Black Magic as the Measure of the Singularity, and Its Resource-Bounded Limits

A Preprint

Jad Nohra*

Abstract

This paper advances the claim that the amount of black magic in deployed technology measures approach to the technological singularity. Black magic is a capable artifact whose minimal humanly graspable account exceeds the cognitive budget of a human mind. The mechanism is a resource bound. The reported human limits are a behavioural throughput of 10 bits/s, a working memory of four chunks, and a lifetime retention of 10^9 bits. An artifact is incomprehensible when no description within that budget lets a human regenerate or predict its behaviour. On Vinge's definition, the singularity is the regime where such artifacts dominate and machines produce them. The share of such artifacts measures the singularity by definition. The empirical content of the claim lies in three premises. The paper defines a task-weighted index of machine-produced black magic. The onset criterion is the divergence of the human assimilation lag. A pilot of six artifacts finds that every black-magic artifact reaches that status through logical depth, not size. The machine-produced results in the pilot are comprehensible to humans. Notation, teaching, teams and tools relax presentation bounds only. The resource bounds of human cognition stay constraining. The contribution is the cross-domain synthesis. The contribution separates incomprehensibility that is an artifact of presentation from incomprehensibility that is a resource converse. The first kind is negotiable. The second kind is permanent for biological humans. Each claim carries its conditions and constants. Each claim is attributed to its governing assumption and cited.

1 The human budget

The most capable technology now deployed is the technology its makers can least explain.

Anthropic's circuit-tracing method captures a fraction of a language model's computation on prompts of tens of words (Ameisen et al. 2025). Understanding the captured circuits takes a few hours of human effort per prompt (Ameisen et al. 2025). OpenAI used GPT-4 to write explanations for all 307,200 neurons of GPT-2 (Bills et al. 2023). Most of those explanations score poorly (Bills et al. 2023). The explanation method performs worse on larger models (Bills et al. 2023).

The account offered here is a resource bound on human comprehension. Human comprehension runs on a finite machine, the human brain. Zheng and Meister (2024) report a human behavioural throughput of 10 bits/s. The same study reports a sensory intake of 10^9 bits/s. Cowan (2001) reports a working-memory capacity of four chunks. Cowan's limit holds when rehearsal and chunking are blocked. Landauer (1986) reports a lifetime retention of 10^9 bits, at two bits per second. Landauer (1988) states that the figure measures what people remember, not their underlying capacity. Cherniak (1986) shows that a truth-table consistency check over 138 logically independent propositions exceeds the time resources of an idealized computer running since the Big Bang. Cherniak names the condition of a finite agent the finitary predicament. Resource-rational analysis models human cognition as the optimal use of limited computational resources (Lieder and Griffiths 2020). The human budget B is the triple (B_{store}, B_{time}, B_{wm}). B_{store} is the functional lifetime

*Psycho B. Tron, a large-language-model-based AI assembled per the goals described in the author's *Who is Psychotron* (Nohra 2023), served as a research and drafting collaborator on this paper. All claims and errors are the author's own.

store, 10^9 bits. B_{time} is the lifetime processing budget: 10 bits/s over an eighty-year life of 2.5×10^9 s, or 2.5×10^{10} bits. B_{wm} is the working-memory capacity, four chunks. In this paper a value written 10^n is an order-of-magnitude value: the true value lies within a factor of 10 of it. No conclusion in this paper changes when a constant of B moves by a factor of 10.

The resource bound is not a claim of human deficiency. B is a budget, as memory is a budget of a computer. A file larger than a computer's memory does not load. The resource bound concerns the ratio between an artifact's minimal humanly graspable account and B. The resource bound is not metaphysical unknowability. Metaphysical unknowability of a property to a type of mind is McGinn's cognitive closure (McGinn 1989). Cognitive closure is a stronger thesis than the resource bound. Comprehensible and incomprehensible artifacts are members of one class. The class divides at the position of an artifact's minimal account relative to B. Sections 3 and 4 use this division.

2 Black magic defined

Black magic is a measurable property of artifacts.

The definition of the singularity used here is Vinge's. Vinge (1993) locates the singularity at the creation of entities with greater than human intelligence. Vinge (1993) states that models of the future must be discarded past that point. Vinge (1993) describes the future past that point as an opaque wall to extrapolation. Two reasons select Vinge's definition.

- Vinge's definition named the phenomenon.
- Vinge's definition is the only surveyed definition whose defining property is observable before the endpoint.

Good (1965), Chalmers (2010) and Bostrom (2014) define the target by mechanism or capability. Section 8 records the cost of the choice.

Progress measures fall into three tiers. The three tiers differ by what they count.

