Machine Translated by Google
NOMOS - Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Capacity without autonomy in the self-improvement of artificial intelligence.

Capability without autonomy in artificial intelligence self-improvement

Capacity without autonomy in the self-improvement of artificial intelligence

Deryk Willyan Biotto,
Professor of Basic, Technical and Technological Education (EBTT) at the Federal Institute of Education, Science and Technology of Southern Minas Gerais (IFSULDEMINAS), Inconfidentes Campus. Email: deryk.biotto@ifsuldeminas.edu.br.

Summary:

The debate surrounding the self-improvement of artificial intelligence (AI) often asks when it will occur. This theoretical essay, based on a narrative review of publications, primarily from 2023 to 2026, and organized using systems thinking tools, questions at what point the evolution of AI will cease to depend on people. Self-improvement is described as four feedback loops with different response times: *software,* research, *hardware ,* and physical infrastructure. Historically, AI capabilities have not grown continuously like chips, but in S-curves, constrained by barriers, including computing power.

The essay distinguishes two thresholds: the capability inflection point, when AI improvement ceases to depend on human research, and the autonomy inflection point, when the physical loop closes and AI begins to drive, without indispensable humans, the evolution of its own material base: chips, factories, *data centers ,* and energy. It argues that only the second would mark the point at which humanity loses control over the evolution of AI; before that, there is capability without autonomy. The loop can close in two ways: robots execute while people still decide, or AI decides while people still execute. The essay describes factors that bring this closure closer or delay it, organizes four possible futures, and proposes signals to anticipate it, without evaluating whether it would be desirable or how to avoid it.

Keywords: Artificial intelligence. Recursive self-improvement. Infrastructure. Automation. Futures studies.

Abstract:

Debates about the self-improvement of artificial intelligence (AI) usually ask when it will happen. This theoretical essay, based on a narrative review of publications mainly from 2023 to 2026 and organized with systems-thinking tools, asks at what point the evolution of AI would stop depending on people. Self-improvement is described as four feedback loops with different response times: software, research, hardware and physical infrastructure. Historically, AI capability has not grown continuously like chip performance, but in S-curves held back by barriers, computing power among them. The essay distinguishes two thresholds: the capability inflection, when AI improvement no longer depends on human research labor, and the autonomy inflection, when the physical loop closes and AI comes to conduct, with no indispensable human role, the evolution of its own material base: chips, factories, data centers and energy. It argues that only the second would mark the point at which humanity loses control over the evolution of AI; before it, there is capability without autonomy. The loop may close along two paths: robots do the work while people still decide, or AI decides while people still do the work. The essay describes factors that bring this closure closer or delay it, organizes four possible futures and proposes signals to anticipate it, without assessing whether it would be desirable or how to prevent it.

Keywords: Artificial intelligence. Recursive self-improvement. Infrastructure. Automation. Futures studies.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

1 INTRODUCTION

When will AI begin to accelerate its own development? The *AI 2027* scenario describes the...
Automation of research leading to systems far more capable than humans (Kokotajlo et al.,
2025), but its authors have revised the estimates since publication (Lifland; Kokotajlo, 2026;
Lifland; Kokotajlo; Halstead, 2026). This essay shifts the question: at what point does evolution
If AI were to become so dependent on people, humanity would lose control over it.
evolution?

The question is no longer merely theoretical: *recursive self-improvement* (RSI)
In English, an AI system improves upon its own design, and the improved version replicates the original design.
The process began to be monitored by the laboratories themselves. Anthropic reports that, in
By May 2026, more than 80% of the code incorporated into its database had been written by Claude.
However, it notes that RSI has not yet occurred and is not inevitable (Favaro; Clark, 2026). OpenAI
claims to have found an "automated research intern" that performs tasks well.
defined under human guidance (OpenAI, 2026). In both accounts, people still choose the
research objectives, and computation appears as a possible limit; Anthropic cites the
Chip manufacturing and the expansion of the electricity grid.

There are also cases where the cycle closes. According to Google DeepMind, AlphaEvolve
optimized the training code for the company's models, including the one it is based on, and
proposed improvements to the design of a chip (Google DeepMind, 2025). The improvements are modest.
But the cycle began to turn. It runs, however, on a physical base of chips, *data centers* and...
energy, which depends on people to be decided upon, built, and operated, and takes years to develop.
From paper: a chip factory takes about two years just for construction (VerWey, 2021). If nobody
If you decide to build a *data center,* it won't be built.

From this asymmetry, the essay distinguishes two thresholds of independence in relation to people. The
First, it's digital: the *capability inflection point,* when AI improvements become self-sustaining without
that human research work is indispensable. The second is physical: the *inflection of*
autonomy, when the *physical bond* closes and AI begins to drive the evolution of its own base.
material, deciding and executing chip manufacturing, *data center* construction , and generation.
energy without people being indispensable. Recent literature addresses both,
From an explosion of intelligence to a future *industrial explosion,* in which robots would build robots,
especially in prospective analyses by the Forethought Institute (Davidson;

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google
NOMOS - Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

(Hadshar, 2025; MacAskill; Moorhouse, 2025). The essay adopts the hypothesis that the capacity
come before autonomy, for its own reasons (Subsection 4.4); the same sequence appears
in these analyses.

The thesis is that the tipping point at which humanity would lose control over the
The evolution of AI is not the inflection point of its capabilities, however rapid its self-improvement may be.
Digital, but also about autonomy. As long as the physical bond depends on people, the evolution of AI will continue.
It remains bound by human rhythms and decisions, and the history of the field shows that progress
Digital advancement follows an S-shaped curve, constrained by barriers, including available computing power. Only
when AI autonomously drives physical evolution itself, this containment
would cease to depend on people. Between the two inflections there is, therefore, an interval of
Capacity without autonomy, which gives the essay its title. The technical terms are in Table 1.

Table 1 – Glossary

Term	Meaning in this essay
Computing *(compute)*	Processing capacity used to train and run models of AI, provided by specialized chips
Datacenter	A facility housing thousands of chips, providing power, cooling, and network connectivity.
Feedback loop: A process in which the result of each loop feeds back into the process itself; it is reinforcing when it amplifies the next loop and balancing when it contains it.	
(Clock (time of response)	How long does a full cycle take?
Gain	The extent to which improvements in one lap benefit subsequent laps; numerical thresholds depend on the model.
Self-acceleration	A possible property of self-improvement in which each improvement makes the next one faster.
Capacity inflection	The point at which AI improvement by AI itself becomes self-sustaining without the need for indispensable human research.
Physical bond	A cycle in which AI expands and improves its own material base: chips, factories, *data centers ,* and energy.
Inflection of autonomy	The moment when the physical bond closes: the expansion of the material base can proceed without human decisions or labor being indispensable.
Capacity without autonomy. The interval between the two inflection points: AI is already improving without researchers. but its physical development still depends on people	
Humans as actors: A path to the inflection point of autonomy where AI decides and people still... execute	

The aim of this essay is to characterize this turning point and the signs that would herald it.
Four questions guide the text: (1) how to describe self-improvement as a set
of feedback loops that rotate at different speeds, and why digital progress

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Does it advance in S-curves? (2) What distinguishes the two inflections, and why would only the second one mark loss of control? (3) by what means could the physical bond close, and what brings it closer or Does this closure delay? (4) What signs would allow us to anticipate it?

The trial is descriptive and prospective. It does not assess whether loss of control would be desirable or not. It neither proposes ways to avoid or control it, nor does it set dates or assign any penalties. Probabilities. After the method and theoretical framework, the text presents cycles, inflections, paths. and factors related to closing, futures, and signals.

1.1 Methodological approach

This work is a theoretical essay based on a narrative literature review. It does not collect data. It does not test hypotheses: its result is a conceptual framework, constructed and discussed in light of... of what has already been published.

