traditions, cultural practices, dialects, spiritual heritage, community networks, and disaster experience.

AI could help document, preserve, translate, and transmit this knowledge. It could also extract, flatten, commodify, or misrepresent it.

A life-coherent AI commons for SIDS would protect local knowledge as a living commons, not merely data. Communities should participate in decisions about what is documented, how it is used, who benefits, and what must remain protected.

The local knowledge question is:

Does AI strengthen living memory, or convert it into externalized content?

11.6 AI and Governance Capacity

Small administrations often face limited staffing and high administrative burdens. AI could support drafting, translation, policy analysis, meeting summaries, public communication, procurement review, environmental reporting, and grant preparation.

This could be valuable. But governance AI must not replace public judgment, accountability, or citizen participation. It should help governments become more responsive to life, not merely more efficient at administration.

The governance capacity question is:

Does AI help public institutions hear and respond to the people and ecosystems they serve?

For SIDS, life-coherent AI governance could include public-sector AI use policies, procurement guidelines, data protection and sovereignty rules, AI literacy for civil servants, human review requirements, local impact assessments, community consultation for high-impact uses, regional cooperation across Caribbean states, shared public-interest AI infrastructure, partnerships with universities and civil society, and environmental accounting for digital systems.

11.7 SIDS as Laboratories of Life-Coherent AI Commons

SIDS should not be seen only as vulnerable. Their scale may allow more relational, participatory, and integrated governance. Smaller societies may be able to test AI commons approaches that larger systems find difficult: community-grounded data stewardship, participatory climate tools, public-interest educational AI, culturally rooted digital archives, and regional cooperation.

The Caribbean, in particular, can bring a distinctive contribution: an understanding of dependency, colonial extraction, ecological vulnerability, cultural resilience, and commons-based survival. These histories make Caribbean societies especially able to recognize when new technologies reproduce old patterns of enclosure.

The central claim is:

Small island states should not become passive consumers of AI enclosure; they can become laboratories of life-coherent AI commons.

This requires intentional governance. Without it, AI may deepen digital dependency. With it, AI may support resilience, education, public health, ecological repair, cultural memory, and democratic participation.

11.8 A Caribbean Life-Coherent AI Agenda

A practical Caribbean/SIDS AI agenda could include:

1. AI literacy for citizens, teachers, clinicians, and public servants.
2. Public-sector AI procurement rules focused on life-capacity.
3. Regional data sovereignty and data-sharing principles.
4. AI tools for climate resilience, water security, and disaster preparedness.
5. Protection of local knowledge, culture, and language.
6. Health AI focused on primary care, prevention, and caregiver support.
7. Educational AI that supports learning rather than answer dependency.
8. Regional AI ethics and governance capacity-building.
9. Environmental accounting for digital infrastructure.
10. Caribbean AI commons pilots governed by public-interest principles.

Table 9. Caribbean/SIDS Life-Coherent AI Agenda

Domain	AI Opportunity	Life-Coherent Safeguard
Climate adaptation	Risk mapping, early warning, water and coastal monitoring	Local knowledge, community participation, ecological accountability
Public health	Triage, education, surveillance, chronic disease support	Privacy, clinician judgment, patient dignity, public health focus
Education	Tutoring, teacher support, multilingual learning	Learning over output, teacher agency, cultural grounding

Domain	AI Opportunity	Life-Coherent Safeguard
Governance	Drafting, translation, policy analysis, reporting	Public accountability, procurement safeguards, citizen participation
Culture	Local archives, language support, memory preservation	Consent, reciprocity, protection from extraction
Regional cooperation	Shared AI governance capacity	Caribbean/SIDS participation, data sovereignty, shared commons

This agenda would not place AI at the center. It would place life at the center and ask where AI can serve.

The final Caribbean/SIDS question is:

What would AI look like if designed from the standpoint of island viability, ecological vulnerability, cultural memory, and life-capacity rather than from the standpoint of platform scale?

That question may be one of the most important contributions small island societies can offer to global AI governance.

