The Motherfucking Intelligence Explosion

An analysis of intelligence improvement loops, validated successor time and the physical limits of recursive acceleration.

DRAFT PREVIEW DENNIS HEDEGREEN ANALYSE DRAFT-LOCAL
HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-MOTHERFUCKING-INTELLIGENCE-EXPLOSION/

Everyone keeps asking when the singularity arrives.

That may be the wrong clock.

Suppose machine intelligence becomes a thousand times more useful across meaningful cognitive tasks over the next few years.

Research accelerates. Software changes faster. Work that once needed a team becomes callable from a device.

Then suppose the curve flattens.

A year later people complain that the thousand-times-better system is slow.

Children grow up with it. Governments regulate it. Somebody invents a terrible pricing tier for it.

It becomes normal.

So where was the singularity?

The word is often treated as an altitude: somewhere above “human”, “AGI” and “superintelligence” sits a line labelled **SINGULARITY**.

But Vernor Vinge’s 1993 formulation was closer to a regime boundary: a point where the old models stop working and a new reality requires new ones.[1]

That is not an altitude.

It is a boundary.

And an intelligence explosion is not the same thing either.

I. J. Good’s older mechanism was simpler. If a machine becomes sufficiently capable at intellectual work, machine design becomes one of the things it can help improve. Better machines could then help design still better machines.[2]

The important object was not an IQ score.

It was the loop.

And that makes the title of this article less gratuitous than it looks.

For almost all of the lineage that produced us, a new human mind arrived through human reproduction.

Mother.

Fucking.

Birth.

Development.

Generation.

Repeat.

A reconstruction of human generation intervals across the past 250,000 years estimates a sex-averaged interval of about 26.9 years.[3] That is not an evolutionary speed limit. Evolution has no target called intelligence. But human biological inheritance still moves through human generations.

A child is not a software update.

Humans had already begun escaping part of that clock long before computers.

A theorem can outlive its inventor. Techniques can be taught. Writing preserves thought. Science lets one group begin where another stopped.

Culture moved much of cognitive accumulation outside genetic inheritance.[4]

So the simple story — biology was slow, then AI made intelligence fast — is wrong.

The stranger transition is this:

Culture made **knowledge** transmissible outside genes.

Machine intelligence may make pieces of **working cognition** directly callable as computation.

Nothing in that claim requires machine consciousness.

A trained model can be instantiated again without another childhood. Copies can operate in parallel and, in machine-checkable domains, propose candidates, run experiments, score results and try again.

But copying is not improvement.

Ten thousand identical copies are still one idea. Copies exploring different candidates under reliable evaluation are something else entirely.

Replication can become search. Variation plus selection can become improvement.

The question is how much **validated search** those copies can perform before the bottleneck moves somewhere else.

That brings us to the quantity I think we should watch:

Successor Time

Successor Time is the wall-clock time required to move from a defined system A to a validated successor B that is measurably better on a target capability relevant to producing the next successor.

A serious measurement has to identify A and B, the target capability, the wall-clock interval, human and machine contributions, resources consumed, and validation showing that B is genuinely better.

We do not yet have a clean public canonical time series. We have proxies.

If Successor Time falls because today's system helps produce tomorrow's system, and tomorrow's system then causes the next Successor Time to fall again, we have evidence for a self-reinforcing acceleration mechanism.

If increasingly capable systems fail to shorten that loop, then we do not have an intelligence explosion.

We have better tools.

Public evidence in 2026 sits between those worlds. AI systems already make validated contributions to bounded software and optimisation tasks, while frontier labs report delegating more AI-development work to AI without claiming full recursive self-improvement.[5][6]

But the public evidence refuses to give us one clean runaway curve.

In August 2026, METR examined public discovery series across cybersecurity, mathematics and algorithmic optimisation. Vulnerability discovery had accelerated sharply. Mathematics showed some evidence of acceleration. Across seven algorithmic-optimisation series, however, METR found no comparable dramatic acceleration.[7]

That result matters: machine cognition does not accelerate every frontier equally.