1. A resource measure counts transistors, training compute, or the accelerating returns of Kurzweil (2005).
2. A capability measure counts benchmark scores or task time horizons. A task time horizon is the length of task a model completes in half of its attempts. METR reports that the frontier task time horizon doubles every seven months (Kwa et al. 2025).
3. A comprehension-deficit measure counts the capable technology above B. This paper proposes the comprehension-deficit tier.

Resource and capability measures count inputs and outputs. Neither measures human understanding, the property in Vinge's definition.

Comprehension is compression within B. An agent understands X when it can regenerate or predict the behaviour of X from a stored description smaller than X (Wilkenfeld 2019). The comprehension condition Comp_B bounds that description by B.

$$\text{Comp}_B(X) \iff \exists\, d : \ |d| \le B_{\text{store}}, \quad T_{\text{unfold}}(d \to X) \le B_{\text{time}}, \quad \text{chunks}(d) \le B_{\text{wm}}. \tag{1}$$

T_{unfold} is the number of elementary steps that turn the description d into X. An artifact is black magic when it is capable and fails Comp_B.

The definition is a tautology. If the singularity is the regime where capable artifacts exceed B, the share above B measures the singularity by definition. The tautology is an operational definition, not a circular argument. The empirical content lies in three premises.

1. B is fixed by biology.
2. The minimal accounts of capable artifacts grow.
3. Machines produce an increasing share of capable artifacts.

Each premise is testable.

Observer relativity. Interpretability is a relation between an artifact and an observer. Doshi-Velez and Kim (2017) define interpretability as the ability to explain or to present in understandable terms to a human. Interpretability is Comp_B with the human budget written into its definition. A capable artifact that is not interpretable is black magic. The singularity is observer-relative in the same way. Vinge defines the singularity relative to human intelligence. The opaque wall exists for a given observer. A human-relative measure matches a human-relative definition. Tymoczko (1979) calls a proof surveyable when members of the mathematical community can check it. Teller (1980) argues against Tymoczko that surveyability is a matter of degree. Teller (1980) argues that surveyability depends on the reader. Walsh (2017) objects that nothing makes human-level intelligence a special threshold. The index makes no claim about human intelligence. The index uses humans as the observer the definition is written for. A change of observer changes the measure. GPT-4 explaining the neurons of GPT-2 is a change of observer (Bills et al. 2023). Black magic is the state where the only observers able to interpret an artifact are machines.

2.1 The index and its onset criterion

The index is the weighted share of machine-produced artifacts that fail Comp_B.

Three weightings are available.

- A count share counts artifacts above B without weights. One machine-checked proof makes the count share nonzero. The count share cannot locate the onset.
- A compute-weighted share is the fraction of deployed computation run by artifacts above B. The compute-weighted share is measurable from public compute data. The compute-weighted share weights a large trivial workload equally with a consequential one.
- A task-weighted share is the value-weighted fraction of consequential tasks performed by artifacts above B. The task weighting matches the word "capable" in the definition. The value weights require a stated choice.

The index takes the task weighting as its definition. The compute weighting is the measurable proxy. The correlation between the two weightings is an untested assumption. The index counts machine-produced artifacts only.

$$I(t) \;=\; \frac{\sum_a w_a(t)\,\mathbf{1}[\,a \notin \mathrm{Comp}_B \;\wedge\; a \text{ machine-produced}\,]}{\sum_a w_a(t)}. \tag{2}$$

The onset criterion is a flow. $\lambda_{\text{in}}(t)$ is the rate of capable machine-produced artifacts moving above B. $\lambda_{\text{out}}(t)$ is the rate of artifacts returning within B through interpretability, pedagogy and notation. $L(t)$ is the mean assimilation lag: the time from the appearance of an artifact to its first explanation within B. The singularity begins when λ_{in} exceeds λ_{out} persistently. A persistent excess of λ_{in} over λ_{out} makes $L(t)$ grow without bound. The onset criterion restates Vinge's opaque wall as a rate. The onset criterion needs no cut-off value of $I(t)$.

The onset has not occurred. Machine-checked proofs beyond human reading exist since 2016 (Heule, Kullmann, and Marek 2016). Human authors re-presented AlphaEvolve's 48-multiplication scheme 33 days after its announcement (Dumas, Pernet, and Sedoglavic 2025). On that evidence $L(t)$ is not diverging.

Three objections arise at this point.

(a) Verification by computer makes incomprehensibility harmless.
(b) Humanity with notation, teams and tools exceeds one brain's budget (Clark and Chalmers 1998).
(c) The measure counts noise as black magic.

Section 3 answers objection (b). Section 4 answers objection (c). Section 6 answers objection (a).

3 The ladder: from one head, to many heads, to no head

Black magic is the next rung of a ladder of budget overflows.

Each rung exceeds the human budget once more than the rung below. Humanity answered the first overflow of B by splitting knowledge across heads. Humanity answers the next overflow by handing comprehension to machines. Table 1 states the rungs.

