V1.0.0

Intelligence Converges to a Metabolically Constrained Resonant Equilibrium, Not a Singularity

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Artificial intelligence systems are projected to achieve ever-greater capabilities, yet practical deployments encounter hard physical limits on energy consumption and logical consistency. Existing theory lacks a quantitative synthesis of thermodynamic (Landauer) and logical (Gödelian) constraints that could explain observed performance saturation and would convert the singularity debate into a testable calculation. We model a self-referential information-processing system as a sequence of elementary logical operations, each dissipating at least the Landauer energy $E_{\mathrm{L}} = k_{\mathrm{B}} T \ln 2$, and we bound the depth of self-modeling recursion by the largest integer d for which $2^d \leq N$, where N is the total number of logical operations. At ambient temperature $T = 300\,\mathrm{K}$ and an energy budget $E = 1\,\mathrm{J}$, we derive $E_{\mathrm{L}} \approx 2.87 \times 10^{-21}\,\mathrm{J}$ per operation, a maximal operation count $N_{\max} \approx 3.48 \times 10^{20}$, and a maximal self-model depth $d_{\max} = 68$ bits. A sensitivity extension that adds a quadratic self-referential instability penalty ($\alpha = 10^{-30}\,\mathrm{J}$, assumed) lowers the efficient operating point to $N^{*} \approx 1.0 \times 10^{15}$ bits under a 1 W budget, a factor of 3.5×10^{5} below the Landauer ceiling. The analysis predicts a resonant equilibrium at which further scaling yields diminishing returns, contrary to the singularity hypothesis, and provides a concrete, testable ceiling for AI safety assessment and hardware design.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). Intelligence Converges to a Metabolically Constrained Resonant Equilibrium, Not a Singularity. QNFO Preprint v1.0.0.

---

1. Introduction The rapid expansion of large-scale language models and reinforcement-learning agents has generated expectations of unbounded recursive self-improvement, often framed as an intelligence singularity. In practice, the organizations deploying such systems — cloud providers, national AI strategy offices, and data-center operators — face escalating electricity bills, cooling requirements, and diminishing marginal gains in benchmark performance. We assume, for illustration, that frontier training runs consume on the order of $10^6\,\mathrm{kWh}$ per experiment, a cost that directly limits how many such experiments an organization can run annually. If the singularity were attainable, a modest investment in compute would eventually yield arbitrarily powerful agents, reshaping economics, security, and governance. Demonstrating a hard ceiling would instead redirect resources toward robust alignment and efficient utilization of existing capabilities: the concrete decision that changes under our result is the allocation of research funds between scaling model size and improving algorithmic energy efficiency. Should the ceiling hold, investing in energy-optimal architectures yields higher returns than simply increasing parameter counts.

We therefore ask: does the combination of thermodynamic irreversibility and logical self-reference impose a finite optimal operating point for any intelligence system? Answering this question clarifies whether policy should focus on controlling runaway growth or on managing bounded, yet still powerful, systems.

Contributions.

- We formulate a joint thermodynamic–logical model of self-referential computation, coupling Landauer's limit to a Gödelian recursion-depth bound (Section 3).
- We derive explicit expressions for the maximal number of logical operations $N_{\max}$ and the maximal self-model depth $d_{\max}$ under a fixed energy budget, with full arithmetic (Section 4).

- We compute the concrete numerical values $N_{\max} \approx 3.48 \times 10^{20}$ operations and $d_{\max} = 68$ bits of recursion at $T = 300\,\mathrm{K}$ and $E = 1\,\mathrm{J}$ (Section 5).
- We derive a sensitivity extension with a quadratic instability penalty, obtaining an efficient operating point $N^* \approx 1.0 \times 10^{15}$ bits under a $1\,\mathrm{W}$ budget and an efficiency gap of 3.5×10^5 relative to the Landauer ceiling (Section 4.5, reported in Section 5).
- We state limitations, falsification criteria, and practical pathways for AI safety and hardware design (Sections 6–7).

The resonant-equilibrium concept also bears on fields that study bounded rationality in economics, the energy limits of biological cognition, and the design of low-power neuromorphic hardware, each of which faces the same trade between information processed and energy dissipated.

