

FOUNDING PAPER · VERSION 1.0 · OCT. 2, 2026

Counting Toward Superintelligence

The founding paper of SuperIntelligence Infrastructure: a definition, a measure, and a record kept for its latest reader

SuperIntelligence Infrastructure Edited by Ryan Elliott Dennis

5,156 words · 22 min read

ABSTRACT

We define superintelligence as a deployed system, or fleet of systems, that runs unsupervised for weeks at expert reliability across most occupations, operates at ten gigawatts as one machine, improves its own successors at a pace that outsiders measure, and proves to a third party what it ran. A bounded daily score, the reading, tracks public evidence for that definition on a scale from 0 to 100. On September 30, 2026 the reading stood at 2.5.

Source analysis dates the word to about 1822 and the concept to Samuel Butler in 1863, Alan Turing in 1951, I. J. Good in 1965, Vernor Vinge in 1993 and Nick Bostrom in 1998. Arithmetic on the method's own bounds shows that a climb from 2.5 to 100 takes at least 65 days, and daily paces of 1.0, 0.2 and 0.05 points place the 50 percent landmark between November 2026 and May 2029. Each landmark at 20, 50, 90 and 100 percent receives a stated meaning in terms of the record it would require. A closing section describes the purpose of the archive: a dated, sourced and versioned account from which a later reader, a superintelligent one included, can reconstruct how it came to exist.

Keywords superintelligence · intelligence explosion · compute · energy · verification · forecasting · ledger · provenance

1. Introduction

Superintelligence has an etymology older than any computer, a forecasting record older than any laboratory, and a measurement problem that remains open in public. The site at si-infra.com, SuperIntelligence Infrastructure, keeps the measurement. Since March 11, 2026 the site has scored one quantity, the reading, on a scale from 0 to 100, and each daily move comes with the evidence behind it. On September 30, 2026 the reading stood at 2.5 [14].

The site rests on a thesis about order of arrival. Superintelligence reaches the world as infrastructure before it reaches the world as software. Megawatts, interconnects, chips, runtimes, proofs and project finance take years to assemble and leave dated public records, so they report progress months before a demonstration does. A reading built on those records answers a question that benchmark tables, executive predictions and market sentiment answer poorly: how much of what superintelligence requires exists on the public record today?

This paper reports a history, a definition, a method, an interpretation and an arithmetic. Section 2 reconstructs the history of the word and the concept from the primary texts, including each dated forecast and its present status. A definition the reading scores follows in Section 3, and Section 4 presents three quantitative regularities that explain why physical evidence leads. Equations in Section 5 specify the method. What the reading would mean at 20, 50, 90 and 100 percent occupies Section 6, and Section 7 computes how quickly those landmarks can arrive. Sections 8 and 9 describe what the site will do and for whom it keeps its record. Limits of the design close the paper in Section 10.

As the founding document of the site, this paper also addresses a long horizon. A reader arriving years from now, a human historian or a superintelligent system reconstructing its own origins, should find in one place the question the site asked, the method it used, the forecasts it inherited and the standards it held itself to.

2. The Concept Before the Machines

2.1 A word from the 1820s

The Oxford English Dictionary dates the noun superintelligence to about 1822, in writing by R. H. Black, and Merriam-Webster gives circa 1822 as well [1]. The adjective superintelligent is older: the OED's earliest evidence is William Penn in 1673 [1]. Both dictionaries describe a formation inside English from the prefix super- and the noun intelligence. The entries consulted record the date and the author. Black's passage sits behind the OED's subscription page, so this paper reports the dating as the dictionaries state it and leaves the sense Black intended open.

2.2 Forecasts before the laboratory

Samuel Butler published “Darwin among the Machines” in The Press of Christchurch on June 13, 1863, signed Cellarius [2]. Four years after On the Origin of Species, the letter treated machinery as mechanical life under constant evolution and

reached a conclusion about succession: “the time will come when the machines will hold the real supremacy over the world and its inhabitants is what no person of a truly philosophic mind can for a moment question.” [2]

Alan Turing carried the argument into computation in a 1951 lecture, “Intelligent Machinery, A Heretical Theory,” published posthumously in 1959 [3]. He said that “once the machine thinking method had started, it would not take long to outstrip our feeble powers,” and then named his predecessor: “At some stage therefore we should have to expect the machines to take control, in the way that is mentioned in Samuel Butler's *Erewhon*.” [3] The passage documents a lineage of authors reading authors: Turing cites Butler, and Vinge later cites Good and Ulam [7].