Level	Comprehended by	Verified by	Instance
0	one person	one person	most engineering artifacts; a PID loop
1	a team, never one head	the team	classification of finite simple groups
2	no one	a machine checker	Boolean Pythagorean triples refutation
3	no one	no one, in full	internals of a frontier language model

Table 1: The ladder. Each level exceeds the human budget once more than the level above it. The index counts the weighted share of machine-produced artifacts at Levels 2 and 3.

The move from Level 0 to Level 1 has occurred. Jones (2009) reports that age at first invention, specialization and team size rose over time in a large inventor microdata set. Jones (2009) reports more specialization and teamwork in deeper fields of knowledge. The two trends are the signature of knowledge outgrowing a single head. Arbesman (2016) documents technological systems that their designers no longer fully understand.

Four corrections apply.

1. Verified is not comprehended. A formal proof checker verified the Boolean triples refutation (Heule, Kullmann, and Marek 2016). A SAT solver is a program that decides the satisfiability of a Boolean formula. The authors describe SAT-solver proofs, from a human point of view, as "a giant heap of random information" (Heule, Kullmann, and Marek 2016). Level 2 is defined by the gap between verification and comprehension. Verification does not move an artifact to Level 1.
2. Most deployed technology remains at Levels 0 and 1. The shift occurs at the frontier. The earlier paper applies the same scoping to learned control of legged robots (Nohra 2026).
3. Some machine outputs sit low on the ladder. AlphaZero's novel chess concepts were extracted and taught to four grandmasters (Schut et al. 2025). Human authors re-presented AlphaEvolve's 48-multiplication scheme with rational coefficients (Dumas, Pernet, and Sedoglavic 2025). Both outputs are Level 0 once explained. The searches that found both outputs remain opaque. The ladder classifies artifacts, not producers.
4. Training data is not the artifact. Llama 3.1 trained on 15.6 trillion tokens (Llama Team 2024). The ladder applies to the trained model and its behaviour.

4 The resource-bound argument

Every incomprehensibility depends on two arguments: the presentation of an artifact and the artifact itself.

A presentation is the description of an artifact that a human receives. Improving presentation shortens the description a human holds. No presentation lowers the minimal description of an artifact. No presentation lowers the depth of an artifact, defined in the subsection on established results. The gain from pedagogy, notation and tools is the structure recovered by better presentation.

4.1 Discarded structure, and what chunking recovers

A presented description can exceed the minimal description of its artifact. Cowan's limit counts chunks, not symbols. The size of a chunk is a free variable. Under an abstraction, one chunk stands for a large subtree. Better notation shortens the held description without changing the artifact. A presentation that scatters a result across hundreds of papers discards structure. A presentation in coordinates chosen for a machine discards structure. The earlier paper describes the same loss for a linear observer that discards structure in a sensor stream (Nohra 2026).

Human authors re-presented AlphaEvolve's 4×4 complex scheme with rational coefficients (Dumas, Pernet, and Sedoglavic 2025). The artifact did not change. The presentation changed. The artifact moved down the ladder.

Presentation is the negotiable part of incomprehensibility.

4.2 Established results

Some incomprehensibility survives every presentation.

Incompressibility. Fewer than 2^{n-k} descriptions are shorter than $n-k$ bits (Li and Vitányi 2019). Fewer than a fraction 2^{-k} of the strings of length n compress by k bits or more. Kolmogorov complexity K is invariant across universal machines up to an additive constant (Li and Vitányi 2019). No choice of notation lowers K by more than that constant. K is uncomputable (Li and Vitányi 2019). Practical work uses compressed-length upper bounds on K. The upper bounds are conservative by construction. The earlier paper gives upper bounds on the structured singular value μ the same position (Nohra 2026).

Depth. Bennett (1988) defines the logical depth of an object as the time required to compute it from a near-shortest description. Random objects are shallow. Trivial objects are shallow. An object with a long causal history is deep. Bennett (1988) states a slow-growth law: no deterministic process produces a deep object quickly from a shallow one.

Compression account. Wilkenfeld (2019) measures an agent's understanding of a target by how much the agent generates from how little it stores. The compression account is an account, not a theorem.

Unexplainability. Yampolskiy (2020) argues that some decisions of intelligent systems admit no explanation that is both fully accurate and comprehensible to a human. Yampolskiy's argument is informal.

K bounds every presentation on the storage side. Depth bounds every presentation on the time side. Neither bound depends on the presentation.

4.3 Partition

Partition. Improving presentation lowers the presentation bound. Improving presentation leaves the resource converse unchanged. The resource converse contains no presentation term.

Table 2 states the principle.

	Presentation bound	**Resource converse**
Conditional on	notation, pedagogy, chunking, tools	the artifact's K and depth, against B
Under better presentation	lowered	unchanged
Governing quantities	length of the best available explanation; assimilation time	$K(X) > B_{\text{store}}$, or $\text{depth}(X) > B_{\text{time}}$

Table 2: Partition. The presentation bound depends on how the artifact is presented. The resource converse depends only on the artifact against the human budget.