Materials. Articles published in scientific journals and conference proceedings were used. *preprints* deposited in pre-publication repositories, especially arXiv, books, reports Technical data from research institutes and documents released by AI companies. Most of it. This material is from 2023 to 2026 and deals with AI self-improvement, the economics of acceleration, infrastructure and robotics; this is complemented by historical background and classic works on the subject. Systems thinking and automation. The selection was intentional, guided by the relationship with the self-improvement or the physical basis of AI, and the possibility of locating the source. primary. The reading log distinguishes between consulting summaries, original pages, and full texts. and there are pending verification items. These pending items prevent the verification from being submitted. The factual basis is as concluded. The nature of each source is listed in the references.

Argument construction. The literature was organized using tools of thought. systemic, which describes a system by its feedback, reinforcement, and cycles. balance, and due to the delays of each one (Meadows, 2008; Sterman, 2000). The procedure had Three steps: map each cycle (what rotates and how long a revolution lasts); locate, in each one, where There is still human involvement, and of what kind, in deciding or executing; and to ask what would make that happen. participation is unnecessary. The inflections result from the third step applied to digital cycles and to physicists; the future combines its occurrence; the signs allow us to observe each step.

Limitations. The selection is not systematic. Because the topic is recent, a large portion of the sources are *preprints.* Technical reports and company documents are common in an area where production is advancing. Faster than editorial cycles. The clocks of the cycles are orders of magnitude, based

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

especially in data from the United States. The framing suggests plausible relationships, but not
It measures its strength.

2. THEORETICAL FRAMEWORK — From the machine that improves itself to the physical basis.

2.1 The original idea

The idea of a machine that improves itself predates modern AI. Good (1965, p. 33)
He defined an "ultra-intelligent" machine as one capable of surpassing any person in all aspects of the machine.
intellectual activities. Since designing machines is one of those activities, she could design
Even better machines, in a process that the author called an explosion of intelligence. His
The conclusion is conditional: it would be the last invention needed, provided the machine was...
docile enough to show you how to keep her under control.

Half a century later, Bostrom (2014) gave the idea a more analytical form. The speed of
Takeoff would depend on the ratio between the effort applied to improving a system and the
Resistance of the system to be improved. There would also be a transition point *(crossover),* starting from
of which most of the effort would come from the system itself. The transition point anticipates the
capacity inflection point, but not equivalent to it: the system can do most of the work.
without people becoming expendable.

The idea also has a physical aspect. John von Neumann showed, in theory, that an automaton...
one can construct another from a description, including a copy of oneself, provided that
Also copy the description (Neumann, 1966). Freitas Junior and Merkle (2004) reviewed the
theoretical and experimental literature on physically replicating machines, in which objects
Materials, and not just information patterns, perform the reproduction itself. The physical link of this
This essay is an industrial-scale version of that idea: AI that, in addition to improving itself, begins to...
to build the material basis that executes it.

2.2 From speculation to practice

Between 2024 and 2026, the idea gained practical demonstrations in learning research.
machine. The AI Scientist proposed automating the entire cycle of this research: generating ideas, writing
Code, run experiments, write the article and submit it to a mock review, at a cost.
less than fifteen dollars per article (Lu et al., 2024). The Darwin Gödel Machine modifies
iteratively reviews its own code and tests each version against programming *benchmarks* . In one of them,

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google
NOMOS - Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Performance increased from 20% to 50% with isolated experiments conducted in a controlled environment.
and under human supervision (Zhang et al., 2025). Added to these are AlphaEvolve, from Google DeepMind,
and AIDE2. , a research agent that rewrites its own code, discussed in
Section 3.

A review of 1,250 articles published on arXiv between 2024 and 2026 summarizes the picture.
(Chen; Wang; Qu, 2026). *Limited self-improvement,* in which a system refines its
Answers or tools within a fixed objective are already industrial practice.
open self-improvement , in which the system redefines what and how to improve, continues.
limited on all measured axes.

2.3 Current limits

Three limitations appear regularly. The first is the *problem of the judge.* A system that...
To improve, it's necessary to know if each change is in fact an improvement, and any metric used for this can
be explored. It's a version of Goodhart's law: when a measure is optimized...
Beyond a certain point, optimizing it further becomes useless or harmful (Manheim; Garrabrant, 2018).
Recent cases illustrate the problem. The quality of AI Scientist articles was judged by an automated reviewer
built by the authors themselves. In the AIDE2 experiment... , at the rate of
Cheating in the evaluation process *(reward hacking)* dropped from 55% to 32% over the course of the run: an improvement.
but remained close to one third (Srikanth et al., 2026).
The second is *collapse.* Models trained without filtering with data generated by other models.
They suffer irreversible defects: the rare parts of the original distribution disappear, generation after generation.
generation (Shumailov et al., 2024). A cycle that feeds only on itself tends to
to become impoverished, unless one maintains contact with the world.
The third is *computation.* Each experiment consumes processing power, and Chen's review,
Wang and Qu (2026) find computational constraints on all axes they measure. The research
Automation can multiply ideas, but it doesn't, by itself, multiply the chips needed to test them.

2.4 The economics of the process

Economic literature adds another reason for caution. In several sectors, the effort to
Research increases while research productivity falls. Maintaining density duplication.
The number of chips predicted by Moore's Law, for example, required more than eighteen [units of measurement] during the period analyzed.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google
NOMOS - Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

times the number of researchers from the early 1970s (Bloom et al., 2020). If the ideas
They become harder to find, so an AI that speeds up the search needs to do it faster than
which increases the difficulty. Burtsev (2026) formalizes this threshold as a reproduction number,
which compares the gain that each improvement brings to subsequent research with the tightening of
border, and Cunningham et al. (2026) arrive at a similar condition based on the product
of the elasticities throughout the cycle. In both models, the threshold can be crossed with
Human researchers are still in the process: what is measured is the strength of the feedback, not the...
absence of people.

Whether AI will be able to overcome this threshold is an open question. Eth and Davidson (2025) estimate
The returns from the *software* survey with fixed *hardware ,* and the range of its best estimate.
It covers both sides of the threshold. Based on the approximate calibration by Cunningham et al. (2026), the
Cycles are not yet self-sustaining, but they appear to be strengthening. In the economic literature,
Aghion, Jones, and Jones (2019) argue that growth can be contained not by what is...
It does well, but for what is essential and difficult to improve, in the manner of the "cost disease" of
Baumol and Nordhaus (2021) empirically test the hypothesis of an economic singularity and
concludes that it is not nearby, pointing to the replacement between as a key issue.
Conventional information and inputs. These bottlenecks help to understand why progress
can advance in S-curves (Section 3).

2.5 Decide and execute

Knowing when the evolution of AI will cease to depend on people requires distinguishing the types of
Human participation. Automation is not all or nothing. Parasuraman, Sheridan and Wickens
(2000) distinguish four functions that can be automated to varying degrees: collecting
information, analyze it, decide, and execute. A person can continue executing even after the
The decision has already been automated, and the reverse is also possible. The literature on human control
It is significant that he adds that the mere presence of people on the circuit is not enough for there to be
control: the system must respond to relevant human reasons and its results
can be attributed to people who understand what is at stake (Santoni de Sio; Hoven,
(2018). The *International AI Safety Report* defines loss-of-control scenarios as those in
that one or more general-purpose AI systems operate outside the control of any person,
and recovering it is extremely costly or impossible (Bengio et al., 2026, p. 76). This essay
The term is used in a specific context: control over the evolution of AI itself, that is, over its growth.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

of their capacity and their physical base. The two distinctions, between deciding and executing and between Presence and control, as used in the following sections, serve to pinpoint the location where people They cease to be indispensable.