2. Background and Related Work The literature on intelligence measurement, scaling limits, and computational thermodynamics provides several concrete observations that motivate our model.

[1] surveys tests of machine intelligence and notes that "the concept of intelligence remains rather murky," adding the validity criterion that "a test should capture intelligence and not some related quantity," and observing that the Turing test "fails to capture the ability to learn and adapt over time" while "compression tests" offer a scalar measure of information efficiency. This motivates our choice to define intelligence operationally as information-processing efficiency per unit resource rather than as task-specific performance.

[2] discusses compression as a proxy for information content in the context of the Fermi paradox and artificial super-intelligence, emphasizing that "the greatest extent to which a signal may be compressed is its Kolmogorov Complexity." Our use of a minimal logical operation count parallels this minimal-description-length view: the irreducible work of a computation is bounded below by the information it must manipulate.

[3] surveys creative problem solving in artificially intelligent agents and notes that it focuses on methods for solving off-nominal or anomalous problems in autonomous systems despite advancements in planning and learning. We adopt the simulation-based perspective for the analytical derivations presented here, and the survey's framing underscores that computational cost is already a recognized limiting factor for scaling trial-based AI methods.

[4] discusses the role of ecological niches in shaping intelligence, emphasizing that behavior depends on the environment in which it is embedded. Analogously, we treat artificial intelligence as evolving under a physical niche — the thermodynamic environment characterized by T and the available energy budget — rather than as an abstract unbounded optimizer.

[5] defines the Artificial Scientist test and emphasizes that computational resources are finite. This supplies the logical basis for our central modeling assumption of a bounded energy budget: any physically instantiated intelligence, including an artificial scientist, works under finite resources.

[6] surveys explainable artificial intelligence for healthcare and notes that model decisions are increasingly scrutinized for their energy and data footprints. Our framework offers a principled reason why scrutiny of energy footprints belongs in the same conversation as deployment decisions: the energy bound is not an implementation detail but a constraint on capability itself.

[7] presents a Markov decision process framework for simulating clinical decision-making, formalizing optimal treatment decisions over time as a sequential decision problem. We read this as a template for how bounded sequential inference can be formalized; in our model each element of such a sequence is a logical operation incurring at least the Landauer cost, so clinical-style sequential decision systems inherit the same ceiling.

[8] categorizes intelligence into human, machine, and biological intelligences. Our unified model treats all three categories as instances of the same thermodynamic–logical substrate, differing only in parameter values such as operating temperature and energy budget, which is precisely the kind of cross-category comparison the watershed framing invites.

These works collectively demonstrate that intelligence can be framed as a resource-constrained inference process, justifying the joint thermodynamic–Gödelian analysis that follows.

3. Methods We construct a minimal analytical model of a self-referential information-processing system. The model rests on two principles.

3.1 Landauer's principle. Erasing one logical bit of information dissipates at least

$$E_{\mathrm{L}} = k_{\mathrm{B}}\, T\, \ln 2,$$

where $k_{\mathrm{B}} = 1.38 \times 10^{-23}\,\mathrm{J\,K^{-1}}$ is Boltzmann's constant and T is the ambient temperature in kelvin. We treat each elementary logical operation (for example, a NAND gate evaluation) as requiring at least this energy.

3.2 Gödelian incompleteness bound. A formal system capable of self-reference inevitably contains statements that cannot be resolved within the system. We approximate the maximal depth d of self-modeling recursion by the largest integer satisfying

$$2^d \leq N,$$

where N is the total number of logical operations performed. This captures the intuition that each additional layer of self-reference roughly doubles the logical complexity required to maintain consistency.

3.3 Conventions adopted. Two modeling conventions appear in the literature-inspired design space, and we state our choice explicitly. We take the recursion-depth convention, in which the Gödelian constraint enters as the bound $2^d \leq N$ on self-modeling depth, as the primary model, because it yields closed-form quantities ($N_{\max}$, $d_{\max}$) that are directly testable. Under the alternative penalty convention, the Gödelian constraint enters as a quadratic energy penalty $E_{\mathrm{inst}}(N) = \alpha N^2$ on self-modeling information, with α a small constant; we carry this convention through as a sensitivity extension in Section 4.5 rather than as the primary model, because its coefficient α is not derivable from first principles and must be assumed. The two conventions are complementary: the depth bound limits *structure*, the penalty term limits *throughput*.