The explosion, if it comes, may not have one speed.

Domains with cheap automated feedback can move differently from those where validation is expensive or physical.

METR's July 2026 "expenditure horizon" work makes the same point from another angle. On NanoGPT, newer agents produced re-validated gains, with four tested models showing positive expenditure horizons of roughly $600 to $3,300. METR still judged autonomous agent optimisation to have had only a minimal effect on overall NanoGPT R&D progress so far.[8]

The mechanism exists in pieces without the loop already running away.

Not:

How smart is it?

But:

How much useful research process can we economically hand to it?

Because cognition has started to become an industrial input.

For most of history, useful intelligence was attached to organisms. Now some cognitive services are callable through an app or API.

The interface feels almost immaterial. It is not. Behind it are accelerators, data centres, networks, cooling, capital and power grids.

We did not make intelligence immaterial.

We made pieces of cognition callable.

And callable cognition can run into capacity.

The International Energy Agency projects global data-centre electricity consumption rising to roughly 945–950 TWh by 2030, about twice the mid-2020s level, with AI a major driver. It also stresses a timing mismatch: data centres can be built faster than much of the electricity infrastructure needed to supply them. [9]

Demand rises faster than capacity.

The superintelligence queues.

The future arrives as a waiting list.

The chip still needs a fab.

The data centre still needs electricity.

A biological experiment still has to wait for biology.

And a scientific conjecture can become cheap while validation remains expensive.

Google DeepMind has framed this as a potential **validation bottleneck**: AI agents may increase hypothesis and candidate generation faster than science can verify which results survive contact with reality.[10]

That distinction may determine what kind of explosion is possible.

Software can sometimes close a loop almost entirely inside machines.

Physics cannot.

Biology usually cannot.

Infrastructure definitely cannot.

Intelligence may escape human generation time without escaping physics.

That is also why “a thousand times smarter” is such a weak description.

A thousand times faster at what — code, proofs, experiments, laboratories, or recognising that the metric itself is wrong?

There is no single useful intelligence dial here. There are rates, costs, parallelism, validation, transfer, hardware and energy.

And one question that cuts through the rest:

How much of the next system was produced by the current system?

If that fraction rises while Successor Time falls, and each successor causally shortens the next loop, then the old phrase becomes deserved.

Intelligence explosion.

Whether it also deserves “singularity” depends on what happens to the rest of the system.

If the transition becomes too fast for our economic, scientific and institutional models to describe, Vinge’s boundary becomes relevant.

But crossing a boundary would not condemn humanity to permanent incomprehension. A new regime can acquire standards, prices, bureaucracies, maintenance windows and bad user interfaces.

A capability can be revolutionary in history and boring in use.

Normalisation does not erase the discontinuity.

It is what happens when a discontinuity survives long enough to acquire routines.

And then there is the other mother in the title.

Mother Earth

Artificial intelligence is often narrated as if some alien substance entered nature from outside.

It did not.

The chain is terrestrial:

Earth → chemistry → life → nervous systems → humans → language → culture → science → semiconductors → computation → machine cognition

The silicon, copper and electricity came from here. The mathematics passed through nervous systems that evolved here. The engineers came through mothers.

Earth planned none of it. Evolution is not a ladder. The planet has no research programme.

But the causal lineage remains.

For billions of years, cognitive systems on this planet were carried forward inside living organisms.

Then one species built a second substrate in which parts of cognition could be copied, called on demand, diversified, evaluated, selected and perhaps increasingly used to construct their successors.

Maybe the loop stalls. Maybe validation dominates. Maybe compute, grids, fabs or institutions remain the binding constraints. Maybe recursive self-improvement turns out to be ordinary automation plus diminishing returns.

Those are the variables we should measure.

The interesting question was never whether a machine becomes impressive enough for us to announce Singularity Day.