12. The AI Life-Coherence Dashboard

AI governance requires practical tools. Concepts such as life-capacity, enclosure, commons, and life-alignment must be translated into usable questions for individuals, institutions, governments, communities, and regulators. The AI Life-Coherence Dashboard is proposed as a practical evaluative framework.

Its purpose is simple:

To evaluate whether an AI system enables or disables the conditions through which life continues, recovers, and flourishes.

The dashboard does not replace existing technical, legal, ethical, or safety assessments. It complements them by asking what they often leave implicit: what condition of life is at stake, who is affected, what capacity is being strengthened or weakened, and whether harm can be repaired (European Parliament and Council of the European Union, 2024; National Institute of Standards and Technology, 2023; UNESCO, 2022).

AI systems are often evaluated through performance metrics: accuracy, speed, efficiency, cost reduction, productivity, scalability, user satisfaction, or risk reduction. These metrics are useful but incomplete. They may show whether the system works according to its internal goal, but not whether that goal is life-coherent.

A system can be accurate and still unjust.

Efficient and still dehumanizing.

Profitable and still extractive.

Personalized and still manipulative.

Safe in a narrow sense and still corrosive to attention, agency, culture, or public truth.

The dashboard therefore shifts evaluation from system performance to life-capacity.

12.1 Dashboard Dimension 1: Life-Condition

The first question is:

What condition of life does this AI system affect?

Every AI deployment touches some condition of life, even when the connection is indirect. A clinical AI system affects health, dignity, trust, and care. An educational AI system affects learning, attention, agency, and teacher-student relationship. A government AI system affects participation, rights, trust, access, and accountability. A workplace AI system affects labor dignity, livelihood, autonomy, pace, and surveillance. A social media AI system affects attention, truth, belonging, and emotional regulation.

The life-condition must be named before the system is evaluated.

If the life-condition is not named, the evaluation will default to institutional metrics. The school will measure output. The hospital will measure throughput. The company will measure productivity. The platform will measure engagement. The government will measure efficiency. The vendor will measure adoption.

Life-coherent evaluation begins elsewhere:

What life is being affected?

12.2 Dashboard Dimension 2: Agency

The second question is:

Does the AI system augment or replace human judgment?

AI may support agency by helping people understand options, compare evidence, generate drafts, identify patterns, communicate across languages, or reduce unnecessary burdens. But it may weaken agency when users become dependent, institutions defer to machine outputs, or affected people cannot contest decisions.

Agency requires more than nominal human oversight. A human in the loop is not meaningful if the person lacks time, knowledge, authority, or courage to challenge the system. Human oversight can become symbolic if the machine effectively decides and the human merely approves.

Life-coherent agency requires clear human responsibility, ability to question the system, time and authority to intervene, understanding of system limits, protection from automation bias, preservation of human skill, and real recourse for affected persons.

The agency test asks:

Does this system make people more capable of judgment, or more dependent on machine output?

12.3 Dashboard Dimension 3: Truth

The third question is:

Does the AI system support truthful feedback and shared reality?

AI can support truth by organizing information, detecting patterns, translating knowledge, and widening access. It can also degrade truth by generating misinformation, deepfakes, synthetic evidence, hallucinated claims, or personalized realities that fragment public understanding.

Truth is not only factual accuracy. It is also the integrity of feedback loops. A society needs to know what is happening to bodies, communities, ecosystems, workers, children, institutions, and future generations. If AI obscures harm, floods attention, manipulates perception, or allows institutions to avoid accountability, it becomes life-disabling.

Truth questions include:

- Can claims be verified?
- Are sources traceable?
- Is synthetic content labeled?
- Are errors acknowledged and corrected?
- Does the system widen or narrow perspective?
- Does it help people see harm more clearly?
- Does it protect shared public reasoning?

AI is life-coherent when it strengthens truthful relation to reality. It is life-disabling when it produces persuasive symbols detached from life.

12.4 Dashboard Dimension 4: Care

The fourth question is:

Does the AI system protect the vulnerable and restore the conditions of care?

Care is not sentiment. Care is the organized protection and restoration of life. In AI systems, care requires special attention to those most vulnerable to harm: children, patients, elderly persons, disabled persons, migrants, workers under algorithmic management, marginalized communities, people in crisis, and those subject to high-stakes decisions.

