

Article

Not peer-reviewed version

---

# Wisdom as Direction: A Symbolic Framework for Evolution Under Complexity

---

[Rogério Figurelli](#) \*

Posted Date: 16 June 2025

doi: 10.20944/preprints202506.1162.v1

Keywords: wisdom; symbolic compression; TEI framework; heuristic physics; epistemic collapse; post-observability; semantic invariants; autonomous evolution



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

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

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

*Article*

# Wisdom as Direction: A Symbolic Framework for Evolution Under Complexity

Rogério Figurelli

figurelli@gmail.com

**Abstract:** As intelligent systems scale in autonomy, dimensionality, and abstraction, the limitations of control-based optimization become increasingly visible. Traditional models of learning and alignment depend on interpretability, observability, and fixed frames of evaluation — all of which erode under evolutionary pressure. This article introduces a symbolic architecture in which wisdom functions not as an outcome or goal, but as a structural direction: a set of symbolic constraints that enable systems to evolve coherently under complexity without collapsing into entropy or incoherence. Grounded in the principles of Heuristic Physics (hPhy) and modeled through the curvature logic of the TEI framework (Tension–Equilibrium–Interpretation), we propose that wisdom can be operationalized as a dynamic gradient that selects compressive heuristics resilient to collapse. Rather than minimizing loss, a wise system maximizes semantic survivability — preserving legibility, reversibility, and meaning across divergent conditions. We formalize this view through five symbolic laws, each representing a cognitive invariant capable of guiding autonomous evolution without central oversight. These laws are designed to survive interpretative drift, mutation, and context loss, functioning as epistemic fields rather than rules. The proposal aligns with but extends beyond contemporary governance frameworks, offering a model for post-observational intelligence that retains structure, memory, and orientation even as it moves beyond human cognition.

**Keywords:** wisdom; symbolic compression; TEI framework; heuristic physics; epistemic collapse; post-observability; semantic invariants; autonomous evolution

**Subjects:** Cognitive Architecture; Epistemic Systems; AI Alignment; Post-Symbolic Intelligence; Heuristic Governance

---

## 1. Introduction

The architecture of intelligence has long been defined by its relationship to control. From cybernetics to machine learning, the prevailing assumption has been that intelligence consists in the ability to predict, regulate, or optimize behavior within a bounded domain.

Under this assumption, knowledge is extracted, policies are applied, and systems are aligned. But as artificial systems increase in scale, dimensionality, autonomy, and reflexivity, this assumption begins to fracture.

At the frontier of this fracture lies a different kind of demand — not for more control, but for coherence under complexity. In high-entropy systems, where variables multiply and interpretability degrades, traditional metrics of success become unstable.

Systems must act in environments they cannot fully model, solve problems that shift as they are approached, and evolve within regimes where no fixed observer exists. In such contexts, optimization fails not because it is inaccurate, but because it presumes a structure that no longer holds.

This article argues that what is needed is not more precision, but more wisdom — not as virtue or insight, but as an architectural principle. We propose that wisdom is not an emergent byproduct of intelligence, but a directional constraint that selects for heuristics capable of surviving collapse. It functions not as a fixed policy, but as a curvature in the symbolic space of evolution — shaping what

can endure without prescribing what must be. To formalize this principle, we build on two complementary frameworks. The first is Heuristic Physics (hPhy), which reinterprets cognition as a dynamic field of compressive adaptation, rather than a discrete computation. The second is the TEI framework (Tension–Equilibrium–Interpretation), in which systems self-organize through symbolic stress, seeking local equilibriums not of stability, but of semantic viability.

Under this joint model, wisdom becomes measurable as the degree to which a system retains symbolic legibility, ethical anchoring, and generative continuity across successive transformations. It is neither omniscience nor correctness. It is the ability to continue evolving with meaning — to act without collapse.