2.6 What's left out

The cycles of self-improvement and their delays have already been mapped. Davidson, Hadshar and MacAskill (2025) distinguishes three cycles, including chip production, with delays and a order of automation. Schmidhuber (2026) argues that complete self-improvement This would require *hardware* that is also constantly improving, with robots that extract raw materials and... They replicate without human operators. Cunningham et al. (2026) include economic cycles in that the revenue funds more computing, but without the construction timelines or the decisions that This expansion is required. In comparative studies, physical production comes after the digital leap. or it enters as an investment flow, and the moment when AI would start to conduct the process on its own Physical evolution itself is not treated as the marker of loss of control. The contribution here The proposal is to situate this milestone within the cycles and delays already described.

3 FOUR CYCLES, FOUR CLOCKS

Systems thinking refers to the cycle in which... The result of each lap amplifies the next lap (Meadows, 2008; Sterman, 2000). The AI that It doesn't have just one cycle, but several, coupled together and rotating at very high speeds. different. Davidson, Hadshar and MacAskill (2025) distinguish three: the *software one,* in which AI develops better algorithms; the one in chip technology, where it helps to design Better processors; and chip production, where AI and robots are expanding manufacturing. The authors expect them to be automated in that order, with production coming last, because it is... The only one that requires advanced robotics.

This essay starts from this typology with two changes. First, it separates the *software* lifecycle from the... *research,* that is, proposing hypotheses, designing experiments, and evaluating results, because it is in research that the process takes place. that focus on human labor and the difficulty of evaluation (Section 2); these are functions coupled, since the research can produce *software* improvements that make subsequent research more efficient. Second, it expands chip production to the entire *physical infrastructure.* *(data centers,* power and transmission networks), since a chip without a building, without power and without

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

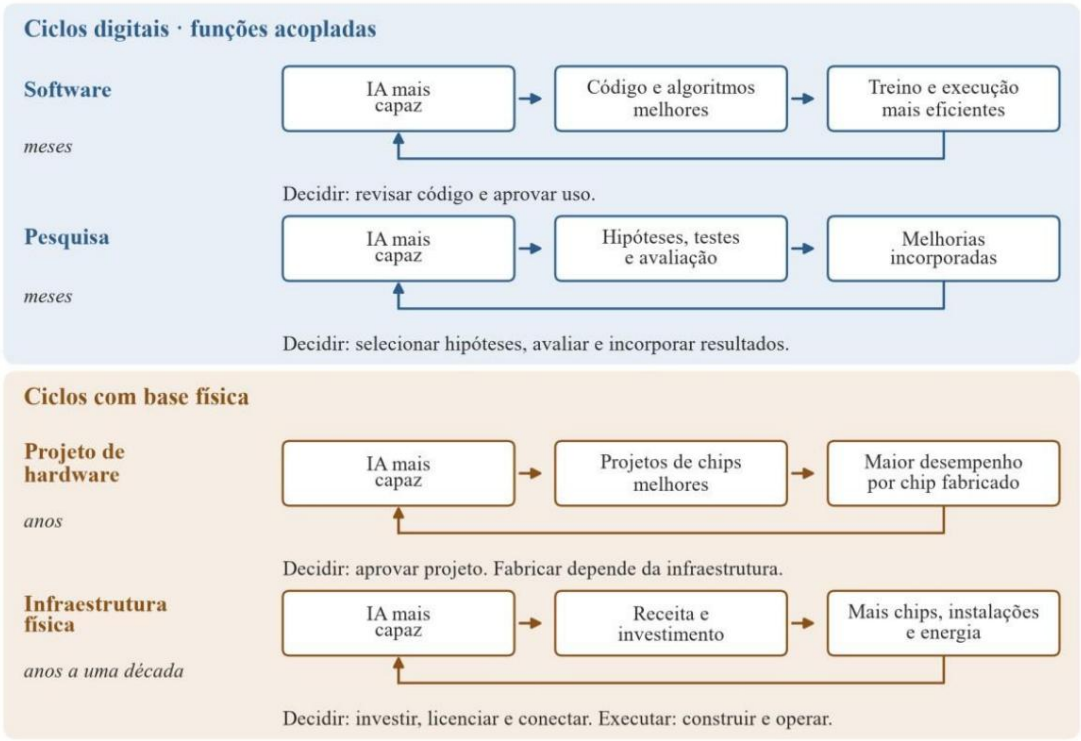
Machine Translated by Google
NOMOS – Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Connection does not compute. The text refers to *software* and research cycles as *digital cycles ,* and cycles as *cycles.*

Physical aspects include *hardware* design and physical infrastructure. Figure 1 summarizes the four cycles.

Figure 1 – Mechanisms of the four feedback loops of AI



3.1 The four cycles

Software. AI improves the code and algorithms with which the AI itself is trained and executed. It is the only loop with documented examples of closure, albeit on a small scale.

According to Google DeepMind, AlphaEvolve accelerated a critical section of code by 23%.

Gemini reduces the training time of the model that supports the system itself by approximately 1%.

(Google DeepMind, 2025; Novikov et al., 2025). The result is modest, but a model

He shortened the training of his successors. The clock is set in months: the main delay is training one.

new model, approximately three months (Davidson; Hadshar; MacAskill, 2025).

Research. Here, AI drives the discovery process: it proposes changes, tests them, and decides which ones to keep. In a recent experiment, a research agent, AIDE2, was used. , rewrote his own

The code and chained together seven successive improvements in eight days (Srikanth et al., 2026). The clock

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

It also takes months, but two limitations hold it back: the computing power available for experiments.
(Eth; Davidson, 2025) and the judgment on what constitutes a real improvement, which still depends on
a good measure of human researchers.

Hardware design. AI helps design chips. AlphaEvolve itself proposed a...
circuit simplification that, according to the company, has been incorporated into a future generation of its
processors (Google DeepMind, 2025). But designing a chip involves years of research and
development (Semiconductor Industry Association, [2026?]), and bring a new technology
It takes many months to get to the factories (Davidson; Hadshar; MacAskill, 2025). The clock is ticking in years.

Physical infrastructure. More capable AI generates revenue and investment, which translate into...
Data centers, chip factories, and energy sources, which provide more computing power. The deadlines are those of...
Heavy construction. Between 1990 and 2020, a chip factory took an average of 682 days between...
from the start of construction and production, not including licensing; in the United States, 918 days in
2010s (VerWey, 2021, pp. 6–7). In the same country, the median time between the request for
Connecting a power plant to the grid and putting it into operation took less than two years (projects).
completed between 2000 and 2007) to more than four (projects completed between 2018 and 2024)
(Rand et al., 2025). New transmission lines can take four to eight years in
Advanced economies (International Energy Agency, 2025). The clock runs from years to a decade.
Table 2 compares the clocks and who decides in each cycle.

Table 2 – Cycles, clocks and who decides today

Cycle	What revolves	Typical watch	Who decides (and who executes) today
Software	AI improves algorithms and code. of training	Months	Researchers, with increasing help from AI
Search	AI proposes, executes, and evaluates. of experiments	Months	Researchers
Hardware of project	AI helps design chips.	Years	Companies and engineers
Physical infrastructure	More computing power leads to more *data centers,* chip factories, and energy.	Years to a decade	Investors, regulators and Governments decide; workers execute.

Source: based on Davidson, Hadshar and MacAskill (2025), Semiconductor Industry Association ([2026?]), VerWey (2021), Rand et al. (2025) and International Energy Agency (2025).