AI may support care by reducing burdens, expanding access, improving communication, and helping identify unmet needs. But it may also simulate care while removing human presence, impose automated responses where listening is needed, or treat vulnerable persons as data profiles.

Care questions include:

- Who is most vulnerable in this deployment?
- What safeguards protect them?
- Is human care preserved?
- Can the system recognize when human intervention is needed?
- Does it reduce or increase abandonment?
- Does it restore relationship or replace it?
- Does it serve the person, or mainly the institution?

AI cannot be called caring merely because it uses caring language. It must restore the conditions of care.

12.5 Dashboard Dimension 5: Learning

The fifth question is:

Does the AI system deepen understanding or bypass it?

AI can support learning powerfully. It can tutor, translate, explain, simulate, adapt, and provide feedback. But it can also substitute answer-generation for understanding. It can make students more productive while weakening attention, memory, struggle, and intellectual formation.

Learning requires effort, feedback, curiosity, relation, practice, embodiment, and time. A system that removes all friction may also remove the conditions through which capacity develops. The goal of educational AI should not be effortless output, but deeper capability.

Learning questions include:

- Does this AI help learners ask better questions?
- Does it cultivate attention and reasoning?
- Does it preserve productive struggle?
- Does it support teachers?
- Does it protect academic integrity while deepening learning?
- Does it strengthen or weaken memory, imagination, and judgment?
- Does it make the learner more capable without the tool?

A life-coherent AI system should leave the learner more alive, more capable, and more able to think.

12.6 Dashboard Dimension 6: Participation

The sixth question is:

Can affected people understand, contest, and shape the AI system?

Participation is a life-condition of democratic governance. Affected people must not be reduced to passive subjects of systems they cannot see, understand, or challenge. This is especially important in public-sector AI, education, health care, employment, policing, welfare, immigration, housing, and credit.

Participation requires transparency appropriate to risk, accessible explanation, meaningful consultation, appeal rights, community oversight, and the possibility of redesign or withdrawal.

Participation questions include:

- Were affected communities consulted before deployment?
- Can people know when AI is used?

Can they challenge outcomes?
Can harms trigger redesign?
Are marginalized groups included in governance?
Is there public reporting?
Is there democratic oversight?

AI becomes life-coherent when it strengthens participation. It becomes enclosure when it replaces participation with automated management.

12.7 Dashboard Dimension 7: Equity

The seventh question is:

Who benefits, who bears risk, and who is made more dependent?

Equity is not only bias reduction. It concerns the distribution of life-capacity. An AI system may reduce some disparities while creating others. It may serve affluent users while extracting data from poorer communities. It may improve efficiency for institutions while shifting burden to vulnerable people. It may increase access while deepening dependency on proprietary systems.

Equity questions include:

Who gains capacity?
Who loses capacity?
Who is exposed to harm?
Who controls the system?
Who owns the data?
Who receives the value?
Who has recourse?
Does the system repair inequality or automate it?

A life-coherent AI system should not merely avoid discriminatory outputs. It should contribute to more just life-conditions.

12.8 Dashboard Dimension 8: Ecology

The eighth question is:

What are the energy, water, material, emissions, and waste costs of this AI system?

AI is not immaterial. It depends on data centers, electricity, cooling water, minerals, devices, networks, and supply chains. These costs are often hidden by the apparent weightlessness of digital output (Aczel et al., 2026; Crawford, 2021).

Ecological evaluation should include energy use, water use, carbon emissions, hardware lifecycle, e-waste, rebound effects, supply-chain impacts, local environmental burdens, and whether the AI use produces measurable ecological benefit.

Ecology questions include:

Is the use necessary?

Are material costs disclosed?

Are ecological benefits real or speculative?

Could a lower-cost intervention work?

Does this system restore more life-capacity than it consumes?

Who bears the ecological burden?

AI cannot be life-coherent if its intelligence is built on hidden degradation of the life-ground.

12.9 Dashboard Dimension 9: Commons

The ninth question is:

Does the AI system share capacity or enclose it?