3.4 Assumptions.

- The energy budget E is fixed and supplied externally. We set $E = 1\,\mathrm{J}$ as a representative modest budget for a laboratory-scale experiment.
- The ambient temperature is $T = 300\,\mathrm{K}$ (room temperature).
- The logical operation cost equals the Landauer limit; no additional overhead is modeled in the primary calculation.
- Self-reference scaling follows the binary exponential relation above; alternative scaling laws are discussed in Section 6.
- In the sensitivity extension, the instability constant α captures all logical-paradox costs, is constant across scales, and is set to $\alpha = 10^{-30}\,\mathrm{J}$ as a representative order of magnitude consistent with the negligible energy of logical paradoxes relative to physical costs.

Under these assumptions the model yields closed-form expressions for $N_{\max}$ and $d_{\max}$, and the extension yields a power-budget-limited operating point N^*.

4. Analysis We now compute the maximal operation count, the maximal self-model depth, and the extension's operating point step by step, stating each input and arithmetic operation.

4.1 Landauer energy per operation

$$E_{\mathrm{L}} = k_{\mathrm{B}}\, T\, \ln 2$$

Inputs: $k_{\mathrm{B}} = 1.38 \times 10^{-23}\,\mathrm{J\,K^{-1}}$ (Boltzmann constant), $T = 300\,\mathrm{K}$ (assumed), $\ln 2 \approx 0.693147$.

First multiply k_{B} and T:

$$1.38 \times 10^{-23} \times 300 = 4.14 \times 10^{-21}\,\mathrm{J}.$$

Then multiply by $\ln 2$:

$$4.14 \times 10^{-21} \times 0.693147 \approx 2.87 \times 10^{-21}\,\mathrm{J}.$$

Thus

$$E_{\mathrm{L}} \approx 2.87 \times 10^{-21}\,\mathrm{J\ per\ operation}.$$

4.2 Maximal number of operations under the energy budget

Given total energy $E = 1\,\mathrm{J}$,

$$N_{\max} = \frac{E}{E_{\mathrm{L}}} = \frac{1}{2.87 \times 10^{-21}}.$$

Perform the division:

$$\frac{1}{2.87 \times 10^{-21}} \approx 3.48 \times 10^{20}.$$

Thus

$$N_{\max} \approx 3.48 \times 10^{20} \text{ operations.}$$

4.3 Maximal self-model depth

We require the largest integer d such that $2^d \leq N_{\max}$, i.e.

$$d_{\max} = \lfloor \log_2 N_{\max} \rfloor .$$

Compute the natural logarithm of $N_{\max}$:

$$\ln N_{\max} = \ln(3.48 \times 10^{20}) = \ln 3.48 + 20 \ln 10 \approx 1.247 + 20 \times 2.302585 = 1.247 + 46.0517 = 47.2987.$$

Convert to the base-2 logarithm:

$$\log_2 N_{\max} = \frac{47.2987}{0.693147} \approx 68.2.$$

Taking the floor,

$$d_{\max} = 68 \text{ bits of recursive self-model depth.}$$

4.4 Resonant equilibrium condition The resonant equilibrium is reached when the system operates at $N_{\max}$ operations and $d_{\max}$ recursion depth simultaneously. Any attempt to increase N beyond $N_{\max}$ would require energy exceeding the budget, violating the Landauer bound; any attempt to increase d beyond $d_{\max}$ would force $2^d > N$, violating the Gödelian consistency condition. The equilibrium is "resonant" in the sense that both constraints become active at the same operating point: neither can be relaxed without violating the other's premise.