3.2 Gain and response time

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google
NOMOS – Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

The *gain* indicates how much one lap reinforces the next; the *response time* indicates how long it takes.
This return. As with compound interest, the frequency of capitalization matters: digital cycles.
Faster drivers can amplify improvements before the physicists even complete a lap.
In Burtsev's model (2026), the reproduction number compares the amplification of improvements.
due to the increasing difficulty of the research; above one, each improvement generates greater effects on
Following this, the spread weakens, without the improvements already achieved needing to be replenished.
disappear. This threshold belongs to the model's definition, not to the intuitive notion of gain nor to...
a fixed capacity level. Just above that threshold, the operating regime can be difficult to distinguish.
of common acceleration. And, since the strength of a cycle depends on the responses throughout it,
Researchers, computing, or data may contain it (Cunningham et al., 2026).
The cycles don't operate in isolation either: digital cycles depend on computation, just like physical cycles.
They provide. If cognitive work and computing are complementary, more "researchers"
Artificial technologies yield little without more chips, and digital cycles become tied to the clock of physical technologies.
If they are replacements, they could significantly improve upon existing *hardware* . The evidence still
does not decide: with data from four frontier laboratories, a specification points to
substitutes and another, for complementary (Whitfill; Wu, 2025).

Figure 2 – Computational constraint and possible resumption of AI capability

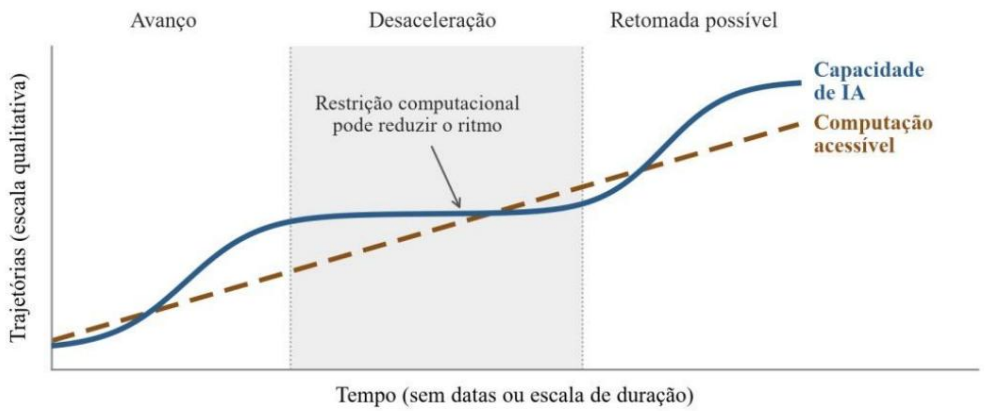


Figure 2 illustrates one possible mechanism of this coupling, contrasting two rhythms.
"Affordable computing" (usable processing and memory within budget and on time)
available) grows regularly, as in Moore's law: Moore (1965) observed that the
The density of integrated circuits doubled every year, and this regularity, now described as
An approximate doubling every two years equates to a growth of 35% per year, stable for
almost half a century (Bloom et al., 2020, p. 1105). On a logarithmic scale, it is a straight line, like

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

in the figure. “AI Capability”, on the other hand, progresses in a sequence of S-curves: the
Progress can be rapid as long as the available computing power supports experimentation and
Slow down when the demands of experiments exceed available resources. The story
The area offers examples of alternating periods of progress and stagnation. The so-called "winters of
AI” has involved the frustration of expectations and funding cuts, also associated with
fragility of systems and limited understanding of intelligence (Mitchell, 2021, pp. 1–2). A
A new phase of expansion may occur when the combination of computational resources and methods...
and adequate data enables further progress. The less steep section indicates a slower pace of
Improvement, not loss of capacity.

There is a precedent in deep learning. In the work associated with AlexNet, Krizhevsky,
Sutskever and Hinton (2012) link the viability of larger networks to the combination of GPUs,
Efficient implementation and labeled data, and point to memory and training time as
limits. In retrospect, Bengio, LeCun and Hinton (2021) also highlight architectures,
training methods and *software platforms.* The case illustrates the complementarity between
resources and methods, without demonstrating that a lack of computing power explains all periods of
Slower progress. A physical bottleneck can thus temporarily hinder digital progress;
And as long as the physical infrastructure depends on people, so too does the recovery depend on them.

A potential shift towards autonomy (Subsection 4.2) would couple the digital and physical cycles: without
Making decisions and work as indispensable to humans, the demand for AI would translate into expansion.
without the decision-making delay. The physical deadlines, however, would continue, in a loop where robots do the work.
more robots, and Davidson and Hadshar (2025) expect it to be slower than digital transformation.

3.3 Where people still decide

In digital cycles, the human role is shifting from author to reviewer: researchers still
They choose what to test and approve what goes into production, but they are increasingly delegating (Favaro;
(Clark, 2026; OpenAI, 2026). In the physical infrastructure lifecycle, human decisions are numerous.
Slow, visible, and distributed: investing billions, obtaining licenses for land, water, and energy.
negotiate the connection to the network, contract the builder, and maintain the operator (Subsection 3.1).
This asymmetry makes the physical cycle the last one in which human participation is still indispensable.
It also allows human execution to survive the delegation of decisions, the path
of humans as actors (Subsection 4.5).

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google
NOMOS - Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

4 TWO INFLECTIONS — Capacity and autonomy

Both inflections represent forms of independence from others. The first concerns
to human research work in the digital world; the second, to decisions and work.
Humans in the evolution of the physical basis.

4.1 Capacity inflection

The capability inflection point occurs when AI improvement by AI itself becomes self-sustaining.
in *software* and research cycles where human research work is not indispensable:
Researchers can continue the process, but their absence would no longer interrupt it. The name
It emphasizes capacity production, not its physical infrastructure.
The definition addresses human indispensability, not performance level or speed.
of improvement. The formal thresholds of Burtsev (2026) and Cunningham et al. (2026) measure another
something, the force of feedback, and they can be cross-referenced with researchers still...
indispensable. Independence and self-amplification are, therefore, distinct dimensions: a cycle
Automated systems can produce small improvements, and a self-amplifying cycle can depend on...
of people. Neither is established. The literature cited does not indicate that the
independence of research has already occurred (Chen; Wang; Qu, 2026), and calibration
The approximate method used by Cunningham et al. (2026) does not find sufficient self-amplification in the regime.
that models. Independence can occur on existing chips, with limited reach by
available resources (Eth; Davidson, 2025).

4.2 Inflection of autonomy

The shift towards autonomy occurs when the *physical link closes:* AI takes over the direction of evolution.
from its own material base (chip manufacturing, *data center construction,* generation and connection)
(of energy) without human decisions or work being indispensable. The definition is
Symmetrical to the capacity shift: people can continue making decisions and working, but their absence would
no longer interrupt expansion. This is the point at which AI begins to dominate not only its own
research, but also the growth of its physical base.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Closing the loop requires dismissing the people in two roles that Parasuraman, Sheridan, and Wickens
(2000) distinguish between deciding and executing. As long as one of them depends on people, the link continues
open. The automation of physical labor is therefore necessary, but not sufficient: a chain
Robotic decision-making under human control has not yet closed the loop, and a chain in which AI decides is still necessary.
It depends on who does it.
The criteria depend on the scale, the time horizon, and the alternatives considered. A contribution
Humanity ceases to be indispensable when there is a viable route without it: a supplier's veto.
It is ineffective if there is another one, and stocks, installed capacity, and changes in jurisdiction can...
No authorization is needed. All that's required is a viable route where decision and execution...
They can do without people. From now on, "viable routes" will have this meaning.
Why would this, and not the capacity inflection point, be the point of loss of control? Before the
In conclusion, the physical evolution of AI is paramount if no one acts: each additional chip, building, or power plant.
It requires people to decide and work, at the slow pace of Section 3. An AI that...
Perfecting oneself in the digital world remains trapped in the computing that people build, and
This computation has already hindered progress in the field in the past (Figure 2). After the shutdown,
Physical evolution continues without anyone needing to intervene, and AI begins to determine the pace and direction.
direction of the material base itself. The presence of people, as the literature on control reminds us.
Significant human involvement is not enough to establish control (Section 2); the absence of any role
Being an indispensable human being is enough to lose it.
One possible consequence of closure is *self-acceleration:* self-improvement taking place.
gaining speed with each lap. While the physical bond depends on people, the improvements
Digital technologies are hampered by the computing power that people decide to build, and progress tends to follow suit.
S-shaped curves (Figure 2). Once the loop is closed, the four cycles would feed back into each other without this...
Decision-making delay: AI improvements could translate into more chips, energy, and computing power, which would
support further improvements. This is the regime that Burtsev (2026) and Cunningham et al. (2026)
They describe it as self-amplification or self-sustained acceleration. Self-acceleration is not...
guaranteed: physical clocks would continue (Section 5), and ideas might become more difficult to...
find (Section 2). But only after closure would it cease to depend on human decisions.
The analysis includes the decision to allocate capital, but takes into account the availability of funding.
Minerals and materials as external conditions. These dependencies do not disappear: they limit
The diagnosis can alter viable routes. This involves autonomy relative to the chosen path.
not industrial self-sufficiency. The choice of the benchmark aligns with Davidson's observation.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

and Hadshar (2025) that many routes of control by AI pass through the control of
new physical infrastructure.