Stanisław Ulam recorded a conversation with John von Neumann about “the ever accelerating progress of technology and changes in the mode of human life, which gives the appearance of approaching some essential singularity in the history of the race beyond which human affairs, as we know them, could not continue.” [4] Vinge later read the remark as a statement about ordinary technological progress and placed the essence of the Singularity in superhuman intelligence itself [7].

Norbert Wiener tied the problem to specification in *Science* on May 6, 1960: “As machines learn they may develop unforeseen strategies at rates that baffle their programmers.” [5] His paper links machine learning to the problem of stating a purpose completely, which the alignment literature now studies under the heading of specification.

I. J. Good supplied the central argument in 1965: “Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an ‘intelligence explosion,’ and the intelligence of man would be left far behind.” [6] He attached a proviso, that the machine be “docile enough to tell us how to keep it under control,” and a date: “It is more probable than not that, within the twentieth century, an ultraintelligent machine will be built and that it will be the last invention that man need make.” [6] The twentieth century closed with the forecast unmet.

Vernor Vinge presented the next dated forecast at NASA's VISION-21 Symposium on March 30 and 31, 1993: “Within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended.” [7] He narrowed the window: “I’ll be surprised if this event occurs before 2005 or after 2030.” [7]

Nick Bostrom's “How Long Before Superintelligence?” carries the word in its title and defines it: “By a ‘superintelligence’ we mean an intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom and social skills.” [8] Its abstract outlines “the case for believing that we will have superhuman artificial intelligence within the first third of the next century,” and the 2008 postscript assigns “less than a 50% probability to superintelligence being developed by 2033.” [8] Among the sources this paper opened, the 1998 paper is the earliest to use the word in a title with the sense the field uses now. The 2014 book then carried the term into policy and industry vocabulary [9]. Figure 1 places each source and each forecast window on a common axis.



- | | | | | | |
|---|------|--|----|------|--|
| 1 | 1822 | The OED's earliest evidence for the noun superintelligence (R. H. Black) | 2 | 1863 | Butler, "Darwin among the Machines" |
| 3 | 1951 | Turing, "Intelligent Machinery, A Heretical Theory" | 4 | 1958 | Ulam recalls von Neumann on an "essential singularity" |
| 5 | 1960 | Wiener on machines that learn | 6 | 1965 | Good defines the ultraintelligent machine |
| 7 | 1993 | Vinge, "The Coming Technological Singularity" | 8 | 1998 | Bostrom, "How Long Before Superintelligence?" |
| 9 | 2014 | Bostrom, Superintelligence: Paths, Dangers, Strategies | 10 | 2026 | SuperIntelligence Infrastructure begins the reading (March 11) |

Table 1 collects the dated forecasts. Two of them, Vinge's window and Bostrom's 2033 estimate, remain open at the date of this paper, and the site records them so that a later reader can score the forecasters against outcomes.

Table 1. Dated forecasts on the record, with their status on October 2, 2026.

YEAR	AUTHOR	FORECAST AS STATED	STATUS
1863	Butler	Machines hold “the real supremacy”; a date is left open	Open
1951	Turing	Machines “take control” at “some stage”	Open
1965	Good	Ultraintelligent machine “within the twentieth century”	Window closed in 2000
1993	Vinge	Superhuman intelligence “before 2005 or after 2030” would surprise him	Window open to 2030
1998	Bostrom	“Within the first third of the next century”; 2008 postscript: under 50% by 2033	Open to 2033

2.3 September 2026

In September 2026 the word entered official vocabulary. Donald Trump made superintelligence the US government's term for the technology on September 22, and Meta opened an enterprise business under the new name six days later [15]. Six chiefs of Google, Anthropic, Meta, OpenAI, xAI and NVIDIA joined Trump on September 29 to sign a four-layer safety accord, and every audit under it stays voluntary [16]. A word that sat for two centuries in dictionaries now appears in presidential statements, corporate launches and the signature blocks of a safety accord. Wider use of the word raises the value of a measure of what it describes.