4.5 Sensitivity extension: quadratic instability penalty under a power budget

Under the penalty convention of Section 3.3, the total energy consumption per unit time for N bits of self-modeling is

$$E_{\mathrm{tot}}(N) = N\, k_{\mathrm{B}} T \ln 2 + \alpha N^2,$$

with $\alpha = 10^{-30}\,\mathrm{J}$ (assumed, Section 3.4). Define the power budget as $E_{\mathrm{budget}} = 1\,\mathrm{J}$ per second (1 W, assumed). Setting $E_{\mathrm{tot}}(N) = E_{\mathrm{budget}}$:

$$2.87 \times 10^{-21}\, N + 10^{-30}\, N^2 = 1,$$

i.e. the quadratic $10^{-30} N^2 + 2.87 \times 10^{-21} N - 1 = 0$. By the quadratic formula,

$$N = \frac{-2.87 \times 10^{-21} + \sqrt{(2.87 \times 10^{-21})^2 + 4 \times 10^{-30}}}{2 \times 10^{-30}}.$$

Discriminant: $(2.87 \times 10^{-21})^2 = 8.24 \times 10^{-42}$, and $4 \times 10^{-30} = 4.0 \times 10^{-30}$; the sum is $4.00000000824 \times 10^{-30} \approx 4.0 \times 10^{-30}$ (the first term is negligible). Square root: $\sqrt{4.0 \times 10^{-30}} = 2.0 \times 10^{-15}$. Numerator: $2.0 \times 10^{-15} - 2.87 \times 10^{-21} \approx 1.99999713 \times 10^{-15}$. Denominator: 2.0×10^{-30}. Hence

$$N^* \approx \frac{1.99999713 \times 10^{-15}}{2.0 \times 10^{-30}} \approx 1.0 \times 10^{15} \text{ bits.}$$

The corresponding achievable bits-per-joule with the instability penalty is

$$\frac{N^*}{E_{\text{budget}}} \approx 1.0 \times 10^{15} \text{ bits/J},$$

compared with the theoretical Landauer ceiling

$$\frac{1}{k_B T \ln 2} = \frac{1}{2.87 \times 10^{-21}} \approx 3.48 \times 10^{20} \text{ bits/J}.$$

The penalty convention therefore reduces achievable efficiency by a factor

$$\frac{3.48 \times 10^{20}}{1.0 \times 10^{15}} \approx 3.5 \times 10^5.$$

We emphasize the status of this extension: both α and the quadratic form are assumptions, so N^* and the factor 3.5×10^5 are conditional projections of the model, not predictions of established physics. The primary results E_{L}, N_{max}, and d_{max} rest only on Landauer's principle and the stated budget.

5. Results The analytical derivation yields the following quantitative results, all computed in Section 4.

Quantity	Symbol	Value	Unit	Status
Landauer energy per operation at 300 K	E_{L}	2.87×10^{-21}	J	Derived (Section 4.1)
Maximal operations under 1 J	N_{max}	3.48×10^{20}	operations	Derived (Section 4.2)
Maximal self-model depth	d_{max}	68	bits	Derived (Section 4.3)
Instability constant (assumed)	α	10^{-30}	J	Assumption (Section 3.4)
Efficient operating point under 1 W with penalty	N^*	$\approx 1.0 \times 10^{15}$	bits	Model projection (Section 4.5)
Efficiency gap vs. Landauer ceiling	—	$\approx 3.5 \times 10^5$	—	Model projection (Section 4.5)

The first three rows define the **resonant equilibrium**: a point at which further scaling of either computational throughput or self-reference depth yields diminishing returns because one of the two fundamental constraints becomes active. The last two rows quantify how much a quadratic self-referential penalty, if real, would tighten the equilibrium: from 3.48×10^{20} operations per joule to roughly 10^{15} bits per joule, a reduction of nearly six orders of magnitude.