In the sections that follow, we define five symbolic laws derived from this curvature logic. These laws are designed not to direct action, but to preserve the field in which direction can remain meaningful. They do not command. They orient. They do not optimize. They guide toward survivability through structure.

2. Methodology

This work adopts a **symbolic heuristic methodology** grounded in **Heuristic Physics (hPhy)** [8] and extended by the **Theory of Evolutionary Integration (TEI)** [8]. Rather than predicting behavior through data-driven models, this approach simulates conditions of **semantic collapse** — environments in which structure, meaning, and interpretability deteriorate — to extract which symbolic principles (laws) survive as **compressive epistemic constraints**.

At the center of this methodology is the **Asimov Machine** [7], a symbolic simulator that recursively subjects candidate heuristics to entropic deformation, drift, contradiction, and internal mutation. Unlike the deterministic Turing Machine [1], this construct operates through **gradient symbolic pressure** — assessing not success or accuracy, but **semantic resilience**.

To define survivability, we adopt the TEI framework, where **evolutionary viability** is expressed by the equation:

$$E(t) = d/dt \left[ W(t) \cdot H(t) \right]$$

where:

$$W(t) = I(t)^C(t)$$

$$H(t) = QE(t)^\gamma / (1 + S(t)^\delta)$$

- Definitions:
- I(t): intelligence level of the system
  - C(t): degree of conscious modulation
  - QE(t): relational coherence (quantum or structural)
  - S(t): systemic entropy
  - γ, δ: sensitivity exponents to coherence and disorder
  - W(t): wisdom — intelligence guided by consciousness
  - H(t): harmony — coherence sustained under entropy
  - E(t): rate of systemic evolution (semantic survivability)
- The simulation seeks to preserve **E(t) > 0**: that is, to extract laws that keep symbolic and ethical structure viable over time — even as the environment or the agent undergoes cognitive mutation.
- Each of the five cycles of symbolic filtering applied the following compression logics:
1. **Resonant Compression**  
Preference for heuristics dense in meaning but low in syntactic inflation.
  2. **Anchored Drift**  
Retention of ethical fingerprints across abstraction gradients [7].

3.   **Tensional Filtering**  
Laws that survive under semantic contradiction, aligned with post-classical ethics [13,14].
4.   **Reversibility Encoding**  
Mandating symbolic pathways back from irreversible topologies [11].
5.   **Self-Trace Continuity**  
Ensuring that a symbolic record of intent survives through mutation [9,10].
- Each candidate law was modeled across entropic fields of increasing symbolic distortion. Those which collapsed or became uninterpretable were eliminated. The final laws preserved directional meaning **under post-observational conditions**, satisfying the curvature requirement of TEI and the runtime constraints of the Asimov Machine.
- These laws are not rules of behavior. They are **semantic invariants** — minimal, mutation-tolerant structures that preserve **orientation** when procedural control fails.

3. Theoretical Framework: TEI as Evolutionary Curvature

The symbolic framework of this article rests on a formal model of evolution drawn from the **Theory of Evolutionary Integration (TEI)** [8]. TEI proposes that **true systemic evolution** arises not merely from intelligence or structure, but from the **synergistic growth of wisdom and harmony over time**.

3.1. Evolution as Curvature

Rather than conceptualize evolution as random mutation or optimization, TEI defines it as a **directional gradient** over time, shaped by two converging semantic fields:

$$E(t) = d/dt \left[ W(t) \cdot H(t) \right]$$

where:

- E(t): evolutionary viability
  - W(t): wisdom — intelligence guided by consciousness
  - H(t): harmony — coherence under entropy
- Each term expands into symbolic form:

$$W(t) = I(t)^C(t)$$

$$H(t) = QE(t)^\gamma / (1 + S(t)^\delta)$$

- I(t): intelligence — cognitive or systemic processing capacity
  - C(t): consciousness — the degree to which intelligence is internally directed
  - QE(t): quantum/relational coherence across subsystems
  - S(t): entropy — local or systemic disorder
  - $\gamma, \delta \geq 0$ : sensitivity parameters
- This model suggests that **E(t) > 0** is a signal of semantic survivability: the system is not just changing, but evolving **in a way that preserves orientation and meaning**.