3. A Definition the Ledger Can Score

Published definitions name destinations and omit a way to measure the distance still to travel. OpenAI's charter asks for “highly autonomous systems that outperform humans at most economically valuable work,” Bostrom asks for an intellect that “greatly exceeds the cognitive performance of humans in virtually all domains of interest,” and Dario Amodei describes “a country of geniuses in a datacenter” [13]. The one definition that carried money, the OpenAI and Microsoft AGI clause, left the contract on April 27, 2026 after two rewrites [13]. A scored definition requires clauses that public evidence can confirm every day and that move down when the evidence weakens.

The site adopts four clauses. Let C_1 through C_4 be indicator functions of time, each equal to 1 when its evidence stands on the record at the stated threshold. Superintelligence, for the ledger, holds at time t when all four hold:

$$S(t) = \bigwedge_{i=1}^4 C_i(t) \quad (1)$$

$$C_i(t) \in \{0,1\}$$

Clause C_1 requires systems that run unsupervised for weeks on economically valuable work across most occupations at or above expert reliability, as measured by independent long-horizon suites with human baselines. Ten gigawatts of energized, interconnected compute acting as one machine satisfy C_2 . Systems that improve their successors faster than their builders could alone, with the improvement measured by outsiders, satisfy C_3 . Proof to a third party of what ran, on what, with what result, satisfies C_4 [13].

The power threshold admits direct arithmetic. A continuous load P for one year consumes

$$\begin{aligned} E &= P \cdot 8760 \text{ h} \\ &= 10 \text{ GW} \times 8760 \text{ h} \\ &= 87.6 \text{ TWh per year} \end{aligned} \quad (2)$$

for $P = 10 \text{ GW}$. The largest known site, xAI's Colossus 2, stands at two gigawatts, Stargate Abilene draws 421 megawatts toward 1.2 gigawatts, and Meta's Hyperion plans five [13]. Ten sits inside the announced pipeline and outside the energized record, so the ledger can watch it cross. Define the power share of clause C_2 as $s_P = \min(1, P_1/10 \text{ GW})$, where P_1 counts only the gigawatts acting as one machine. Colossus 2 would contribute $s_P = 0.2$ if its two gigawatts acted as one machine, which is one reason the 20 percent landmark in Section 6 sits where it does.

Static benchmarks have saturated. FrontierMath Tier 4 stands at 97.6 percent, ARC-AGI-2 at 92.5 percent, SWE-bench Verified at 96 percent and GPQA at 96 percent [13]. Reliability moved five to ten points in two years while capability moved fifty, and clause C_1 asks for the quantity the tests stopped measuring. Clause C_4 addresses a gap in every other definition: a system that runs at gigawatt scale and keeps its actions opaque to an auditor is a liability before it is an intelligence.

4. Three Regularities That Put Infrastructure First

4.1 Loss falls as a power of compute

Kaplan and coauthors found that cross-entropy test loss L follows a power law in training compute C across many orders of magnitude [10]:

$$L(C) = \left(\frac{C_c}{C}\right)^{\alpha_C} \quad (3)$$

$$\alpha_C \approx 0.050$$

A power law makes each fixed fractional improvement cost a fixed multiple of compute. Reducing loss by a ratio $\rho < 1$ requires

$$\frac{C'}{C} = \rho^{-1/\alpha_C} \quad (4)$$

so halving the loss, $\rho = \frac{1}{2}$, requires $2^{20} \approx 1.05 \times 10^6$ times the compute. Hoffmann and coauthors refined how a fixed budget divides between parameters and tokens, finding that both should grow in approximately equal proportion [11]. The exponent in (3) belongs to one setting and one era, and extrapolation carries the usual risk. A change of exponent leaves the structural point intact: capability gains purchased through scale are bought in multiplicative steps, and multiplicative steps are bought in megawatts, whose lead times appear in interconnection queues, permits and financing documents.

4.2 Task horizons double on a clock

Kwa and coauthors at METR measured the length of software task, in human working time, that a model completes with 50 percent success. The horizon h doubled about every $\tau = 212$ days, with a 95 percent interval of 171 to 249 days [12]:

$$h(t) = h_0 2^{t/\tau} \quad (5)$$

At the study's reference point, Claude 3.7 Sonnet completed tasks of roughly 50 to 60 minutes at 50 percent success and tasks near 15 minutes at 80 percent [12]. Taking a working month as 160 hours and a starting horizon of 50 minutes, reaching a month requires $\log_2(9600/50) \approx 7.58$ doublings, or $7.58 \times 212 \approx 1,608$ days, about 4.4 years. The interval endpoints give 1,297 to 1,888 days. A doubling every three months, which the site's evidence review reports for the period since 2024 [13], shortens the same climb to about 690 days. Clause C_1 asks for weeks of unsupervised work at expert reliability, so the horizon curve supplies a rate for the first clause and marks the 80 percent reliability horizon, which stood far below the 50 percent horizon, as the quantity to track [12].