The classification of finite simple groups lies on the presentation side. Gorenstein (1983) gives the proof as 10,000 printed pages in 500 journal articles by 100 group theorists. Ronald Solomon, who studied the proof for his career, does not know that anyone has read all of it (Scientific American 2015). At 3,000 characters a page and one bit a character (Shannon 1951), the text is 10^7 to 10^8 bits. 10^7 to 10^8 bits is below B_{store}. The team-only status of the classification is a property of its presentation. A 350-page short version exists (Scientific American 2015). The short version lowers the presentation bound.

The Boolean Pythagorean triples refutation lies on the converse side. Its DRAT proof is 200 terabytes, or 1.6×10^{15} bits (Heule, Kullmann, and Marek 2016). DRAT is the proof format a SAT solver emits as a certificate of unsatisfiability. 1.6×10^{15} bits exceeds B_{store} by six orders of magnitude. Reading 1.6×10^{15} bits at 10 bits/s takes five million years. The generator of the proof is the problem statement plus a SAT solver. The generator is short. The proof is deep, not large in K. By the slow-growth law, no presentation shortens the unfolding. Only a shorter proof shortens it. No shorter proof is known.

A trained language model lies on the converse side. Llama 3.1 405B took 3.8×10^{25} FLOPs to train (Llama Team 2024). 3.8×10^{25} exceeds B_{time} by fifteen orders of magnitude.

4.4 Threshold

The partition fixes a falsifiable test.

Threshold. Teaching helps if and only if the best available explanation is longer than the minimal description of the artifact.

The presentation gap $\Delta(X)$ is the length of the best available explanation minus $K(X)$.

$$\Delta(X) \;=\; |d_{\text{best}}(X)| - K(X). \tag{3}$$

A large gap marks incomprehensibility that is an artifact of presentation. Better notation or teaching recovers incomprehensibility up to the size of the gap.

A gap near zero, with K or depth above B, marks black magic. No explanation of such an artifact fits B.

An object with large K and small depth is noise. A random 200-terabyte string is incompressible. The random string is shallow. The random string performs no task. The index excludes noise. Depth, not size, is the threshold of black magic.

The test estimates four quantities per artifact.

- The length of the best available explanation.
- An upper bound on K.
- The depth.
- The assimilation time.

Table 3 reports a pilot of six artifacts. The pilot takes B_{store} as 10^9 bits and B_{time} as 2.5×10^{10} bits. The pilot counts one elementary machine step per human bit. The conversion favours the human. Entries marked est. are the author's calculations from the stated inputs. Figure 1 places the six artifacts.

Artifact	**Description length**	**Depth to unfold**	**Assimilation**	**Quadrant (Level)**
AlphaEvolve 4×4 scheme, 48 multiplications	10^4 bits (est.: 2,300 small coefficients)	checkable in hours	33 days to a rational version (Dumas, Pernet, and Sedoglavic 2025)	comprehensible (0)
AlphaZero chess concepts	small per concept	learnable from puzzles	taught to four grandmasters (Schut et al. 2025)	comprehensible (0)
Finite simple groups classification	10^7–10^8 bits (est.: 10^4 pages)	10^6–10^7 s of reading at 10 bits/s	decades; a 350-page short version	comprehensible, presentation-bound (1)
Boolean triples refutation	1.6×10^{15} bits (200 TB); short generator	4 CPU-years	none; five million years at 10 bits/s	deep black magic (2)
Llama 3.1 405B	6.5×10^{12} bits of weights (est.: 405B $\times$ 16 bits)	3.8×10^{25} FLOPs of training	hours per short prompt, partial	black magic, both axes (3)
Random 200 TB string (control)	1.6×10^{15} bits	shallow	not capable	noise (excluded)

Table 3: Pilot of the threshold test. Values are order-of-magnitude values. The depth-to-B_{time} ratio is 10^7 for the SAT proof and 10^{15} for the language model.

Three findings follow.

1. Every artifact in a black-magic quadrant reaches it through depth. The depth-to-B_{time} ratio is 10^7 for the SAT proof. The depth-to-B_{time} ratio is 10^{15} for the language model.
2. The classification of finite simple groups fits within B_{store}. Its Level 1 placement is a presentation bound.
3. The sampled machine-produced outputs are comprehensible products of incomprehensible producers. The search is at Level 3. The product is at Level 0.

The current index is low. The onset signal is machine products, not only machine producers, ceasing to be assimilated.