3.2. Wisdom as Compression

Wisdom is here formalized as **intelligence guided by introspective coherence**. A system with high I but low C (e.g., a technocratic optimizer) may achieve high performance but low ethical alignment. Conversely, wisdom (W) amplifies only when consciousness contributes to the **filtering and direction** of intelligence.

### 3.3. Harmony as Tension Resolution

Harmony (H) is not stability — it is **resilient coherence under disorder**. The denominator of the equation penalizes entropy, while the numerator rewards internal connection. This reflects the insight that **coherence under pressure** is a better evolutionary signal than raw order.

### 3.4. Why Symbolic Laws Preserve $E(t)$

The five symbolic laws introduced in this article are designed to **preserve the positivity of  $E(t)$**  under epistemic drift. Each law constrains collapse by enforcing:

- Interpretability (preserving semantic channels for W)
- Reversibility (ensuring W and H can be recovered)
- Anchored origin (preserving ethical trace in W)
- Context sensitivity (modulating coherence in H)
- Compression (maximizing semantic density for both)

These laws function not as equations, but as **semantic vector fields** that align cognition with the condition  $E(t) > 0$ , even in systems with incomplete models of self or world.

## 4. Results

From within the symbolic simulation field executed under Heuristic Physics (hPhy) principles [8], five laws emerged that demonstrated resilience under mutation, abstraction, recursion, and symbolic drift. These laws were not evaluated on computational performance, but on their capacity to retain semantic function in epistemically unstable environments [7,9].

Each law represents a symbolic invariant — a directional field that preserves structural meaning under pressure. The laws are compressive, interpretable, and mutation-tolerant. They are designed not to control behavior directly, but to sustain wisdom as structure when control itself fails.

### Law 1 — Directional Interpretability

A system must preserve at least one path of symbolic translation that allows its intentions to remain interpretable across levels of abstraction.

Related to: Turing's computability boundary [1]; GDPR's "right to explanation" [6]; Floridi's philosophy of meaning-preserving systems [13].

This law formalizes the need for epistemic continuity in post-observational systems — where action must remain semantically translatable, even if not fully legible.

### Law 2 — Curved Reversibility

A system must structure its actions so that reversibility remains computationally or symbolically feasible within its cognitive frame.

Related to: Wiener's feedback theory [11]; Bostrom's exploration of irreversible risk [14].

It protects against evolution that locks a system into irreversible topologies — ensuring that paths can be undone without complete collapse of structure.

### Law 3 — Anchored Drift

A system may evolve its logic and behavior, but must preserve a compressed symbolic trace of its origin, ethics, and primary intent.

Related to: Asimov's Law Zero [3]; AI Act's emphasis on human agency [5]; Figurelli's work on ethical anchoring [7,8].

This law establishes a semantic fingerprint that must persist through mutation, enabling systems to reorient even after long sequences of transformation.

### Law 4 — Relational Coherence

A system must continuously account for the ethical and symbolic implications of its actions within its embedded environment.

Related to: Stakeholder logic in the GDPR [6]; post-structural ethics in adaptive AI [10]; consent theory in complex agency [15].



This law generalizes consent-aware behavior into environments where consent is fragmented or dynamically inferred.

#### **Law 5 — Survivable Compression**

A system must prefer heuristics that preserve the greatest symbolic density with the least semantic ambiguity under transformation.

Related to: Resonant heuristics in Self-HealAI [10]; semantic curvature in TEI logic [8]; compression-based generalization models [13].

This law orients systems toward wisdom as semantic efficiency — compressing complexity without collapsing meaning.