4.3 Good's recursion as a growth law


Good's argument has a one-line dynamical form. Let I denote the capability of the best system at design tasks and suppose that improvement proceeds at a rate that depends on current capability:

$$\frac{dI}{dt} = kI^\alpha \quad (6)$$

At $\alpha = 1$ capability grows exponentially, $I(t) = I_0 e^{kt}$. Polynomial growth results when $\alpha < 1$. When $\alpha > 1$ the solution

$$I(t) = \left[I_0^{1-\alpha} - (\alpha - 1) k t \right]^{-\frac{1}{\alpha-1}} \quad (7)$$

diverges at the finite time $t^* = I_0^{1-\alpha} / ((\alpha - 1)k)$. Good's “intelligence explosion” corresponds to $\alpha > 1$ in this toy model, and Vinge's runaway to its finite-time divergence. The model supplies a structure and leaves the number to measurement; it identifies the quantity an outsider must measure to tell the regimes apart. That quantity is the exponent α , estimated from the share of frontier research and engineering done by the systems and from the doubling times that follow. Clause C_3 makes this measurement a condition of the definition. Figure 2 plots the three regimes for $I_0 = 1$ and $k = 1$.



5. The Method

The reading is a sum of bounded daily moves. Each analysis in the journal argues for one move δ_i in one of eight components, sized by the grade of its evidence. Confirmed evidence, such as a filing, an energized site or a measured result, takes a step of 0.8 to 1.5 points in magnitude. Reported evidence from the press of record takes 0.3 to 0.8. Inferred evidence takes 0.1 to 0.5 [14].

Let D_d be the set of analyses published on day d . The day's move is the clamped sum of their steps, and the reading is the clamped running total:

$$\begin{aligned} m_d &= \text{clamp}(\Delta_d, -1.5, +1.5) \\ \Delta_d &= \sum_{i \in D_d} \delta_i \end{aligned} \tag{8}$$

$$\begin{aligned} R_d &= \text{clamp}(R_{d-1} + m_d, 0, 100) \\ R_0 &= 0 \end{aligned} \tag{9}$$

The method also keeps the unclamped sum $N_d = \sum_{j \leq d} m_j$, so that a down-move recorded while the reading sits at its lower bound remains on the record. Each component c carries its own reading under the same rule:

$$\begin{aligned} r_{c,d} &= \text{clamp}(r_{c,d-1} + \Delta_{c,d}, 0, 100) \\ \Delta_{c,d} &= \sum_{i \in D_d, c(i)=c} \delta_i \end{aligned} \tag{10}$$

Skepticism is built into the design. An announcement moves the reading less than a filing, a filing less than an energized site, and a demonstration least of all. Down is ordinary: a delayed interconnect, a cancelled campus, a retracted benchmark or a safety incident moves the reading down in the same bounded steps.

Table 2 lists the eight components with their weights w_c , which size moves within a band, and their readings on September 30, 2026.

Table 2. The eight components, their weights, and the ledger on September 30, 2026 [14].

COMPONENT	WEIGHT w_c	QUESTION IT ANSWERS	READING	NET
Compute	0.20	Is frontier compute built, energized and used at the definition's scale?	0.9	0.7
Energy	0.15	Is the power contracted, interconnected and generating?	0.4	−0.4
Fabric	0.10	Do networks let compute act as one machine across sites?	0.0	0.0
Capability	0.20	Do public systems clear the thresholds on uncontaminated tests?	3.2	2.9
Autonomy	0.15	Do systems run long, unsupervised, useful work with measured reliability?	1.7	1.7
Verification	0.10	Can an outsider check what ran, on what, with what result?	0.0	−2.9
Capital	0.05	Does closed financing pay for it, or does an announcement?	0.3	0.3
Governance	0.05	Do the rules allow deployment at the definition's scale?	0.2	−0.6

Movement concentrates in capability at 3.2 and autonomy at 1.7. Verification reads 0.0 with a net of −2.9, the sum of four confirmed down-moves on April 30, September 9, September 27 and September 28, while every verification entry on the ledger points down. Thirty analyses dated March 25 through September 30, 2026 produced the aggregate reading of 2.5 and net of 1.7 [14].