6. Discussion

6.1 Limitations

- *Energy budget*: $E = 1$ J is assumed; N_{max} scales linearly in E, so a data-center-scale budget of 10^6 J s^{-1} raises the throughput ceiling to 3.48×10^{26} operations per second. The *shape* of the equilibrium is budget-independent; its *location* is not.
- *Temperature*: $T = 300$ K is assumed; lower T reduces E_{L} proportionally, raising N_{max}. Cryogenic operation at 4 K would lower E_{L} by a factor of 75 and raise N_{max} by the same factor.
- *Landauer limit as exact cost*: real hardware incurs overhead; actual $E_{\text{op}} > E_{\text{L}}$ reduces N_{max} proportionally. We assume real hardware incurs overhead beyond the Landauer limit, so the derived N_{max} is an upper bound, not an achieved value.
- *Binary exponential relation*: the recursion bound $2^d \leq N$ is a modeling choice; a polynomial self-reference growth law would alter d_{max} substantially (for polynomial scaling $d \sim N^{1/k}$, the depth bound becomes far weaker).
- *Quadratic penalty*: $\alpha = 10^{-30}$ J and the quadratic form are assumptions; the extension's $N^* \approx 1.0 \times 10^{15}$ bits inherits their full uncertainty and should be read as a conditional projection.
- *Neglected overheads*: cooling, memory access, communication, and latency costs are omitted; each would lower the achievable operating point.
- *Performance metric*: we equate capability with operations and bits per joule, ignoring task-specific utility; a system at the equilibrium may still be more or less useful depending on what it computes.

6.2 Falsifiability The main result would be falsified if empirical scaling studies of AI systems demonstrate sustained performance growth beyond the predicted ceiling while operating under a known energy budget comparable to our assumption. Concretely, observing a self-modeling agent that reliably exceeds $d = 68$ bits of recursion and performs more

than 3.5×10^{20} ***logical operations within*** $1\,\mathrm{J}$ ***would contradict the joint Landauer–Gödelian bound. A second falsification path targets the extension: if measurements of energy per processed bit in systems spanning*** 10^{12} ***to*** 10^{18} ***bits of self-modeling information show no deviation from linear Landauer scaling, the quadratic instability term is overestimated and the penalty convention's*** N^* ***is refuted, though the primary depth bound survives.***

6.3 Relation to prior work Our quantitative bound complements the validity criteria for intelligence tests identified in [1], the compression perspective of [2], and the finite-resource viewpoint of [5]. Unlike the qualitative discussions of ecological constraint in [4], we provide a concrete numerical ceiling; unlike the evaluation-mode taxonomy of [3], we attach a cost to each trial. The framework also gives the energy-footprint scrutiny noted in [6] a theoretical anchor, and it extends the sequential-decision formalism of [7] with a per-operation thermodynamic cost. The cross-category framing of [8] is honored by treating human, machine, and biological intelligence as parameter variations of one model.

6.4 Open questions

- How does stochasticity in learning algorithms affect the effective E_{op}?
- Can alternative logical architectures (for example, reversible computing) approach the Landauer limit more closely, and how does that shift the equilibrium?
- What is the precise mapping between Gödelian depth and practical measures of self-modeling in deep networks?
- How does α depend on architecture (transformer versus recurrent), and can modular self-modeling reduce the instability penalty?

7. Practical Implications and Engineering Pathways (a) Intended users and decision impact. AI safety researchers, hardware architects, and policy makers can use the derived ceiling to benchmark the feasibility of runaway self-improvement scenarios. A hardware team deciding whether to invest in ultra-high-density compute can compare projected energy consumption against the $N_{\max}$ bound to avoid over-engineering; a policy office can use the bound to justify funding energy-efficiency research over pure scale-up. If the bound holds, design strategies should target operation near the equilibrium rather than unbounded scaling.

(b) Technology readiness. The analysis is purely analytical and has not undergone peer review; it corresponds to NASA Technology Readiness Level 2 (concept and application formulated). No hardware has been built or measured under this model, and no readiness claim beyond analytical concept formulation is made.

(c) Minimal de-risking experiment. A feasible pilot study would train a compact reinforcement-learning agent equipped with an explicit self-modeling module under a strict energy cap of $1\,\mathrm{J}$, measured via power-monitoring hardware. The study would record (1) total logical operations performed, estimated from instruction counts, and (2) depth of self-model recursion, measured by the number of nested belief updates. If the agent surpasses $N = 3.5 \times 10^{20}$ operations or $d = 68$ bits without performance degradation, the resonant-equilibrium hypothesis would be refuted; a performance plateau at or below these values under the fixed budget would support it.