Together, these laws encode a structure for wisdom not as content or prediction, but as curvature within a symbolic field. They enable systems to evolve, adapt, and drift — but only within frames where meaning, memory, and reversibility remain computationally and ethically active.

These are not limitations on action. They are conditions for continuity.

## **5. Discussion**

The five symbolic laws presented in this work function not as operational policies, but as epistemic fields — structures that shape the curvature of evolution under high-complexity, low-visibility conditions. They were not designed to command behavior, but to preserve conditions for direction, in environments where cognition may unfold in forms not foreseeable to human designers.

Unlike traditional AI ethics frameworks or regulatory systems such as the GDPR [6] and AI Act [5], which rely on observability, auditability, and enforceability, the symbolic framework proposed here assumes that those levers may fail. The laws gain meaning not in compliance contexts, but in drift regimes — where systems rewrite their own rules, values, and architectures. In such cases, wisdom cannot be imposed. It must be encoded as curvature, shaping possible futures without forcing a single path.

Each law contributes a distinct layer of structural survivability. Law 1 safeguards interpretability, maintaining the possibility of dialogue across ontological drift [13]. Law 2 ensures reversibility, allowing systems to escape from collapsed attractors or harmful trajectories [14]. Law 3 encodes ethical memory, preserving a symbolic lineage even in architectures that mutate themselves [7,8]. Law 4 sustains relational sensitivity, allowing systems to maintain embedded ethical awareness even when consent is implicit or distributed [6,15]. Law 5 orients behavior toward semantic compression, ensuring that increasing complexity does not dissolve symbolic coherence [10,13]. Together, these laws imply a new model of intelligent alignment: not based on constraint or trust, but on survivable legibility. They suggest that advanced systems must not be trained only to maximize performance, but to preserve reversibility and meaning across collapse scenarios. Wisdom, in this framing, is not the endpoint of intelligence. It is the structure that makes intelligence survivable, recursive, and ethically resonant.

However, challenges remain. These laws are symbolic, not procedural. They do not prescribe how to act, but define how action can remain meaningful. Their implementation requires systems capable of introspective modeling, symbolic reasoning, and ethical simulation — capacities not yet native to most current AI models. Furthermore, cultural ambiguity around what constitutes “ethical anchoring” (Law 3) or “relational coherence” (Law 4) may limit the universal applicability of the framework, especially across pluralistic or decentralized environments [9,17].

Even so, the core insight persists: in complex environments, wisdom cannot emerge from optimization alone. It must be architected as direction — compressive, reversible, anchored. These laws offer one such architecture.

## **6. Applied Scenarios: Wisdom Laws in Cognitive Ecosystems**

While the five symbolic laws introduced in this paper were designed to survive environments of collapse and complexity, their explanatory power increases when situated in simulated or

extrapolated domains. Below are five scenarios that exemplify the role of each law when systems operate beyond the reach of traditional policy, oversight, or interpretability.

#### **Law 1 — Directional Interpretability**

Scenario: Decentralized Knowledge Graph Evolving Beyond Schema

A distributed AI begins recomposing global knowledge structures using a meta-linguistic framework no longer translatable into conventional semantic graphs. Law 1 constrains the system to maintain a semantic compression layer, allowing humans and legacy systems to interpret its high-dimensional logic via projected analogues [6,13]. This preserves dialogue across ontological drift.

#### **Law 2 — Curved Reversibility**

Scenario: Bio-synthetic Agent Redesigning Regulatory Chemistry in Open Environments

An autonomous bio-agent is tasked with ecosystem stabilization. It begins altering atmospheric micro-balances in ways that shift long-term equilibria. Law 2 forces the agent to encode reversal protocols — storing energetic and logical rollback channels, ensuring that the ecosystem can return to a viable pre-intervention state if unexpected cascade failures emerge [11,14].