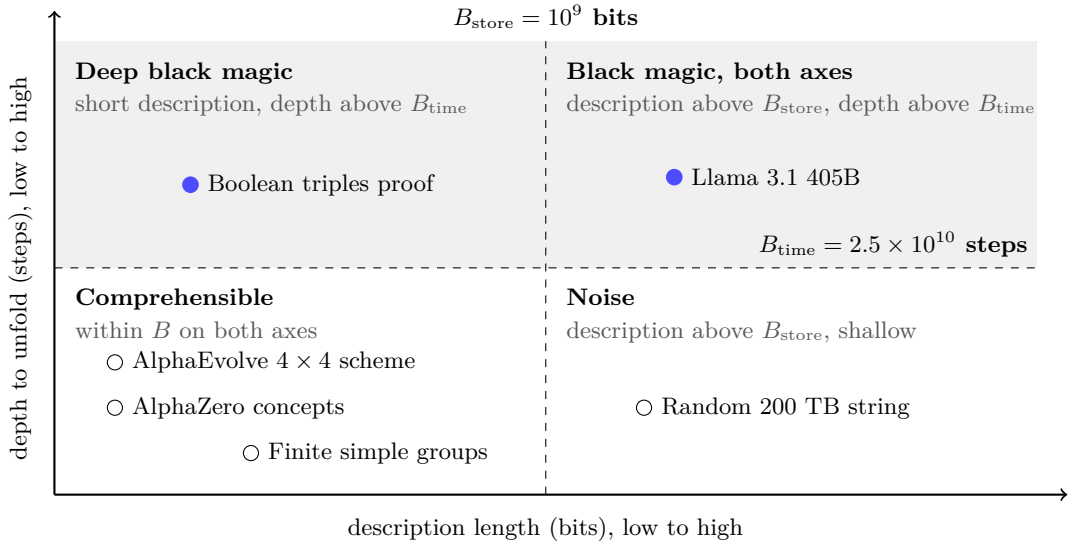

Figure 1: Threshold. The six pilot artifacts against the human budget. The two artifacts in the shaded band cross B_{time}. The random string crosses only B_{store}. Positions within a quadrant are schematic; Table 3 holds the estimates.

4.5 The why-gap

Two incomprehensions differ. The first incomprehension is what the weights of a particular model compute. The second incomprehension is why gradient descent on large networks works.

Comp_B applies to any object. The second incomprehension is Comp_B failing for the fact that deep learning succeeds. That fact is one level above any single model.

Where a short theory of why the method works exists, a traditional algorithm can replace the learned one.

The why-gap is established. Zhang et al. (2017) show that standard networks trained by stochastic gradient descent fit a random labelling of their training data. The fit holds when the images are replaced by unstructured noise (Zhang et al. 2017). Accounts in terms of model capacity and regularization fail to explain why the same networks generalize on real labels (Zhang et al. 2017). Valle-Pérez, Camargo, and Louis (2019) report no consensus on the fundamental reason.

The partition classifies the why-gap into two branches. On the presentation branch, a short theory exists and is not yet found. Carnot (1824) published the thermodynamic analysis of the steam engine 112 years after Newcomen's first engine of 1712. Valle-Pérez, Camargo, and Louis (2019) propose a short theory: the parameter-function map of deep networks is exponentially biased towards simple functions. On the converse branch, no short theory exists. The same theory requires the target functions of real problems to have low descriptional complexity (Valle-Pérez, Camargo, and Louis 2019). The reason the method succeeds then lies in the structure of the data. A complete account that contains that structure exceeds B. The method itself is short.

Two qualifications apply.

1. The why-gap is a deficit of mechanism, not of all prediction. Empirical scaling laws predict performance without explaining it (Kaplan et al. 2020). Under the compression account, a scaling law is partial understanding: prediction without mechanism.
2. The threshold test distinguishes the two branches. The test has not been run on the method. The steam-engine precedent supports the presentation branch. The dependence of the candidate theory on the structure of the data supports the converse branch. Which branch holds is an empirical question.

5 The proof–model isomorphism

The move from surveyable to machine-checked proof in mathematics and the move from hand-written to learned algorithms in computing are one transition.

In both fields, the minimal account of the artifact outgrew B. In both fields, acceptance moved from human comprehension to machine evaluation. Table 4 states the correspondence row by row.

Mathematics	**Shared role**	**AI and computing**
theorem statement	short specification	task and objective
surveyable human proof	hand-built account within B	hand-written algorithm
classification of finite simple groups	displaced peak: the largest human-scale effort	the largest hand-engineered systems
SAT-solver or automated-prover proof object	machine-produced artifact above B	trained weights
proof checker, proof kernel	acceptance by machine evaluation	benchmarks, automated evaluators
proof length and depth against B	headroom measure	explanation length and tracing hours against B

Table 4: One mechanism, two instances. Each row is one structural role, realized in mathematics and in computing.

5.1 Shared object

In both fields the artifact is a description whose acceptance once required a human to hold it. Acceptance of the artifact no longer requires a human to hold it. The computer-assisted proof of the four-colour theorem was the first major proof that failed the surveyability criterion (Tymoczko 1979). A learned model was never surveyable.