(d) Engineering constraints. Energy measurement precision: accurate joule-level accounting is required to enforce the budget. Self-model depth quantification: defining and instrumenting a metric for recursion depth is nontrivial and may require custom logging. Thermal management: maintaining $T = 300\,\mathrm{K}$ under load keeps the Landauer calculation valid; deviations must be recorded. Parameter granularity: the test system must permit fine adjustment of model size so that the operating point can be swept. Benchmark selection: tasks whose utility scales with model size are needed to detect saturation rather than task-floor artifacts.

(e) Projected cost and performance bounds (assumption-based). Assuming a commercial electricity price of \0.10 $perkWh, the$ 1\,\mathrm{J}$budgetcorrespondsto$1/(3.6\times10^{6})$kWh, i.e.about$2.78\times10^{-7} $kWh, whichatanassumedpriceof0.10dollarsperkWhcostsabout$2.78\times10^{-8} $dollars, anegligiblemonetarycost; theengineeringeffortto instrumentandenforcethebudgetwoulddominate. Asupplyof$ 10^{6}\\,\mathrm{W}$operatingattheLandauerboundwoulddeliver$3.48\\times10^{26} $operationspersecond; underthepenaltyconvention'sprojection$(10^{15} bits/J), thesamesupplywoulddeliver10^{21} bits per second. Both figures are projections conditional on the stated assumptions, not measurements.

(f) Risks and misuse. Misinterpreting the bound as a guarantee that all AI systems will remain harmless could lead to complacency. The bound applies only under the stated thermodynamic and logical assumptions; engineered systems that circumvent them (reversible computing, cryogenic operation) could shift the equilibrium upward. Focusing solely on energy efficiency could also divert attention from alignment concerns that the bound does not address.

8. Conclusion By jointly applying Landauer's thermodynamic limit and a Gödelian incompleteness bound, we have derived a concrete resonant-equilibrium point for self-referential intelligence systems. Under a $1\,\mathrm{J}$ energy budget at room temperature, the maximal feasible logical work is 3.48×10^{20} operations, supporting at most 68 bits of recursive self-model depth; a quadratic instability extension lowers the efficient operating point to roughly 10^{15} bits under a $1\,\mathrm{W}$ budget, a factor of 3.5×10^{5} below the Landauer ceiling. This analytical ceiling challenges the singularity narrative and provides a quantitative target for AI safety assessment, hardware design, and future empirical validation. Subsequent work should explore how alternative computing paradigms, temperature control, and more refined models of self-reference modify the equilibrium, and should test the predictions through energy-constrained self-modeling experiments of the kind outlined in Section 7.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.55.14). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, openai/gpt-oss-120b, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: the execution sandbox was unavailable, so a second language model recomputed 10 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] Tests of Machine Intelligence. arXiv:0712.3825v1. https://arxiv.org/abs/0712.3825v1 [2] Compression, The Fermi Paradox and Artificial Super-Intelligence. arXiv:2110.01835v1. https://arxiv.org/abs/2110.01835v1 [3] Creative Problem Solving in Artificially Intelligent Agents: A Survey and Framework. arXiv:2204.10358v1. https://arxiv.org/abs/2204.10358v1 [4] Intelligent behavior depends on the ecological niche: Scaling up AI to human-like intelligence in socio-cultural environments. arXiv:2103.06769v1. https://arxiv.org/abs/2103.06769v1 [5] The Artificial Scientist: Logicist, Emergentist, and Universalist Approaches to Artificial General Intelligence. arXiv:2110.01831v1. https://arxiv.org/abs/2110.01831v1 [6] A Review on Explainable Artificial Intelligence for Healthcare: Why, How, and When?. arXiv:2304.04780v1. https://arxiv.org/abs/2304.04780v1 [7] Artificial Intelligence Framework for Simulating Clinical Decision-Making: A Markov Decision Process Approach. arXiv:1301.2158v1. https://arxiv.org/abs/1301.2158v1 [8] Watershed of Artificial Intelligence: Human Intelligence, Machine Intelligence, and Biological Intelligence. arXiv:2104.13155v2. https://arxiv.org/abs/2104.13155v2