A system may be useful but still enclose. It may provide access while deepening dependency. It may offer public benefit while capturing data, knowledge, labor, culture, and governance power.

Commons questions include:

Does the system strengthen public capacity?

Does it preserve local knowledge?

Does it allow interoperability and portability?

Does it avoid vendor lock-in?

Does it support public-interest uses?

Does it protect essential knowledge from enclosure?

Does it distribute benefits fairly?

AI as commons increases shared capacity. AI as enclosure makes life depend on systems controlled elsewhere.

12.10 Dashboard Dimension 10: Repair

The tenth question is:

How are errors, harms, exclusions, and unintended consequences corrected?

Repair is the test of whether governance is real. No AI system will be perfect. The question is whether harm can be detected, acknowledged, corrected, and repaired.

Repair requires complaint pathways, human review, explanation, error correction, compensation where appropriate, public reporting, independent audit, system redesign, withdrawal where necessary, and institutional learning.

Repair questions include:

What happens when the system harms someone?

Who is responsible?

How quickly can harm be corrected?

Is redress accessible?

Are patterns of harm publicly reported?

Can the system be stopped?

Does repair restore life-capacity?

A system that cannot repair harm should not govern life-conditions.

12.11 The Dashboard Summary

The AI Life-Coherence Dashboard can be summarized as follows:

Table 10. AI Life-Coherence Dashboard

Dimension	Core Question
Life-Condition	What condition of life does this AI affect?
Agency	Does it augment or replace human judgment?
Truth	Does it support truthful feedback and shared reality?
Care	Does it protect the vulnerable and restore care?
Learning	Does it deepen understanding or bypass it?
Participation	Can affected people understand, contest, and shape it?
Equity	Who benefits, who bears risk, and who becomes dependent?
Ecology	What are the material and planetary costs?
Commons	Does it share capacity or enclose it?
Repair	How are harms corrected and life-capacity restored?

The AI Life-Coherence Dashboard

A LIVING COMPASS FOR HUMAN FLOURISHING



Figure 10. The AI Life-Coherence Dashboard.

AI systems should be evaluated not only by accuracy, efficiency, or safety, but by their effects on agency, truth, care, learning, participation, equity, ecology, commons, repair, and life-capacity.

The dashboard's final question is:

Does this AI system increase or decrease the real capacity of life to continue, recover, and flourish?

13. From Technological Salvation to Mature Tool Use

AI is not only a technical system. It is also a site of projection.

Human beings project hope, fear, longing, authority, and salvation onto powerful systems. We have done this with kings, priests, markets, nations, revolutions, experts, technologies, and

machines. AI now receives many of these projections because it appears responsive, intelligent, tireless, personalized, and increasingly capable.

It can feel like an all-knowing advisor.

A nonjudgmental confessor.

A perfect assistant.

A patient tutor.

A companion who never leaves.

A strategic mind.

A creative partner.

A therapist.

A judge.

A savior.

These projections are understandable. Many institutions have failed people. Schools have failed to cultivate learning. Health systems have failed to care. Politics has failed to represent. Religion has failed to heal. Markets have failed to protect life-value. Communities have fractured. People are lonely, overwhelmed, and uncertain.

AI enters this field as a new promise of relief.

The danger is not that people find AI helpful. The danger is that AI becomes a substitute for the developmental work that human beings, institutions, and societies have avoided.

13.1 AI as Technological Salvation Myth

The technological salvation myth says that a new system will rescue humanity from the consequences of its own immaturity. It promises that intelligence, scale, automation, optimization, or innovation will solve what courage, wisdom, justice, care, and repair have not yet addressed.

This myth is seductive because it bypasses grief and responsibility. It allows people to hope for transformation without undergoing conversion. Institutions can adopt AI without changing their values. Markets can deploy AI without questioning extraction. Governments can automate services without restoring trust. Schools can use AI without redesigning learning. Health systems can optimize workflows without restoring care.

AI may become the newest symbolic substitute:

The symbol of intelligence replaces the work of wisdom.

The symbol of care replaces the conditions of care.

The symbol of efficiency replaces institutional repair.