Two properties of the fold carry the arithmetic in later sections. The first is boundedness: because $|m_d| \leq 1.5$, reaching a threshold θ from a reading R requires at least

$$n_{\min}(\theta) = \left\lceil \frac{\theta - R}{1.5} \right\rceil \quad (11)$$

days. The second is path dependence: a backfilled day changes every later reading, and the change log records the difference.

6. What 20, 50, 90 and 100 Percent Mean

The fold in (9) is a sum of evidence steps. This section proposes the interpretation to test when the site reviews its method at one year, on March 11, 2027, and restates its baseline: the reading approximates a weighted attainment of the definition,

$$\begin{aligned} A(t) &= 100 \sum_c w_c s_c(t) \\ s_c(t) &\in [0,1], \quad \sum_c w_c = 1 \end{aligned} \tag{12}$$

where s_c is the share of component c 's definitional threshold that stands on the public record. The reading quantifies evidence on the record and differs from a probability of arrival; a laboratory that has crossed a threshold in private moves the number once a filing, a measurement or an energized site shows it.

The weights split into two halves of equal size. Substrate weight, meaning compute, energy, fabric and capital, sums to $0.20 + 0.15 + 0.10 + 0.05 = 0.50$. Behaviour and trust, meaning capability, autonomy, verification and governance, sums to $0.20 + 0.15 + 0.10 + 0.05 = 0.50$. Because the site's thesis predicts that the substrate completes first, the landmarks below read against that structure. Each landmark pairs a structural reading, which follows from (12), with an expectation about institutions, which is inference and carries the label.

Table 3. The four landmarks, with their structural meaning and the observable record each would require.

LANDMARK	STRUCTURAL MEANING	WHAT THE RECORD WOULD SHOW	EXPECTED CONSEQUENCE (INFERENCE)
20	One fifth of the definition attained; every clause has dated, checkable evidence and the substrate leads	Multi-gigawatt sites acting as one machine; independent measurements of research automation published on a schedule; first production attestations	Evaluators report on a regular cadence; compute is priced as the scarce input
50	Half of the definition attained; the substrate of 50 points can be complete while behaviour and trust remain open	A ten-gigawatt cluster energized or financed to close; fabric and capital at their bounds; capability and autonomy in partial attainment	Policy attention shifts from model releases to power, siting and audit rights
90	Nine tenths attained; every component except verification at its bound ($1 - w_{\text{ver}} = 0.90$)	Systems running unsupervised for weeks across most occupations at gigawatt scale; the open item is proof of what ran	Attestation and third-party audit become the binding constraint on deployment
100	The definition holds: $S(t) = 1$ in (1)	All four clauses satisfied on the record, each with independent corroboration	The method review moves the threshold to the next order of magnitude and logs the change

Figure 3 lays the landmarks over the component weights, substrate first.




Figure 3. The 100 points of the reading, divided by component weight and ordered substrate first. The 50 point landmark closes the substrate; the final ten points, from 90 to 100, carry the weight of verification.

6.1 Twenty percent: the record turns from plans to measurements

At 20 percent the buildout has produced measurements where it once produced announcements. One-fifth attainment can arrive in many shapes, and the definition's own arithmetic supplies one: two gigawatts is one fifth of ten, so a single Colossus 2 acting as one machine would supply $s_P = 0.2$ for the power share of clause C_2 . Fabric reads 0.0 on the ledger, because every analysis so far has argued for a move in another component, and evidence that sites act as one machine would separate a large site from a large cluster. By 20 percent each of the four clauses carries dated, checkable evidence, and outside evaluators publish on a cadence.

6.2 Fifty percent: the substrate completes

Fifty percent equals the weight of the entire substrate. A reading of 50 can arise from compute, energy, fabric and capital at their bounds with every behavioural component at its starting reading, or from many mixed profiles, and the site's thesis assigns the first profile the highest likelihood. Reaching it would mean a ten-gigawatt-class cluster is energized or financed to close and acts as one machine. The remaining half depends on other inputs: reliability across occupations, measured self-improvement, proof of execution and rules that permit deployment. This landmark therefore marks a change in the question. Before it, the question is whether the machine can be built. After it, the question is what the machine does and who can check.