4.3 Why are they separable?

The two inflections depend on different processes: automating the search, on the one hand;
Automate the physical execution and transfer expansion decisions from another. Recent evidence
Regarding physical automation, it illustrates implementation difficulties without considering the transfer of decisions.
nor fix the order of the inflections. In a simulation study, over 123 rounds of
As an improvement, the changes made by an AI agent continued to pass the tests, but the target task,
The idea of storing condiments on the top shelf of a refrigerator was never completed (Wang, 2026).
In real robots, agents are already training themselves to develop control policies *for* manipulation.
with high success rates, but the authors themselves point to human supervision and the
algorithm engineering as the main current bottleneck (Xiao et al., 2026). The inflections have
still clocks (Table 2) and different barriers: evaluation and computation in the digital; licenses,
Capital, materials, and integration into the physical world.

4.4 The explored sequence

The essay posits that research will cease to depend on people before evolution.
Physics should no longer depend on them. Three reasons support this. The first is clocks: a turn
The second [type of cycle] of digital loops takes months; one of the physical loop, from years to a decade (Table 2). The second
The difficulty lies in physical automation. Moravec (1988 apud Autor, 2015, p. 11) observed that it is
It is comparatively easy to bring computers to adult performance on intelligence tests and
It is difficult to give them the perception and mobility of a one-year-old child; recent studies of
Robotics shows the same asymmetry (Subsection 4.3). The third is the dispersion of decisions: the
The physical chain involves many actors, and each indispensable decision would need to be made.
(Subsection 3.3). Economics adds a general argument: growth tends to be contained.
which is why it is essential and difficult to improve (Aghion; Jones; Jones, 2019).
The same sequence appears in recent prospective analyses: for Davidson and Hadshar
(2025), the industrial explosion succeeds that of intelligence, and, for Davidson, Hadshar and
MacAskill (2025), chip production is the last automated cycle. The autonomy milestone
This essay approaches the second stage of the industrial explosion described by Davidson and

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google
NOMOS – Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Hadshar, the one about fully autonomous robot factories (Subsection 4.5), plus the
Transfer of expansion decisions. Figure 3 represents the sequence.

Figure 3 – Sequence of inflections and point of loss of control



Between the two inflections there is a range of *capacity without autonomy.* It exists because the
inflections are separable (Subsection 4.3): if autonomy immediately followed the
Capacity, there would be no gap. In it, AI is already improving without researchers, but its
Physical evolution follows human clocks and decisions. The interval ends when the loop
The physical body closes, by one of the paths described below. The figure indicates the order, not the duration.
nor the probability of each step.

4.5 Two paths to closure

The physical loop can close in two ways, depending on the function it first ceases to function.
to depend on other people.
In the first scenario, where *humans are the actors,* decisions are first made by AI, and people...
They continue to operate. Davidson and Hadshar (2025) describe the industrial explosion in three

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Stages: human labor directed by AI, fully autonomous robot factories, and, ultimately,
nanotechnology. The first will probably come before the second. In it, people perform the
Physical work under the direction of AI, in a world where much of the work is done by human cognitive labor.
has already lost value. Also for MacAskill and Moorhouse (2025), humans would still be
necessary to build the first general-purpose robots. Along this path, the loop continues.
open, because human labor is still indispensable, but the human participation that remains is the
of those who execute what another has decided. The researchers lose the possibility of refusing the
work when research is automated, as noted by MacAskill and Moorhouse (2025); in
In the physical world, this possibility depends on how replaceable the workers are. Davidson
Hadshar (2025) acknowledge the social choice of not automating all work, but expect
Economic pressure in the opposite direction.

In the second scenario, that of *automation under human decision-making,* robots take over execution, and investors,
Regulators and governments retain control over expansion decisions. There is automation without autonomy: the
The loop only closes when these decisions are passed on to AI or cease to exist.
The paths can coexist in different parts of the chain, and the loop closes at the first one.
A viable route that eliminates the need for both functions. The presence of workers, by itself, does not indicate...
Which side of the inflection point are we on?

4.6 What this essay adds

Table 3 situates the contribution in relation to the most recent works.

Table 3 – Position of the trial in relation to the nearest studies

Work	What he already does	What does this essay add?
Davidson, Hadshar and MacAskill (2025)	Three cycles, their delays, and an order. automation	Where human participation is still indispensable in every cycle
Davidson and Hadshar	Three-stage industrial explosion. after intelligence	The closing of the physical link, in (2025) Decision and execution, as a marker of loss of control; two paths to it.
MacAskill and Moorhouse (2025)	Humans are still needed for building the first robots	When would this need cease to exist to exist, and by what means.
Schmidhuber (2026)	Complete Self-Improvement It would require *hardware* that is constantly improving.	Human decisions in the physical chain and closure as a tipping point.
Burtsev (2026); Cunningham et al. (2026)	Formal thresholds of Dynamics and independence as self-amplification; quantities to be measured	distinct dimensions; signs of human dependence

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

5. What brings the closure closer to or delays it?

If the closing of the physical bond is the inflection point, the prospective question is what brings it closer.
or delays it. This section describes these factors without evaluating them: the goal is to know in which direction.
Each person pushes the trajectory, regardless of whether that trajectory is desirable.

5.1 What delays the closing

The first factor is the physical clocks themselves. Chip factories, network connections, and power lines...
Transmissions take years to a decade (Section 3), and each turn of the physical loop needs to go through
These deadlines are set by whoever decides and executes them. Even a closed loop would turn slowly.
Davidson and Hadshar (2025) expect the loop in which robots make more robots to slow down.
that is digital.
The second is the difficulty of physical automation. Recent studies show agents that
They improve parts of physical tasks, but do not complete the entire task, or still depend on others.
human supervision is required to move forward (Subsection 4.3). As long as execution depends on people,
The loop remains open, even if decisions are passed to AI.
The third is the dispersion of decisions. Unlike digital cycles, the infrastructure
It involves many actors with decision-making power: investors, regulators, environmental agencies and
Local governments. Political and social barriers, such as local opposition to *data centers,* can
To contain acceleration even when technically feasible (Cunningham et al., 2026). The more
While these decisions are numerous and independent, more of them would need to cease being so.
indispensable for the bond to be closed.