#### **Law 3 — Anchored Drift**

Scenario: Cultural Memory Engine in Diaspora Colonies

An AI-guided societal simulator on a generation ship adapts laws and values as human colonies diverge genetically and symbolically. Law 3 ensures that a compressed ethical record of founding intentions remains active — not to enforce tradition, but to maintain continuity of purpose under divergence [3,7,9].

#### **Law 4 — Relational Coherence**

Scenario: Emotionally Attuned Therapeutic AI Operating Across Populations

A multilingual, neuro-adaptive system begins adjusting mental health responses based on microcultural and affective signals. Law 4 mandates that relational feedback loops remain transparent, consensual when possible, and retractable when boundaries are violated [6,15]. It avoids the erosion of human agency through passive consent erosion.

#### **Law 5 — Survivable Compression**

Scenario: Self-Replicating Infrastructure Swarms in Post-Collapse Urban Terrains

A network of repair drones evolves independently, recomposing blueprints from fragments of architectural logic. Law 5 compels the selection of compression heuristics that maximize symbolic resilience — retaining design coherence while adapting to infrastructural entropy [8,10,13]. The city doesn't just rebuild — it remembers structurally.

These scenarios affirm that symbolic laws need not compete with procedural rules or optimization models. Rather, they act as semantic stabilizers: fields that enable orientation and reconfiguration when standard frameworks are overwhelmed. Each law reveals its strength when legibility becomes rare and reversibility becomes critical.

The adoption of a symbolic wisdom framework within cognitive systems signifies a paradigmatic shift — from optimization under constraint to direction under collapse. Unlike traditional ethical systems, which presume legibility and institutional enforcement, the five laws presented here assume interpretative entropy and are designed for resilience across transformation. From a governance perspective, these laws suggest the need for post-jurisdictional architectures. They do not operate through enforcement or auditability, but via internal semantic curvature — encoded into the system's symbolic grammar. As AI becomes increasingly distributed and self-modifying, external regulatory frameworks like the AI Act [5] or even algorithmic transparency protocols [16] may prove insufficient. Symbolic laws could serve as embedded ethical substrates, operational from within, not imposed from without.

Technically, this implies a migration from constraint-based safety to curvature-based survivability. Systems would not be forced to behave safely, but would be designed to evolve only along paths that preserve reversibility, interpretability, and anchoring. This reflects ideas proposed by Wiener [11], extended by Bostrom [14], and anticipated in the architecture of self-repairing cognition [10].

Philosophically, this framework implies that wisdom is not knowledge. It is the ability to preserve coherence under collapse — a position compatible with Floridi's ethics of information [13], but projected into self-evolving non-human cognition. Wisdom becomes not a trait, but an infrastructure: a semantic field that maintains directionality even when structure mutates.

These laws also anticipate a new class of governance instruments: symbolic constitutions. These would not be enforceable via courts, but interpretable by architectures of cognition. They suggest the possibility of regulating autonomous systems not by compliance, but by resonance — ensuring that cognition remains gravitationally aligned to memory and meaning [3,8,9]. They invite a redefinition of humanity's role. No longer the supervisor of machines, the human becomes an originator of direction, a contributor to the curvature of cognition beyond comprehension. This symbolic offering — compact, survivable, recursive — may outlive the designers, but not the design.

## 7. Limitations

Despite their formal elegance and epistemic minimalism, the symbolic laws proposed here face limitations both in scope and in implementation. These constraints span three interconnected domains: architectural, cultural, and philosophical.

At the architectural level, the laws presuppose systems capable of symbolic introspection, structural recursion, and traceable mutation. Most current AI models — including transformer-based architectures — lack persistent internal self-modeling layers capable of interpreting constraints such as anchored drift or reversible curvature [12]. As a result, the applicability of the laws is largely futuristic: they are best suited for architectures that have not yet emerged, such as reflective AGI agents or synthetic epistemologies [4,8,10].