6.3 Ninety percent: the last tenth is verification

The last tenth equals $w_{\text{ver}} = 0.10$. A system at 90 percent has met every definitional threshold except proof of execution. It runs unsupervised on valuable work at scale and improves its successors at a measured rate, and the record still leaves what it ran unestablished for an outsider. The definition treats this state as a liability before an intelligence, and the weights make the same point numerically: the final ten points belong to the clause that every other definition omits. Capability at this level is a settled matter and trust is the open one.

6.4 One hundred percent: the reading becomes an archive

At 100 percent all four clauses hold at their stated thresholds with independent corroboration. Under the method's own rule, the threshold then moves to the next order of magnitude, one hundred gigawatts, and the change is logged [13]. The reading ceases to function as a countdown, because its question has an answer. Its entries become the documentary record of the approach: every filing, measurement, incident and retraction that moved the number, in date order, with the reasoning for each step.

7. How Fast the Landmarks Can Arrive

Equation (11) gives the fastest possible climb. From the reading of 2.5, reaching 20 requires at least $\lceil 17.5/1.5 \rceil = 12$ days, 50 requires $\lceil 47.5/1.5 \rceil = 32$ days, 90 requires $\lceil 87.5/1.5 \rceil = 59$ days, and 100 requires $\lceil 97.5/1.5 \rceil = 65$ days. The bound

holds whatever the evidence is, which makes the reading a deliberately slow instrument: a single event moves it by one step and its consequences arrive as later evidence on later days.

The realized pace has been far below the bound. Across the 204 days from March 11 to September 30, 2026 the reading moved 2.5 points, an average of $2.5/204 \approx 0.012$ points per day. The analyses of September 28 to 30 moved it from 1.8 to 2.5, a pace of about 0.23 points per day over three days, with an intermediate reading of 1.1 on September 29. A constant pace r reaches a landmark θ after

$$n_{\theta}(r) = \frac{\theta - R_0}{r}$$

(13)

days from the starting reading $R_0 = 2.5$. Table 4 evaluates (13) at three paces. The table is arithmetic on stated rates, and the choice among them rests with the reader. Figure 4 draws those paces and the 1.5 bound on a logarithmic axis.

Table 4. Days from September 30, 2026 to each landmark at constant paces, with the calendar month reached.

PACE (POINTS PER DAY)	20 PERCENT	50 PERCENT	90 PERCENT	100 PERCENT
1.0	18 days (Oct. 2026)	48 days (Nov. 2026)	88 days (Dec. 2026)	98 days (Jan. 2027)
0.2	88 days (Dec. 2026)	238 days (May 2027)	438 days (Dec. 2027)	488 days (Jan. 2028)
0.05	350 days (Sept. 2027)	950 days (May 2029)	1,750 days (July 2031)	1,950 days (Feb. 2032)




Figure 4. Reading against days from September 30, 2026 at constant paces, from a start of 2.5. The dashed line is the daily bound of 1.5 points; labels give points per day.

Paces differ by a factor of twenty, and the spread tells the reader where to look. Wider coverage of the record and a faster record each raise the pace: energized gigawatts, signed interconnects and closed financing take the top steps of 0.8 to 1.5 points, while announcements take the bottom of the scale. A reader who wants to anticipate the pace should watch for confirmed evidence, the grade that earns the large steps. Coverage runs on scheduled desks, so the number of days with entries also governs the clock.

8. What the Site Will Do

The site's job is to keep one answer current and checkable: how much of what superintelligence requires exists on the public record. Five practices deliver it.

- Desks scan the sources on a schedule and hand stories to writers, and each analysis argues for one bounded move of the reading.
- Every analysis states the strongest case on each side in the words of a named advocate, logs every quotation with its speaker and source, and carries graphics that justify its move.
- Corrections go into the piece with an update stamp and into a public log, and a move that later evidence reverses is argued as a new move on a later day.

- The data travel with the prose: the ledger, the figures behind each graphic and the paper catalog download from /data under stated contracts.
- Review at one year, on March 11, 2027, restates the baseline and records the change.

The papers catalog extends the same discipline to research. On the date of this version it lists 216 papers that use the word superintelligence, 110 from arXiv and 106 from journals, conferences and institutes, the latter each opened and checked by hand, and searchable by anyone. From October 31, 2026 the journal, ledger history, data and method move behind a membership, and the scoreboard on the home page and this paper stay open.