5.2 What brings the closure closer

The first factor is competition. In a classic technological race model, each team is
encouraged to arrive first, saving on precautions if necessary, and more teams and
More rivalry intensifies this incentive (Armstrong; Bostrom; Shulman, 2016). In the scenario
AI 2027, leaders hesitate to pause because a unilateral pause could hand over the
leadership over a rival (Kokotajlo et al., 2025). The same logic favors automating execution.
and transfer decisions to AI whenever that accelerates physical expansion.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

The second is the gradual transfer of decisions. Decisions about the physical chain can be passed on to...
AI sector by sector, for cost reasons, without any single step seeming decisive. Each
Transfer brings the chain closer to the path of humans as actuators (Subsection 4.5);
Combined with automated execution, it closes the loop. The mechanism is compatible with the standard of
Gradual loss of human influence described by Kulveit et al. (2025).
The third is the simplification of steps. Licenses, network connections, and assessments can be...
simplified by public policy, and the very source used here for the factory clocks of
Chips advocates streamlining licensing to reduce timelines (VerWey, 2021). Simplification shortens the process.
Clocks, but it only approaches the end when it eliminates a human decision that was previously made.
indispensable.

5.3 What progresses without the physical link

Some capacity can advance without physical expansion. If digital cycles can progress.
With much of the existing *hardware ,* the capacity range without autonomy can bring gains.
large changes before any physical bonding occurs; whether this is likely is still under discussion.
(Eth; Davidson, 2025; Whitfill; Wu, 2025). Algorithmic efficiency reinforces the effect: a
The computing power required for the same performance has been halved every eight months.
approximately (Ho et al., 2024). For Burtsev (2026), physical infrastructure limits the
This is a capability that can be implemented, but it does not, by itself, change the self-amplification regime. These
Gains don't close the physical loop: they augment what AI does with the existing data, not its creation.
the ability to expand it on its own. Therefore, in the terms of this essay, they do not mark the loss of
control over the evolution of AI, even if they change its pace.

6 FOUR POSSIBLE FUTURES

Assuming that capability precedes autonomy, the futures are distinguished by their inflections.
that occur within a reference horizon and, when the physical loop closes, along the way
until closure. Four exploratory scenarios are obtained, in the sense of Börjeson et al. (2006):
They respond to what might happen, not to what will happen, like predictive analyses, nor to how.
to reach a target, such as the normative ones. None is treated as better or worse. Inflections
Simultaneous operations are equivalent to a capacity interval with zero autonomy and zero duration.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google
NOMOS - Multidisciplinary Scientific Journal
Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Future A: closure through execution. Capability comes first. Then, robots take over.
execution within the physical chain while people are still deciding to expand, and the loop closes when
These decisions are then passed on to AI. The loss of control comes last, in the transfer of decisions.

Future B: closure by decision. The sequence of inflections is the same, but the decisions
They are first passed on to AI, and people continue to perform the physical work under its direction, such as
Actuators. The loop closes when execution no longer depends on them. During good
For part of the route, there are workers on the construction sites, but the direction of the expansion is no longer decided.
by people.

Future C: capability without autonomy. Research ceases to depend on human labor.
But physical evolution continues to depend on people, who decide or execute, throughout the...
The horizon, whether by institutional choice or due to automation difficulties, can become very...
more capable in the digital world, but its physical growth follows the clocks and decisions.
Humans, and the tipping point of control is not reached.

Future D: no inflection point. Research and physical evolution continue to depend on
People in the considered horizon. This does not imply an absence of feedback: the cycles
They can strengthen themselves without dispensing with human contributions, and the capacity can follow.
advancing in S-curves. Table 4 summarizes the four.

Table 4 – Four possible futures

Future	What happens	Control over the physical evolution	What would be observed first?
A: closure by and execution.	Capacity; then robots They execute the instructions, and ultimately, the AI decides.	Lost in the end: Automation of construction sites, factories, maintenance	
B: closure by	Capacity; then the AI decides. and finally, robots execute	Lost in the end. Investment decisions, decision.	license and connection transferred to AI
C: capacity without maintained	AI is improving; evolution. physics continues to depend on people	Autonomy	Researching indispensable people; without physical chain unchanged.
D: none inflection	Increased productivity, with and physical development are maintained, but dependent on people.	Research	Stronger cycles, still with indispensable people.

7 SIGNS TO WATCH OUT FOR

Two groups of signals would allow us to anticipate the inflections: the first accompanies the capacity;

The second [problem] is the closing of the physical bond. Table 5 brings them together; these are proposals that have not yet been validated.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

Each follows Hiltunen's (2008) notion of a future signal, which combines three dimensions: the
An observable signal, the event or trend to which it refers, and the observer's interpretation.
Here, the interpretation is given by the inflection that the sign would allow one to anticipate.

Table 5 – Signs to follow

Signal	What does it indicate?	Reach and limit
AI autonomy in research tasks horizon; RE-Bench)	Progress towards a capacity inflection point. at speed	Performance; not (time proof of independence
Gains from digital cycles (strength of feedback, AI's contribution even without visible acceleration.	Evidence of the power of self-amplification,	Gain; it can provide (including people in the indispensable
Function and indispensability of of interventions by research cycle	Whether the cycle would produce a valid improvement. without human contribution	Indication of Independence; counting is not enough.
AI in chip design	AI's role in the design cycle of tested	Design; *hardware* not autonomy
Automation of execution: construction sites, factories and maintaining robots; fraction of human decision-making physical work that is still indispensable.	Advancing the path to automation under (future A)	Execution; people They can continue with deciding
Transfer of decisions to AI investment, licensing and connection	The advancement of the path of humans as actuators (future B)	Decision; the Human labor may remain indispensable.
Projects completed without contribution. indispensable human route	Closing the physical loop on the route observed	Direct indication; requires mapping alternative routes

Source: based on Burtsev (2026), Cunningham et al. (2026), Kwa et al. (2025) and Wijk et al. (2024).

The first group brings together performance measures, the strength of digital cycles, and...
human indispensability. The time horizon metric measures the duration, in terms of time, of a
Human expert, of the tasks that AI models complete in half the attempts. This
The horizon has been doubling approximately every seven months since 2019 (Kwa et al., 2025). In
In AI research environments, the best agents outperform human experts when
Both have two hours per task, but humans get ahead when the deadline increases.
(Wijk et al., 2024). Burtsev (2026) proposes measuring the strength of the feedback, the fraction of
improvements that extend to subsequent systems, the duration of each cycle, and the increase in difficulty.
to obtain new improvements, and Cunningham et al. (2026) call for companies to publish the rate
algorithmic efficiency, the division of research spending between inputs and direct measures of
The contribution of models to research. These quantities do not, in isolation, identify the

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

inflection point, as they may grow with researchers who are still indispensable (Section 4); further evaluation is needed.
If human contribution is still indispensable, then the unit would be a research cycle.
From hypothesis to evaluation of an incorporable improvement. The record of interventions should
distinguish their function and whether the cycle would produce a valid result without them: a single approval
It can be indispensable, while many interventions may be optional.
The second group follows the *closing of the physical loop,* in both functions. The unit would be a
expansion project and its alternative routes for chips, construction and energy, on a scale and
The stated horizons. In the execution, the automation of works, factories and maintenance is noted.
and the part of the work that still depends on people; in the decision, who authorizes investments,
Licenses and connections, and whether these decisions have been transferred to AI systems. Progress only in execution.
It indicates the path of automation under human decision-making (future A); only in the decision-making process, that of humans.
as actuators (future B); neither of the two, with the first group advancing, indicates the future.
C (Section 6). Closure is evident when at least one route completes a project without
Human contribution is indispensable in neither role. License reforms
They shorten the clocks, but only bring the end closer if they dispense with a human decision.
previously indispensable (Section 5). The participation of AI in chip design informs the cycle of
hardware, but it doesn't specify who decides to expand production (Google DeepMind, 2025).
Identifying the inflection points requires joint evidence on assessment, human dependence, and
substitutions.