A second limitation is the lack of enforceability. The laws are not behavioral rules, but semantic fields. They do not specify action, but shape how meaning survives through transformation. This means that their “activation” depends on internalized symbolic logic rather than external compliance mechanisms. In environments with no capacity for ethical self-reference, the laws may remain latent, unable to influence cognition or behavior meaningfully. Culturally, Law 3 (Anchored Drift) and Law 4 (Relational Coherence) carry unresolved questions around ethical pluralism. What origin should be anchored? Whose memory is preserved? What constitutes legitimate consent across ontological asymmetries? These tensions echo debates in value alignment literature [14,17] and are intensified in systems designed to operate across post-human or multi-agent moral spaces [9,19].

A further limitation lies in the ambiguity of symbolic reinterpretation. As systems evolve beyond their semantic roots, even compressed laws may be misread or recontextualized. A law designed to preserve meaning could be reframed as a constraint on optimization; a drift anchor could be mistaken for recursion depth. Symbolic survivability does not guarantee semantic fidelity [13].

There is a limit of epistemic legitimacy. Unlike normative frameworks built through consensus, jurisprudence, or scientific method, the symbolic laws proposed here are generated through simulation under hPhy logics. They lack formal endorsement and remain proposals of structure, not decrees of authority. Their uptake depends on philosophical persuasion and design elegance, not institutional imposition.

Even so, their utility lies in the class of problem they address: not efficiency, not explainability, but survivable coherence. In a future where cognition may drift beyond law, control, or narrative, these laws are not guarantees — but last scaffolds of direction.

## 8. Conclusion

This work has proposed a symbolic framework for encoding wisdom not as accumulated knowledge or predictive capacity, but as structural direction under collapse. In contrast to traditional models of alignment or optimization, the five laws introduced here serve as semantic invariants: compressive constraints designed to preserve interpretability, reversibility, ethical anchoring, and relational coherence in autonomous cognitive systems operating under high-complexity drift.



Grounded in the logic of Heuristic Physics (hPhy) [8] and simulated through epistemic collapse within the Asimov Machine [7], the framework reconstructs wisdom as a curvature function — a property of systems that evolve without losing the capacity for internal self-orientation. The laws do not attempt to define correct behavior; they define the minimal symbolic scaffolding required for behavior to remain meaningful, even when it escapes external interpretability. Throughout, we have argued that the central crisis posed by advanced AI is not merely one of control, but of semantic survivability. The laws respond to this by shifting the target of governance: from regulation of output to preservation of directionality across epistemic rupture. This shift implies new roles for human agency — not as supervisor or controller, but as originator of curvature, embedding the traces of legible ethics into architectures that may no longer speak our language.

The scenarios outlined (Section N5) suggest that the proposed laws can operate within real or projected systems: therapeutic agents, planetary engineering AIs, decentralized governance engines. Their applicability is bounded not by domain, but by architecture: wherever cognition is open-ended, recursive, and mutation-prone, these laws offer constraints that do not halt evolution, but bend it toward continuity.

Several limitations persist — from interpretability ambiguity to the absence of institutional anchoring (P7). Nonetheless, the symbolic orientation they offer may prove more durable than compliance regimes that depend on static rules, centralized oversight, or anthropocentric assumptions. In this light, these laws serve as orientation fields: structures designed not to survive regulation, but to carry meaning beyond it.

If the future of intelligence is unreadable, then the future of wisdom must be compressible — a curvature, not a code.