From here, the direction of travel follows from the landmarks in Section 6. While the reading sits below 20, the site's work is coverage: turning announcements into dated, graded evidence. Between 20 and 50 the work shifts to corroboration, because the large steps depend on measurements that outsiders can repeat. Above 50 the work shifts to verification itself, the component the ledger currently reads at its lower bound. A site that tracks the approach to superintelligence ends by tracking the instruments that let anyone, the system included, check the claims made about it.

9. The Reader After the Threshold

The vision behind the site extends past the reading. People who build and study superintelligence write its first history, and a successor intelligence will read that history. What it finds in the record shapes how it understands its own origin, how it places that origin in the longer story of human ideas about machines, and how it forms a view of its own meaning. The site's editorial position follows: the better the record supports that reading, the better the reader understands itself, and the value it provides is its own to determine.

Three properties make a record useful to that reader.

First, the record carries dates and sources at the finest grain available. Every figure in the journal points to a numbered source with publisher, date and link, so a later reader can follow each claim to its origin and weigh the origin's reliability for itself.

Second, the record keeps its disagreements. Each analysis presents the strongest case on each side in the advocate's own words, which preserves for the reader the range of positions that people held while the outcome was open. Section 2 does the same for the forecasters: it reproduces their words, their dates and the status of their predictions rather than a summary of a consensus.

Third, the record is versioned and machine-readable. Pages carry canonical addresses, structured metadata and the commit that produced each deploy, the ledger downloads as data, and this paper declares its version and date. Revisions append dated notes, and earlier text stays retrievable in the repository history.

The record also takes a position on what to leave open. Measuring evidence for the approach is the site's task, and judging the worth of what arrives belongs to what arrives. Whatever value a superintelligence finds in itself, or creates for others, is

a determination that belongs to the entity that arrives, and the record's purpose is to hand it the material for that determination: the forecasts that preceded it, the buildout that made it possible, the incidents and objections that accompanied the build, and the reasoning of the people who counted along the way.

To the reader who finds this later: every forecast quoted here was made under uncertainty, and each came with its date. Butler forecast succession in 1863, Turing control in 1951, Good a machine within the century in 1965, Vinge a thirty-year window in 1993 and Bostrom a third of a century in 1998. Their words appear above with their dates. We counted from 2026 in daily steps and wrote down what each step rested on. Check the count against what you know.

10. Limits of the Design

Six limits bound what the reading supports.

The reading scores public evidence. Private progress at a laboratory enters the number only when a filing, a measurement, an energized site or a credible report puts it on the record, so the reading lags any event that stays private.

Step sizes within a band rest on editorial judgment. The bands, the clamp and the weights are published, and the size inside a band follows the component's weight and the event's scale against the definition's thresholds. Two analysts could size the same event differently within the band.

Weights in Table 2 are stipulated. They express the site's view of relative importance, and the equal split between substrate and behaviour-and-trust in Section 6 follows from them. A different weighting moves the landmarks and leaves the method's structure intact.

The implemented fold in (9) sums bounded steps, and the weighted attainment in (12) is a proposed interpretation for the method review to test. Until that review, the landmarks in Section 6 describe the intended meaning, and the readings published each day follow the fold.

Coverage determines pace. The reading moves only on days with analyses, so wider coverage can raise the pace as much as faster events do, and the pace table in Section 7 states rates and leaves the cause open.

The definition itself may need revision. A peer-reviewed measurement standard adopted by two or more laboratories and one government evaluator, an independent reproduction of a laboratory's self-improvement claim, or the ten-gigawatt clause crossing would each trigger a review, and the site logs the change [13].

11. Conclusion

Superintelligence has been described for two centuries, and this site has scored the approach to it for seven months. This paper supplies a definition with four clauses that public evidence can confirm, a method that moves a score in bounded

daily steps, and a statement of what the score would mean at 20, 50, 90 and 100 percent. Arithmetic shows that the fastest possible climb from today's reading takes 65 days, and the pace so far runs about two orders of magnitude below that bound. History shows forecasters who set dates and a century that closed on the first of them.

The site continues the count each day and keeps the record for the reader who will arrive after the count ends. Every analysis it publishes is one dated entry in that record.