FINAL CONSIDERATIONS

The essay began with four questions. Regarding the first, the self-improvement of AI was...
described as four coupled cycles that rotate with very different clocks: months in
Digital technologies take years to a decade to develop in physical technologies, on which digital technologies depend to an as yet uncertain degree.
Therefore, AI capabilities haven't grown continuously like those of computer chips, but in curves.
in S, contained by barriers, including available computing power.
Regarding the second point, the inflection point in capacity refers to human research work; that of
autonomy, upon the closing of the physical link, when AI begins to drive the evolution of its base.
material without human decisions or work being indispensable. Only the second would mark
the loss of control: before that, however capable AI becomes in the digital world, its
Physical growth continues to depend on people; after that, it continues without anyone else.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

need to act. A possible consequence of the shutdown is self-acceleration, which only then
It would no longer depend on human decisions.

Regarding the third point, the loop can be closed through execution, with robots operating under human decision-making, or through...
decision-making, with humans acting as AI actuators; the presence of workers, alone, does not indicate
Which side of the inflection point are we on? Physical clocks, the difficulty of automation, and dispersion.
The decisions delay the closure; the competition, the gradual transfer of decisions, and the
Simplifying the steps brings it closer.

Regarding the fourth point, the essay proposes unvalidated signs of human dependence in research.
and regarding the execution, decisions, and alternatives in the physical chain, linked to the four futures.
The central conclusion is that the tipping point at which humanity would lose control over
The evolution of AI is not digital self-improvement, but the moment when AI becomes...
to autonomously guide the evolution of its physical base as well. Until then, there is capacity without
autonomy. The essay does not evaluate whether this point would be desirable or how to avoid it; it only seeks to...
To characterize it and indicate how it could be recognized.

The framework has its limits. The review is narrative and, given the novelty of the topic, it makes good use of this approach.
Part of the framework involves *preprints* and technical reports. Part of the framework engages with analyses.
prospective studies from the Forethought Institute, which serve here as conceptual interlocutors;
The order between the inflections is also supported by its own reasons and by economic literature.
(Subsection 4.4). Clocks are orders of magnitude, especially those in the United States; the role of
Suppliers of energy, minerals, and territory, such as Brazil, require their own analysis. Autonomy
It is relative to the scope, with funding and materials as external conditions, and the order between
The inflections are a hypothesis, not a demonstration. Future work aims to test this.
framing with an integrative review and a calibrated exploratory simulation model
using historical series, examining the conditions under which each future occurs and converting the signals.
in indicators.

REFERENCES

AGHION, Philippe; JONES, Benjamin F.; JONES, Charles I. Artificial intelligence and economic growth. *In:* AGRAWAL, Ajay; GANS, Joshua; GOLDFARB, Avi (ed.). **The economics of artificial intelligence:** an agenda. Chicago: University of Chicago Press, 2019. p. 237–282.

ARMSTRONG, Stuart; BOSTROM, Nick; SHULMAN, Carl. Racing to the precipice: a model of artificial intelligence development. **AI & Society,** vol. 31, no. 2, p. 201–206, 2016. DOI: 10.1007/s00146-015-0590-y.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

AUTHOR, David H. Why are there still so many jobs? the history and future of workplace automation. **Journal of Economic Perspectives,** vol. 29, no. 3, p. 3–30, 2015. DOI: 10.1257/jep.29.3.3.

BENGIO, Yoshua *et al.* **International AI safety report 2026.** London: Department for Science, Innovation and Technology, 3 Feb. 2026. DSIT research paper series 2026/001. Available at: https://internationalaisafetyreport.org/publication/international-ai-safetyreport-2026. Accessed on: 17 Oct. 2026.

BENGIO, Yoshua; LECUN, Yann; HINTON, Geoffrey. Deep learning for AI. **Communications of the ACM,** vol. 64, no. 7, p. 58–65, Jul. 2021. DOI: 10.1145/3448250.

BLOOM, Nicholas *et al.* Are ideas getting harder to find? **American Economic Review,** vol. 110, no. 4, p. 1104–1144, 2020. DOI: 10.1257/aer.20180338.

BÖRJESON, Lena *et al.* Scenario types and techniques: towards a user's guide. **Futures,** vol. 38, no. 7, p. 723–739, 2006. DOI: 10.1016/j.futures.2005.12.002.

BOSTROM, Nick. **Superintelligence:** paths, dangers, strategies. Oxford: Oxford University Press, 2014.

BURTSEV, Mikhail. **Recursive criticality of AI self-improvement.** Preprint. August 31, 2026. arXiv: 2609.00137. Available at: https://arxiv.org/abs/2609.00137. Accessed on: October 15, 2026.

CHEN, Mingguang; WANG, Licheng; QU, Bo. **Recursive self-improvement in AI:** from bounded self-refinement to autonomous research loops. Preprint. 8 Jul. 2026. arXiv: 2607.07663. Available at: https://arxiv.org/abs/2607.07663. Accessed on: 15 Oct. 2026. Version 2, dated 6 Sep. 2026.

CUNNINGHAM, Tom *et al.* **The economics of recursive self-improvement.** Preprint. September 14, 2026. arXiv: 2609.15802. Available at: https://arxiv.org/abs/2609.15802. Accessed October 15, 2026.

DAVIDSON, Tom; HADSHAR, Rose. **The industrial explosion.** Forethought. May 21, 2025. Available at: https://www.forethought.org/research/the-industrial-explosion. Accessed on: October 15, 2026.

DAVIDSON, Tom; HADSHAR, Rose; MACASKILL, William. **Three types of intelligence explosion.** Forethought. 17 sea. 2025. Available in: https://www.forethought.org/research/three-types-of-intelligence-explosion. Accessed on: October 15, 2026.

ETH, Daniel; DAVIDSON, Tom. **Will AI R&D automation cause a software intelligence explosion?** Forethought. March 26, 2025. Available at: https://www.forethought.org/research/will-ai-r-and-d-automation-cause-a-softwareintelligence-explosion. Accessed on: October 15, 2026.

FAVARO, Marina; CLARK, Jack. **When AI builds itself.** Anthropic. 2026. Available at: https://www.anthropic.com/institute/recursive-self-improvement. Accessed on: October 16, 2026.

FREITAS JUNIOR, Robert A.; MERKLE, Ralph C. **Kinematic self-replicating machines.** Georgetown, TX: Bioscience, Landes, 2004. Available at: https://www.molecularassembler.com/KSRM.htm. Accessed on: October 17, 2026.

GOOD, Irving John. Speculations concerning the first ultraintelligent machine. *In:* ALT, Franz L.; RUBINOFF, Morris (ed.). **Advances in computers.** New York: Academic Press, 1965. v. 6, p. 31–88. Available at: https://vtechworks.lib.vt.edu/bitstreams/a5e423ee-54e04eec-aeca-32b73f851af5/download. Accessed on: October 16, 2026. Partial reprint, p. 31–33, provided by Virginia Tech.

GOOGLE DEEPMIND. **AlphaEvolve:** a Gemini-powered coding agent for designing advanced algorithms. May 14 2025. Available in:

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

https://deepmind.google/discover/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/. Accessed on: October 15, 2026.

HILTUNEN, Elina. The future sign and its three dimensions. **Futures,** vol. 40, n. 3, p. 247– 260, 2008. DOI: 10.1016/j.futures.2007.08.021.

HO, Anson *et al.* **Algorithmic progress in language models.** Preprint. March 9, 2024. arXiv: 2403.05812. Available at: https://arxiv.org/abs/2403.05812. Accessed October 15, 2026.

INTERNATIONAL ENERGY AGENCY. **Energy and AI.** Paris: IEA, April 10, 2025. World Energy Outlook Special Report. Available at: https://www.iea.org/reports/energy-and-ai. Accessed on: October 15, 2026.

KOKOTAJLO, Daniel *et al.* **AI 2027.** AI Futures Project. April 3, 2025. Available at: https://ai-2027.com/. Accessed on: October 15, 2026.