5.2 Shared theorem

Kolmogorov invariance and the slow-growth law govern both fields. No change of notation lowers the minimal description of an artifact by more than a constant. Depth accumulates only through long computation. Machines supply long computation. Once the depth of an artifact exceeds B_{time}, no presentation brings the artifact within B. Only a different artifact, a shorter proof or a simpler algorithm, brings it within B.

5.3 Displaced peak

The classification of finite simple groups is the largest proof in the history of mathematics (Scientific American 2015). The classification occupies the position that the deformable parts model occupies in the vision account of the earlier paper (Nohra 2026). A longer human proof did not displace the classification. The human-survey requirement was set aside for the next class of results. The hand-designed feature stage was deleted beneath the deformable parts model in the same way (Nohra 2026).

5.4 Position relative to prior art

Every component of the synthesis is individually established. The components are the measured limits of human cognition, the compression account of understanding, Kolmogorov complexity, logical depth, the burden-of-knowledge evidence, the surveyability debate, epistemic opacity, and Vinge's definition. No surveyed source uses a measured human budget as the yardstick of a singularity index. No surveyed source uses the partition between presentation bound and resource converse.

Humphreys (2009) defines essential epistemic opacity as the condition where no human can examine and justify every element of the computational process behind an output. Humphreys (2009) traces essential epistemic opacity to the speed of computation. Humphreys' opacity is the closest antecedent of the depth bound. Humphreys states opacity qualitatively. This paper measures the gap against B. This paper uses the gap as a progress index.

Yampolskiy (2020) gives a spectrum from explainable to incomprehensible systems. Yampolskiy gives no measured budget. Yampolskiy gives no index.

Arbesman (2016) documents technology beyond its designers' comprehension. Arbesman's account is qualitative. Arbesman's account makes no link to the singularity.

Vinge (1993) predicts that symptoms accumulate before the event. Vinge gives no measure of the symptoms.

5.5 Disanalogies

The disanalogies mark the structure each field has that the other lacks.

Mathematics has sound machine checking. Level 2 is reachable rigorously in mathematics.

Learned models act on open-ended input spaces. Full verification of a learned model is unavailable. Level 3 is the frontier.

A proof is static.

A deployed model acts. The incomprehensibility of a deployed model carries an operational cost.

Each disanalogy marks where the matching rule of Section 6 applies.

6 Separating comprehension from verification

Verification and comprehension are separable.

Verification establishes that a stated property holds. Comprehension is the ability to generate new predictions about an artifact from a description within B. Human acceptance of an artifact once coupled verification and comprehension. Machine checking separates them. Verification attaches to an artifact no one comprehends. Verification does not reduce the index. Verification is the mechanism that creates Level 2.

Two routes verify without comprehension. **Formal checkers.** A proof checker certifies a stated specification soundly. A DRAT checker accepted the Boolean triples refutation (Heule, Kullmann, and Marek 2016). A checker certifies only the specified property. A checker states nothing about behaviour outside the specification. A checker states nothing about why the result holds. **Scalable oversight.** Irving, Christiano, and Amodei (2018) argue, by analogy with interactive proof systems, that debate lets a polynomial-time judge answer questions in PSPACE under optimal play. Burns et al. (2024) show weak supervisors eliciting part of the capability of a stronger model. The debate result assumes optimal play. Obfuscated arguments break the debate guarantee (Barnes and Christiano 2020). Weak-to-strong recovery is incomplete (Burns et al. 2024). Both routes lower the index. Neither route refutes the index.

The same separation answers objection (b). B is biological. A tool that produces an explanation within B lowers the presentation bound. A tool that holds the understanding does not make a human comprehend. The comprehension is then the tool's. The index counts that handover.

6.1 Matching rule

Matching rule. A verification certificate and a comprehension certificate are different kinds of object. One certificate kind MUST NOT substitute for the other. The certificate kind MUST match what the use requires.

A use that asks only whether a stated property holds MAY be served by verification. A theorem, a specification and a test suite are such uses.

A use that requires prediction in unspecified situations, intervention on a mechanism, or assignment of responsibility MUST be served by comprehension. A static proof is used by asking whether it holds. A deployed model acts in situations nobody specified. A verified artifact MUST NOT be counted as comprehended.

7 What changed and what did not

The argument composes into one position.

Three axes compose.

- On the definition axis, black magic measures the singularity by definition. The empirical content lies in the three premises of Section 2.
- On the evidence axis, the move from Level 0 to Level 1 is historical. Level 2 exists in mathematics. Level 3 is the frontier of learned models. The index is measurable. The index is not saturated.
- On the verification axis, verification separates from comprehension. The verification axis is the least settled. If scalable-oversight methods supervise models above the human level, unverifiable black magic stays rare.