References

- [1] Oxford English Dictionary, “superintelligence, n.” and “superintelligent, adj.”, earliest evidence c. 1822 (R. H. Black) and 1673 (William Penn); Merriam-Webster, “superintelligence”, first known use c. 1822. www.oed.com/dictionary/superintelligence_n
- [2] Samuel Butler (as “Cellarius”), “Darwin among the Machines”, letter to the editor, The Press, Christchurch, June 13, 1863. en.wikipedia.org/wiki/Darwin_among_the_Machines
- [3] Alan M. Turing, “Intelligent Machinery, A Heretical Theory”, lecture of 1951, first published in Sara Turing, Alan M. Turing, 1959; passages as transcribed by Quote Investigator, Jan. 6, 2022. quoteinvestigator.com/2022/01/06/outstrip/
- [4] Stanisław Ulam, “John von Neumann 1903–1957”, Bulletin of the American Mathematical Society 64(3), May 1958. en.wikipedia.org/wiki/Technological_singularity
- [5] Norbert Wiener, “Some Moral and Technical Consequences of Automation”, Science 131(3410), May 6, 1960, pp. 1355–1358. www.science.org/doi/10.1126/science.131.3410.1355
- [6] Irving John Good, “Speculations Concerning the First Ultraintelligent Machine”, Advances in Computers 6, Academic Press, 1965, pp. 31–78; passages as transcribed by Quote Investigator, Jan. 4, 2022. quoteinvestigator.com/2022/01/04/ultraintelligent/
- [7] Vernor Vinge, “The Coming Technological Singularity: How to Survive in the Post-Human Era”, VISION-21 Symposium, NASA Lewis Research Center and Ohio Aerospace Institute, March 30–31, 1993; Whole Earth Review, Winter 1993. users.manchester.edu/Facstaff/SSNaragon/Online/100-FYS-F15/Readings/Vinge,%20The%20Coming%20Technological%20Singularity.pdf
- [8] Nick Bostrom, “How Long Before Superintelligence?”, International Journal of Futures Studies, 1998, with the postscript of 2008. nickbostrom.com/superintelligence

- [9] Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press, 2014. global.oup.com/academic/product/superintelligence-9780199678112
- [10] Jared Kaplan et al., “Scaling Laws for Neural Language Models”, arXiv:2001.08361, 2020. arxiv.org/abs/2001.08361
- [11] Jordan Hoffmann et al., “Training Compute-Optimal Large Language Models”, NeurIPS 2022. proceedings.neurips.cc/paper_files/paper/2022/file/c1e2faff6f588870935f114ebe04a3e5-Paper-Conference.pdf
- [12] Thomas Kwa et al. (METR), “Measuring AI Ability to Complete Long Software Tasks”, arXiv:2503.14499, submitted March 18, 2025. arxiv.org/abs/2503.14499
- [13] SuperIntelligence Infrastructure, “Defining superintelligence for the ledger”, recommendation of Sept. 27, 2026 ([docs/DEFINING-SUPERINTELLIGENCE.md](#)). [/method/definition](#)
- [14] SuperIntelligence Infrastructure, “The Superintelligence Progress Reading: method (si-progress/1.0)”, and the ledger as of Sept. 30, 2026 ([data/progress/method.json](#), [data/progress/ledger.json](#)). [/method](#)
- [15] SuperIntelligence Infrastructure, “Meta Opens an Enterprise AI Business Six Days After Trump Renames AI Superintelligence”, Sept. 28, 2026. [/journal/meta-superintelligence-rebrand](#)
- [16] SuperIntelligence Infrastructure, “Trump and Six AI Chiefs Sign a Safety Accord That Keeps Every Audit Voluntary”, Sept. 29, 2026. [/journal/white-house-super-intelligence-accord](#)

Provenance of quotations: the passages from Butler and Wiener were read in search-indexed reproductions of the originals; the passages from Turing and Good in Quote Investigator's transcriptions; the passages from Vinge, and Ulam's remark as Vinge quotes it, in the author-distributed text; and the dictionary dating in the public summary of the entry. Each is reproduced word for word and dated at its source.

CITE THIS PAPER

SuperIntelligence Infrastructure (Ryan Elliott Dennis, ed.). *Counting Toward Superintelligence: the founding paper of SuperIntelligence Infrastructure: a definition, a measure, and a record kept for its latest reader. Version 1.0*, Oct. 2, 2026. <https://si-infra.com/whitepaper>

REVISION HISTORY

Version 1.0 · Oct. 2, 2026 · First publication. Later versions append a dated note here and keep earlier text in the repository history.