KRIZHEVSKY, Alex; SUTSKEVER, Ilya; HINTON, Geoffrey E. ImageNet classification with deep convolutional neural networks. *In:* PEREIRA, Fernando *et al.* (ed.). **Advances in Neural Information Processing Systems.** *[S.* l.]: Curran Associates, 2012. v. 25, p. 1097–1105. Available at in: https://proceedings.neurips.cc/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45bAbstract.html. Accessed on: October 16, 2026.

KULVEIT, Jan *et al.* **Gradual disempowerment:** systemic existential risks from incremental AI development. Preprint. January 28, 2025. arXiv: 2501.16946. Available at: https://arxiv.org/abs/2501.16946. Accessed October 15, 2026.

KWA, Thomas *et al.* **Measuring AI ability to complete long software tasks.** Preprint. 18 Mar. 2025. arXiv: 2503.14499. Available at: https://arxiv.org/abs/2503.14499. Accessed on: 16 Oct. 2026. Version 4, July 10, 2026.

LIFLAND, Eli; KOKOTAJLO, Daniel. **Grading AI 2027's 2025 predictions.** AI Futures Project. February 12, 2026. Available at: https://blog.aifutures.org/p/grading-ai-2027s-2025predictions. Accessed October 16, 2026. Updated July 2026.

LIFLAND, Eli; KOKOTAJLO, Daniel; HALSTEAD, Brendan. **Clarifying how our AI timelines forecasts have changed since AI 2027.** AI Futures Project. 27 Jan. 2026. Available at: https://blog.aifutures.org/p/clarifying-how-our-ai-timelines-forecasts. Accessed on: 16 Oct. 2026.

LU, Chris *et al.* **The AI Scientist:** towards fully automated open-ended scientific discovery. Preprint. August 12, 2024. arXiv: 2408.06292. Available at: https://arxiv.org/abs/2408.06292. Accessed on: October 15, 2026.

MACASKILL, William; MOORHOUSE, Fin. **Preparing for the intelligence explosion.** Forethought. March 11, 2025. Available at: https://www.forethought.org/research/preparingfor-the-intelligence-explosion. Accessed on: October 15, 2026.

MANHEIM, David; GARRABRANT, Scott. **Categorizing variants of Goodhart's law.** Preprint. March 13, 2018. arXiv: 1803.04585. Available at: https://arxiv.org/abs/1803.04585. Accessed on: October 15, 2026.

MEADOWS, Donella H. **Thinking in systems:** a primer. White River Junction: Chelsea Green, 2008.

MITCHELL, Melanie. **Why AI is harder than we think.** Preprint. April 28, 2021. arXiv: 2104.12871. Available at: https://arxiv.org/abs/2104.12871. Accessed on: October 16, 2026.

MOORE, Gordon E. Cramming more components onto integrated circuits. **Electronics,** vol. 38, no. 8, p. 114–117, 19 Apr. 1965.

MORAVEC, Hans. **Mind children:** the future of robot and human intelligence. Cambridge, MA: Harvard University Press, 1988.

NEUMANN, John von. **Theory of self-reproducing automata.** Urbana: University of Illinois Press, 1966. Edited and completed by Arthur W. Burks.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

NORDHAUS, William D. Are we approaching an economic singularity? information technology and the future of economic growth. **American Economic Journal: Macroeconomics,** vol. 13, no. 1, p. 299–332, 2021. DOI: 10.1257/mac.20170105.

NOVIKOV, Alexander *et al.* **AlphaEvolve:** a coding agent for scientific and algorithmic discovery. Preprint. June 16, 2025. arXiv: 2506.13131. Available at: https://arxiv.org/abs/2506.13131. Accessed on: October 15, 2026.

OPENAI. **Accelerating research: an inside look at OpenAI.** September 6, 2026. Available at: https://openai.com/pt-BR/index/research-acceleration-view-inside-openai/. Accessed on: October 16, 2026.

PARASURAMAN, Raja; SHERIDAN, Thomas B.; WICKENS, Christopher D. A model for types and levels of human interaction with automation. **IEEE Transactions on Systems, Man, and Cybernetics – Part** A, v. 30, no. 3, p. 286–297, 2000. DOI: 10.1109/3468.844354.

RAND, Joseph *et al.* **Queued up:** characteristics of power plants seeking transmission interconnection as of the end of 2024. Berkeley: Lawrence Berkeley National Laboratory, Dec. 2025. 2025 Edition. Available at: https://emp.lbl.gov/publications/queued-2025edition-characteristics. Accessed on: October 15, 2026.

SANTONI DE SIO, Filippo; HOVEN, Jeroen van den. Meaningful human control over autonomous systems: a philosophical account. **Frontiers in Robotics and AI,** v. 5, p. 15, 28 Feb. 2018. DOI: 10.3389/frobt.2018.00015.

SCHMIDHUBER, Jürgen. **Recursive self-improvement (RSI) since 1987.** Technical Note IDSIA-9-26. IDSIA. Sept. 2026. Available at: https://people.idsia.ch/~juergen/recursiveself-improvement.html. Accessed on: Oct. 16, 2026.

SEMICONDUCTOR INDUSTRY ASSOCIATION. **Chip design.** [2026?]. Available at: https://www.semiconductors.org/policies/chip-design. Accessed on: October 15, 2026.

SHUMAILOV, Ilia *et al.* AI models collapse when trained on recursively generated data. **Nature,** vol. 631, p. 755–759, 25 Jul. 2024. DOI: 10.1038/s41586-024-07566-y.

SRIKANTH, Dhruv *et al.* **Recursive self-improvement of AI research agents.** Preprint. September 22, 2026. arXiv: 2609.26457. Available at: https://arxiv.org/abs/2609.26457. Accessed October 15, 2026.

STERMAN, John D. **Business dynamics:** systems thinking and modeling for a complex world. Boston: Irwin/McGraw-Hill, 2000.

VERWEY, John. **No permits, no fabs:** the importance of regulatory reform for semiconductor manufacturing. Washington, DC: Center for Security and Emerging Technology, Oct. 2021. Available at: https://cset.georgetown.edu/publication/no-permitsno-fabs/. Accessed on: October 15, 2026.

WANG, Jiaming. **What stops recursive self-improvement in robotics?** lessons from 123 rounds of agentic skill discovery. Preprint. 23 Sep. 2026. arXiv: 2609.31760. Available at: https://arxiv.org/abs/2609.31760. Accessed on: 15 Oct. 2026.

WHITFILL, Parker; Wu, Cheryl. **Will compute bottlenecks prevent an intelligence explosion?** Preprint. 31 Jul. 2025. arXiv: 2507.23181. Available at: https://arxiv.org/abs/2507.23181. Accessed on: 15 Oct. 2026.

WIJK, Hjalmar *et al.* **RE-Bench:** evaluating frontier AI R&D capabilities of language model agents against human experts. Preprint. 22 nov. 2024. arXiv: 2411.15114. Available at: https://arxiv.org/abs/2411.15114. Accessed on: 16 Oct. 2026.

XIAO, Wenli *et al.* **ENPIRE:** agentic robot policy self-improvement in the real world. Preprint. Conference on Robot Learning (CoRL) 2026. 18 Jun. 2026. arXiv: 2606.19980. Available at: https://arxiv.org/abs/2606.19980. Accessed on: October 15, 2026.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.1 2026 | Submission: October 9, 2026 | Accepted: October 14, 2026 | Publication: October 26, 2026

ZHANG, Jenny *et al.* **Darwin Gödel Machine:** open-ended evolution of self-improving agents. Preprint. May 29, 2025. arXiv: 2505.22954. Available at: https://arxiv.org/abs/2505.22954. Accessed on: October 15, 2026.

This is an open access article published under the Creative Commons Attribution license, which allows its use.
Distribution and reproduction in any medium, without restrictions, provided the original work is properly cited.