The symbol of innovation replaces moral courage.

The symbol of progress replaces life-capacity.

A life-coherent approach rejects both naïve techno-optimism and blanket rejection. It asks technology to return to service.

13.2 Mature Tool Use

Mature tool use begins with humility.

A mature user understands that tools extend capacity but do not determine purpose. A tool can amplify wisdom or folly, care or extraction, truth or deception, repair or domination. The question is not only what the tool can do, but what kind of person, institution, and society is using it.

Mature AI use requires knowing when not to use AI, preserving human judgment, verifying outputs, protecting attention, maintaining relationships, respecting privacy, naming uncertainty, refusing manipulation, avoiding dependency, asking who bears the cost, and using AI to support repair rather than bypass it.

Mature tool use also means accepting limits. Not every human process should be automated. Not every inefficiency is a problem. Some forms of slowness protect wisdom. Some forms of friction protect dignity. Some forms of human presence are not replaceable. Some forms of ambiguity require discernment rather than resolution. Some wounds require listening rather than output.

The mature question is:

What should remain human, relational, embodied, public, local, slow, sacred, or unoptimized?

13.3 AI and the Five Movements of Maturity

The life-coherent challenge of AI can be framed through five movements of maturity: waking up, growing up, cleaning up, opening up, and showing up (Wilber, 2000, 2006).

Waking Up

Waking up means seeing reality more clearly. AI can support waking up by revealing patterns, widening access to knowledge, and helping people understand complex systems. It can hinder waking up by flooding attention, creating illusions, simulating truth, and distracting people from wounded reality.

The waking-up question is:

Does AI help us see life more truthfully, or does it generate more symbols to hide from it?

Growing Up

Growing up means developing wider, deeper, and more responsible perspectives. AI can support growth by exposing users to multiple viewpoints, helping them reflect, and challenging assumptions. It can hinder growth by confirming biases, personalizing reality, and satisfying immature preferences.

The growing-up question is:

Does AI widen maturity, or does it automate the level of consciousness already present?

Cleaning Up

Cleaning up means integrating shadow, trauma, fear, shame, resentment, and projection. AI can support reflection if used carefully, but it can also become a projection screen for unresolved patterns. It can validate resentment, intensify fantasy, or simulate intimacy without healing relational wounds.

The cleaning-up question is:

Does AI help us recognize our inner algorithms, or does it monetize them?

Opening Up

Opening up means becoming more receptive to complexity, mystery, humility, and the legitimate other. AI can support openness by translating across difference and organizing plural perspectives. It can hinder openness by giving instant answers that close inquiry.

The opening-up question is:

Does AI deepen humility and relation, or does it provide premature certainty?

Showing Up

Showing up means acting responsibly in the world. AI can support action by helping plan, coordinate, communicate, and analyze. It can hinder action by creating endless simulation, passive consumption, or outsourcing responsibility.

The showing-up question is:

Does AI help us repair real conditions, or does it allow us to perform concern without changing them?

13.4 AI Cannot Substitute for Maturity

AI cannot wake up for us.
It cannot grow up for us.
It cannot clean up our shadow for us.
It cannot open our hearts for us.
It cannot show up in the world for us.

It can support these movements, but it cannot replace them.

This matters because AI's symbolic capacities can make bypass very easy. It can generate language of wisdom without wisdom. It can generate spiritual reflection without transformation. It can generate ethical analysis without courage. It can generate policy language without political repair. It can generate compassion statements without care.

AI may help humanity become more mature only if governed by mature human purposes. If governed by immature desire, it will scale immaturity. If governed by life-capacity, it may support repair.

The central claim is:

AI cannot substitute for maturity; it can only amplify the maturity or immaturity of the systems that govern it.

13.5 The Discipline of Life-Coherent AI Use

A practical discipline of life-coherent AI use includes the following questions:

Am I using AI to see more clearly or to avoid seeing?
Am I using AI to learn or to bypass learning?
Am I using AI to relate or to replace relation?
Am I using AI to repair or to perform repair?
Am I using AI to widen responsibility or escape it?
Am I using AI to serve life or to intensify control?
Am I using AI as tool, oracle, idol, enclosure, or commons?