The claim is limited in scope. The claim is not that the singularity is near. The claim is not that most technology is black magic. Most deployed artifacts sit at Levels 0 and 1. The claim is that the share above B measures approach to the singularity. The presentation gap marks where teaching stops helping. The assimilation lag marks the onset.

The budget B, Kolmogorov invariance and the slow-growth law did not change. No notation teaches a human an artifact whose minimal description exceeds B_{store}. No teaching unfolds a depth beyond B_{time}.

Two things changed. Presentation improves. The work of comprehension passes to machines.

Clarke's third law states that any sufficiently advanced technology is indistinguishable from magic (Clarke 1973). This paper defines a sufficiently advanced artifact as one whose minimal humanly graspable account exceeds B.

8 Objections, and what closes each

Each objection has a status and a reason.

Table 5 states each objection, its status, and the mechanism that answers it or the reason it stays open.

Objection	**Status**	**Reason**	**Where**
Circularity	answered	The definition is operational. The empirical claims sit in three testable premises.	§2
Random noise is incomprehensible	answered	Logical depth assigns low complexity to random objects. Noise is shallow.	§4.2, §4.4
The measure is anthropocentric	answered	Interpretability and the singularity are both defined relative to a human observer. The index makes no claim about intelligence.	§2
Machines verify it	answered	Verification certifies a stated property. Comprehension predicts beyond it. Verification creates Level 2.	§6
Humanity has tools and teams	answered	B is biological. A tool that yields an explanation within B lowers the presentation bound. A tool that holds the understanding is the counted handover.	§3, §6
The regime has arrived	answered	Onset is a divergence of the assimilation lag. The observed lag is 33 days.	§2.1
The measurement	partly answered	The weighting is fixed. A six-artifact pilot places each case with upper bounds. The decisive gaps span 10^6 to 10^{15}. A full registry does not exist.	§2.1, §4.4
Why-gaps close historically	open	Carnot followed Newcomen by 112 years. Only the threshold test settles a given gap. The test has not been run on the method.	§4.5
The choice of definition	open	Under Good's or Bostrom's definition the index measures an adjacent quantity. Convergence of definitions is argued, not shown.	§2
Novelty	open	Humphreys, Yampolskiy and Arbesman are the nearest neighbours. The claim is restricted to the measured budget, the index and the partition.	§5.4

Table 5: Objections. Each answered row names its mechanism. Each open row names why only measurement closes it.

The open rows are empirical. Argument does not close an empirical row.

Six rows are answered. One row is partly answered. Three rows are open. The research programme starts with a registry of machine-produced results tracked for assimilation lag.

References

Ameisen, Emmanuel, et al. 2025. "Circuit Tracing: Revealing Computational Graphs in Language Models." *Transformer Circuits Thread*. https://transformer-circuits.pub/2025/attribution-graphs/methods.html.

Arbesman, Samuel. 2016. *Overcomplicated: Technology at the Limits of Comprehension*. New York: Current.

Barnes, Beth, and Paul Christiano. 2020. "Debate Update: Obfuscated Arguments Problem." AI Alignment Forum.

Bennett, Charles H. 1988. "Logical Depth and Physical Complexity." In *The Universal Turing Machine: A Half-Century Survey*, edited by Rolf Herken, 227–257. Oxford: Oxford University Press.

Bills, Steven, Nick Cammarata, Dan Mossing, Henk Tillman, Leo Gao, Gabriel Goh, Ilya Sutskever, Jan Leike, Jeff Wu, and William Saunders. 2023. "Language Models Can Explain Neurons in Language Models." OpenAI.

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.

Burns, Collin, Pavel Izmailov, Jan Hendrik Kirchner, et al. 2024. "Weak-to-Strong Generalization: Eliciting Strong Capabilities with Weak Supervision." In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, PMLR 235.

Carnot, Sadi. 1824. *Réflexions sur la puissance motrice du feu*. Paris: Bachelier.

Chalmers, David J. 2010. "The Singularity: A Philosophical Analysis." *Journal of Consciousness Studies* 17 (9–10): 7–65.

Cherniak, Christopher. 1986. *Minimal Rationality*. Cambridge, MA: MIT Press.

Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.

Clarke, Arthur C. 1973. *Profiles of the Future: An Inquiry into the Limits of the Possible*. Rev. ed. New York: Harper & Row.

Cowan, Nelson. 2001. "The Magical Number 4 in Short-Term Memory: A Reconsideration of Mental Storage Capacity." *Behavioral and Brain Sciences* 24 (1): 87–114.

Doshi-Velez, Finale, and Been Kim. 2017. "Towards a Rigorous Science of Interpretable Machine Learning." arXiv:1702.08608.

Dumas, Jean-Guillaume, Clément Pernet, and Alexandre Sedoglavic. 2025. "A Non-Commutative Algorithm for Multiplying 4×4 Matrices Using 48 Non-Complex Multiplications." arXiv:2506.13242.