The question is whether the machinery that produces intelligence has changed its clock.

Mother Earth took billions of years to produce a species capable of building cognition outside the reproductive cycle that produced it.

Now we get to find out how fast the second clock can run.

References

[1]: Vernor Vinge, “The Coming Technological Singularity” (1993). https://cseweb.ucsd.edu/~goguen/misc/singularity.html (https://cseweb.ucsd.edu/~goguen/misc/singularity.html)

[2]: I. J. Good, “Speculations Concerning the First Ultraintelligent Machine,” *Advances in Computers*, vol. 6, pp. 31–88. Publisher page: https://www.sciencedirect.com/science/article/pii/S0065245808604180 (https://www.sciencedirect.com/science/article/pii/S0065245808604180)

[3]: Wang et al., “Human generation times across the past 250,000 years,” *Science Advances* (2023). https://pmc.ncbi.nlm.nih.gov/articles/PMC9821931/ (https://pmc.ncbi.nlm.nih.gov/articles/PMC9821931/)

[4]: Morgan & Feldman, “Human culture is uniquely open-ended rather than uniquely cumulative,” *Nature Human Behaviour* 9, 28–42 (2025); Andersson & Tennie, “Zooming out the microscope on cumulative cultural evolution,” *Humanities and Social Sciences Communications* 10, 402 (2023). https://www.nature.com/articles/s41562-024-02035-y (https://www.nature.com/articles/s41562-024-02035-y) ; https://www.nature.com/articles/s41599-023-01878-6 (https://www.nature.com/articles/s41599-023-01878-6)

[5]: Anthropic Institute, “When AI builds itself” (2026). Anthropic states that a growing share of AI development is delegated to AI systems, while also stating that full recursive self-improvement has not been reached and is not inevitable. https://www.anthropic.com/institute/recursive-self-improvement (https://www.anthropic.com/institute/recursive-self-improvement)

[6]: OpenAI, “The AI policy window is open. We need to act.” (9 September 2026). OpenAI states that fully autonomous recursive self-improvement is not happening today, while AI already accelerates parts of research used to develop and align later models. https://openai.com/index/ai-policy-window/ (https://openai.com/index/ai-policy-window/)

[7]: METR, “Have We Seen an Acceleration in Discoveries?” (14 August 2026). https://metr.org/notes/2026-08-14-llm-contribution-to-discoveries/ (https://metr.org/notes/2026-08-14-llm-contribution-to-discoveries/)

[8]: METR, “Expenditure Horizon: Measuring Optimization Ability, with an Application to NanoGPT” (21 July 2026). https://metr.org/blog/2026-07-21-expenditure-horizon/ (https://metr.org/blog/2026-07-21-expenditure-horizon/)

[9]: International Energy Agency, *Energy and AI* and *Key Questions on Energy and AI*. https://www.iea.org/reports/energy-and-ai/executive-summary (https://www.iea.org/reports/energy-and-ai/executive-summary) ; https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary (https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary)

[10]: Google DeepMind, “Conjecture Machines: AI agents and the new validation bottleneck in science” (July 2026). https://deepmind.google/public-policy/conjecture-machines-ai-agents-and-the-new-validation-bottleneck-in-science/ (https://deepmind.google/public-policy/conjecture-machines-ai-agents-and-the-new-validation-bottleneck-in-science/)

RELATION MEMORY

NEARBY Closing the Loop Is Not Waking Up · Both examine recursive self-improvement. The earlier article separates operational capability from consciousness; this article proposes measuring validated successor time without a consciousness claim.

NEARBY The People Who Understand Bottlenecks · The earlier article maps physical infrastructure and strategic dependencies. This article asks how grids, fabs and validation constrain the speed of intelligence-improvement loops.

CANONICAL URL: HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-MOTHERFUCKING-INTELLIGENCE-EXPLOSION/
PDF VIEW GENERATED: 2026-09-12T19:11:01.336Z