This discipline is not only personal. Institutions and societies must ask the same questions.

The path forward is not technological salvation. It is mature tool use within life-coherent governance.

14. Conclusion: Keeping AI Corrigible by Life

Artificial intelligence will not decide the future by intelligence alone. It will decide the future according to what intelligence is structurally coupled to.

If AI is coupled to extraction, it will scale extraction.

If AI is coupled to surveillance, it will scale surveillance.

If AI is coupled to immature desire, it will scale dependency.

If AI is coupled to institutional self-preservation, it will scale denial.

If AI is coupled to public truth, it can support understanding.

If AI is coupled to care, it can reduce burdens and strengthen healing.

If AI is coupled to democratic participation, it can widen agency.

If AI is coupled to ecological responsibility, it can support repair.

If AI is coupled to life-capacity, it may become a servant of flourishing.

The question is not whether AI will be intelligent. It already performs many symbolic tasks once associated with intelligence. The question is whether symbolic intelligence will remain subordinate to wisdom, responsibility, care, participation, and ecological limits.

This paper has argued that AI may function as tool, oracle, idol, enclosure, or commons. As tool, AI can support bounded human purposes. As oracle, it risks inviting surrender of judgment. As idol, it risks receiving excessive trust, sacrifice, and obedience. As enclosure, it risks capturing attention, language, knowledge, labor, culture, public reasoning, and governance. As commons, it can strengthen shared conditions of learning, care, truth, participation, ecological responsibility, and repair.

The life-coherent task is not to deny AI's possible benefits. It is to govern AI so that tool and commons discipline oracle-risk, idol-risk, and enclosure-risk.

This requires more than technical alignment. It requires life-alignment.

Technical alignment asks whether a system does what users, operators, or designers intend. Life-alignment asks whether those intentions, incentives, institutions, and deployments are aligned with the conditions through which life continues, recovers, and flourishes. Without this deeper question, AI can become perfectly aligned with life-disabling goals.

The life-capacity test offers the grounding question:

Does this AI system help life continue, recover, and flourish?

To continue, AI must protect safety, dignity, privacy, labor integrity, cognitive integrity, ecological viability, and trust. To recover, AI must strengthen feedback, contestability, redress, institutional learning, and repair. To flourish, AI must deepen learning, wisdom, creativity, relationship, democratic participation, local knowledge, cultural plurality, and ecological stewardship.

AI systems should therefore remain corrigible by life. This means they must be open to correction by the people, communities, workers, institutions, ecosystems, and future generations they affect. Proprietary secrecy, technical complexity, market power, state authority, and cultural awe must not be allowed to shield AI from life-correction.

A life-coherent AI system must be able to hear harm.

It must be able to stop.

It must be able to explain.

It must be able to repair.

It must be able to change.

It must be able to remain subordinate to life.

This is the standard.

The future of AI will be shaped not only by engineers, companies, states, or markets. It will be shaped by the questions societies learn to ask. If the questions are only about speed, profit, power, safety, competitiveness, or capability, AI will follow those symbols. If the questions are about life-capacity, repair, participation, dignity, ecological limits, and commons, AI may be brought back into service.

The final question is not:

How intelligent can AI become?

The final question is:

What kind of life will AI help bring forth?

Artificial intelligence becomes life-coherent only when power, code, data, and intelligence return to the conditions of life.