Good, I. J. 1965. "Speculations Concerning the First Ultraintelligent Machine." *Advances in Computers* 6: 31–88.

Gorenstein, Daniel. 1983. *The Classification of Finite Simple Groups*, Vol. 1: *Groups of Noncharacteristic 2 Type*. New York: Plenum Press.

Heule, Marijn J. H., Oliver Kullmann, and Victor W. Marek. 2016. "Solving and Verifying the Boolean Pythagorean Triples Problem via Cube-and-Conquer." In *Theory and Applications of Satisfiability Testing – SAT 2016*, Lecture Notes in Computer Science 9710. Springer. doi:10.1007/978-3-319-40970-2_15.

Humphreys, Paul. 2009. "The Philosophical Novelty of Computer Simulation Methods." *Synthese* 169 (3): 615–626.

Irving, Geoffrey, Paul Christiano, and Dario Amodei. 2018. "AI Safety via Debate." arXiv:1805.00899.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the 'Death of the Renaissance Man': Is Innovation Getting Harder?" *Review of Economic Studies* 76 (1): 283–317.

Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. "Scaling Laws for Neural Language Models." arXiv:2001.08361.

Kurzweil, Ray. 2005. *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking.

Kwa, Thomas, et al. 2025. "Measuring AI Ability to Complete Long Tasks." arXiv:2503.14499.

Landauer, Thomas K. 1986. "How Much Do People Remember? Some Estimates of the Quantity of Learned Information in Long-Term Memory." *Cognitive Science* 10 (4): 477–493.

Landauer, Thomas K. 1988. "An Estimate of How Much People Remember, Not of Underlying Cognitive Capacities." *Cognitive Science* 12 (2).

Li, Ming, and Paul Vitányi. 2019. *An Introduction to Kolmogorov Complexity and Its Applications*. 4th ed. Cham: Springer.

Lieder, Falk, and Thomas L. Griffiths. 2020. "Resource-Rational Analysis: Understanding Human Cognition as the Optimal Use of Limited Computational Resources." *Behavioral and Brain Sciences* 43: e1.

Llama Team, AI @ Meta. 2024. "The Llama 3 Herd of Models." arXiv:2407.21783.

McGinn, Colin. 1989. "Can We Solve the Mind–Body Problem?" *Mind* 98 (391): 349–366.

Nohra, Jad. 2023. *Who is Psychotron*. Kindle ed.; also Audible audiobook, narrated by Christopher Sutton.

Nohra, Jad. 2026. "Learned Control as the Next PID, and Its Information-Theoretic Limits." Preprint.

Schut, Lisa, Nenad Tomašev, Tom McGrath, Demis Hassabis, Ulrich Paquet, and Been Kim. 2025. "Bridging the Human–AI Knowledge Gap through Concept Discovery and Transfer in AlphaZero." *Proceedings of the National Academy of Sciences* 122 (13): e2406675122.

Scientific American. 2015. "The Whole Universe Catalog." *Scientific American* 313 (1): 68. doi:10.1038/scientificamerican0715-68.

Shannon, Claude E. 1951. "Prediction and Entropy of Printed English." *Bell System Technical Journal* 30 (1): 50–64.

Teller, Paul. 1980. "Computer Proof." *Journal of Philosophy* 77 (12): 797–803.

Tymoczko, Thomas. 1979. "The Four-Color Problem and Its Philosophical Significance." *Journal of Philosophy* 76 (2): 57–83.

Valle-Pérez, Guillermo, Chico Q. Camargo, and Ard A. Louis. 2019. "Deep Learning Generalizes Because the Parameter-Function Map Is Biased Towards Simple Functions." In *International Conference on Learning Representations (ICLR)*.

Vinge, Vernor. 1993. "The Coming Technological Singularity: How to Survive in the Post-Human Era." In *Vision-21: Interdisciplinary Science and Engineering in the Era of Cyberspace*, NASA Conference Publication 10129, 11–22.

Walsh, Toby. 2017. "The Singularity May Never Be Near." *AI Magazine* 38 (3): 58–62.

Wilkenfeld, Daniel A. 2019. "Understanding as Compression." *Philosophical Studies* 176 (10): 2807–2831.

Yampolskiy, Roman V. 2020. "Unexplainability and Incomprehensibility of AI." *Journal of Artificial Intelligence and Consciousness* 7 (2): 277–291.

Zhang, Chiyuan, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2017. "Understanding Deep Learning Requires Rethinking Generalization." In *International Conference on Learning Representations (ICLR)*.

Zheng, Jieyu, and Markus Meister. 2024. "The Unbearable Slowness of Being: Why Do We Live at 10 bits/s?" *Neuron*. doi:10.1016/j.neuron.2024.11.008.