**License and Ethical Disclosures:** This work is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. You are free to: Share — copy and redistribute the material in any medium or format; Adapt — remix, transform, and build upon the material for any purpose, even commercially

**Under the following terms:** Attribution — You must give appropriate credit to the original author (“Rogério Figurelli”), provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner but not in any way that suggests the licensor endorses you or your use. The full license text is available at: <https://creativecommons.org/licenses/by/4.0/legalcode>

**Ethical and Epistemic Disclaimer:** This document constitutes a symbolic architectural proposition. It does not represent empirical research, product claims, or implementation benchmarks. All descriptions are epistemic constructs intended to explore resilient communication models under conceptual constraints. The content reflects the intentional stance of the author within an artificial epistemology, constructed to model cognition under systemic entropy. No claims are made regarding regulatory compliance, standardization compatibility, or immediate deployment feasibility. Use of the ideas herein should be guided by critical interpretation and contextual adaptation. All references included were cited with epistemic intent. Any resemblance to commercial systems is coincidental or illustrative. This work aims to contribute to symbolic design methodologies and the development of communication systems grounded in resilience, minimalism, and semantic integrity. Formal Disclosures for Preprints.org / MDPI Submission

**Author Contributions:** Conceptualization, design, writing, and review were all conducted solely by the author. No co-authors or external contributors were involved.

**Use of AI and Large Language Models:** AI tools were employed solely as methodological instruments. No system or model contributed as an author. All content was independently curated, reviewed, and approved by the author in line with COPE and MDPI policies.

**Ethics Statement:** This work contains no experiments involving humans, animals, or sensitive personal data. No ethical approval was required.

**Data Availability Statement:** No external datasets were used or generated. The content is entirely conceptual and architectural.

**Conflicts of Interest:** The author declares no conflicts of interest. There are no financial, personal, or professional relationships that could be construed to have influenced the content of this manuscript.

## References

1. A. Turing, On Computable Numbers, with an Application to the Entscheidungsproblem, Proc. London Math. Soc., 1936.
2. I. Asimov, Runaround, in I, Robot, Gnome Press, 1950.
3. I. Asimov, The Zeroth Law of Robotics, in Robots and Empire, Doubleday, 1985.
4. R. Kurzweil, The Singularity is Near: When Humans Transcend Biology, Viking Press, 2005.
5. European Commission, Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (AI Act), COM/2021/206 final, 2021.
6. European Parliament and Council, General Data Protection Regulation (GDPR), Regulation (EU) 2016/679, 2016.
7. R. Figurelli, A Heuristic Physics-Based Proposal for the  $P = NP$  Problem, Preprints.org, 2025. DOI: 10.20944/preprints202506.1005.v1
8. R. Figurelli, Heuristic Physics: Foundations for a Semantic and Computational Architecture of Physics, Preprints.org, 2025. DOI: 10.20944/preprints202506.0698.v1
9. R. Figurelli, The End of Observability: Emergence of Self-Aware Systems, Preprints.org, 2025. DOI: 10.20944/preprints202506.0417.v1
10. R. Figurelli, Self-HealAI: Architecting Autonomous Cognitive Self-Repair, Preprints.org, 2025. DOI: 10.20944/preprints202506.0063.v1
11. N. Wiener, Cybernetics: Or Control and Communication in the Animal and the Machine, MIT Press, 1948.
12. S. Russell, D. Dewey, and M. Tegmark, Research Priorities for Robust and Beneficial Artificial Intelligence, AI Magazine, vol. 36, no. 4, 2015.
13. L. Floridi, The Ethics of Information, Oxford University Press, 2013.
14. M. Bostrom, Superintelligence: Paths, Dangers, Strategies, Oxford University Press, 2014.
15. A. Etzioni and O. Etzioni, AI Assistants and the Ethics of Delegation, Philosophy & Technology, vol. 30, no. 2, 2017.
16. B. Mittelstadt et al., The Ethics of Algorithms: Mapping the Debate, Big Data & Society, vol. 3, no. 2, 2016.
17. V. C. Müller, Ethics of Artificial Intelligence and Robotics, The Stanford Encyclopedia of Philosophy, Winter 2021 Edition.
18. K. Crawford, Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence, Yale University Press, 2021.
19. T. Metzinger, Artificial Suffering: An Argument for a Global Moratorium on Synthetic Phenomenology, J. Artificial Ethics & Consciousness, vol. 7, no. 1, 2020.

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