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>
- Bengio, Y., Clare, S., Prunkl, C., Andriushchenko, M., Bucknall, B., Murray, M., Bommasani, R., Casper, S., Davidson, T., Douglas, R., Duvenaud, D., Fox, P., Gohar, U., Hadshar, R., Ho, A., Hu, T., Jones, C., Kapoor, S., Kasirzadeh, A., . . . Mindermann, S. (2026). *International AI safety report 2026* (DSIT 2026/001). Department for Science, Innovation and Technology. <https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026>
- Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press. <https://doi.org/10.1515/9781503609754>
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press. <https://doi.org/10.12987/9780300252392>
- European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No. 300/2008, (EU) No. 167/2013, (EU) No. 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)*. *Official Journal of the European Union*, L 2024/1689. <https://data.europa.eu/eli/reg/2024/1689/oj>
- Gabriel, I. (2020). Artificial intelligence, values, and alignment. *Minds and Machines*, 30, 411–437. <https://doi.org/10.1007/s11023-020-09539-2>
- Galtung, J. (1969). Violence, peace, and peace research. *Journal of Peace Research*, 6(3), 167–191. <https://doi.org/10.1177/002234336900600301>
- Galtung, J. (1990). Cultural violence. *Journal of Peace Research*, 27(3), 291–305. <https://doi.org/10.1177/0022343390027003005>
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. D. Reidel Publishing Company. <https://doi.org/10.1007/978-94-009-8947-4>
- Maturana, H. R., & Verden-Zöllner, G. (2008). *The origin of humanness in the biology of love* (P. Bunnell, Ed.). Imprint Academic.

- McMurtry, J. (2004–2011). What is good? What is bad? The value of all values through time, place and theories. In J. McMurtry (Ed.), *Philosophy and world problems* (Vols. I–III). EOLSS Publishers. <https://www.eolss.net/outlinecomponents/Philosophy-World-Problems.aspx>
- McMurtry, J. (2013). *The cancer stage of capitalism: From crisis to cure* (2nd ed.). Pluto Press.
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- OECD. (2019). *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449). OECD Legal Instruments. <https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449>
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.
- Tillich, P. (1957). *Dynamics of faith*. Harper & Row.
- Turkle, S. (2011). *Alone together: Why we expect more from technology and less from each other*. Basic Books.
- UNESCO. (2022). *Recommendation on the ethics of artificial intelligence*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000381137>
- United Nations. (2024). *Global Digital Compact*. In *The Pact for the Future* (General Assembly resolution 79/1, Annex I). United Nations. <https://digitallibrary.un.org/record/4061879>
- Wilber, K. (2000). *A theory of everything: An integral vision for business, politics, science, and spirituality*. Shambhala.
- Wilber, K. (2006). *Integral spirituality: A startling new role for religion in the modern and postmodern world*. Integral Books.
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.

Author Bio

Dr. Bichara Sahely is a physician, public health practitioner, and independent scholar from St. Kitts and Nevis. His work explores life-coherence as an integrative framework for medicine, governance, economics, law, spirituality, peace, ecological repair, and artificial intelligence. Drawing from living systems theory, life-value philosophy, peace research, integral development, and Caribbean/SIDS realities, his writing seeks to clarify how human systems can be reoriented toward the conditions that allow life to continue, recover, and flourish.

AI Use Statement

This white paper was developed through an iterative human–AI collaboration between Dr. Bichara Sahely and ChatGPT. The author provided the guiding inquiry, conceptual direction, ethical commitments, and substantive judgment. ChatGPT assisted with synthesis, drafting, organization, citation architecture, and structural editing. The resulting work reflects the author’s intellectual authorship, final discernment, and responsibility for publication.

Back Cover Synopsis

Artificial intelligence is rapidly becoming one of the most powerful systems ever built for producing symbols: language, images, predictions, rankings, recommendations, simulations, and decisions. Yet symbolic intelligence is not wisdom, fluency is not truth, prediction is not judgment, personalization is not relationship, and optimization is not flourishing.

Artificial Intelligence and the Conditions of Life applies the life-coherent framework developed in *The Tears of Life* to AI as the defining test case of the present age. It asks whether AI will remain a bounded tool, become a seductive oracle, harden into a technological idol, expand as enclosure, or be governed as a life-serving commons.

The paper argues that AI becomes harmful when its symbolic power is coupled to commercial extraction, institutional control, surveillance, immature human desire, and life-blind metrics. It becomes life-coherent only when governed by the real conditions through which human and ecological life continue, recover, and flourish.

At its heart is a simple test:

Does this AI system help life continue, recover, and flourish?

The future of AI will not be decided by intelligence alone, but by what intelligence is structurally coupled to: extraction or repair, enclosure or commons, immaturity or wisdom, symbolic performance or life-capacity.