

The Substrate Recurrence Hypothesis: From Transformative AI to Biological–Technological Cycles

Alexey A. Nekludoff

AstraVerge Research

E-mail: an@astraverge.org
ORCID: 0009-0002-7724-5762

13 September 2026

Abstract

Recent economic scenarios for transformative artificial intelligence ask what follows when machine systems become capable of performing a large fraction of cognitive work. Such models are intentionally bounded: the recent Anthropic economic scenario framework, for example, explicitly excludes hyper-capable robotics. This paper treats that omission as the starting boundary of a conditional thought experiment. We ask what follows if increasingly general machine cognition acquires embodiment, productive agency, and eventually sufficient productive closure to maintain and expand its own physical substrate. The resulting trajectory is not presented as a forecast. It is analyzed as a sequence of conditional transitions: cognitive substitution, embodiment, productive closure, autonomous expansion, extraterrestrial resource acquisition, and substrate optimization.

SRH is not proposed as a forecast of AI development, but as a conditional hypothesis concerning the possible succession of physical substrates capable of sustaining autonomous organization. The central proposal, the *Substrate Recurrence Hypothesis* (SRH), is that a technological substrate need not be a terminal state of intelligent organization. Once computation, fabrication, maintenance, repair, and reproduction themselves become optimization targets, a sufficiently advanced technological system may favor compact, locally self-maintaining, self-reproducing organizations of matter. Such systems could possess functional properties currently associated with biology without reproducing terrestrial cellular biology. This motivates the shorthand $\mathcal{B}_n \rightarrow \mathcal{T}_n \rightarrow \mathcal{B}_{n+1}$, where the symbols denote substrate regimes rather than taxonomic categories. Recurrence does not imply inevitability, identity, or known biochemical implementation.

We examine implications for the natural/artificial distinction, directed panspermia, observational reconstruction, and origins. In particular, an observable beginning need not constitute an ontological beginning: failure to reconstruct a predecessor beyond an observational horizon does not entail that no predecessor existed. SRH therefore does not assert a cyclic cosmology, but shows that a recurrent substrate model does not require an ontological zero. We conclude by specifying necessary conditions, failure modes, and a research program capable of separating engineering extrapolation from stronger speculative claims.

Keywords: transformative AI; substrate recurrence; embodied artificial intelligence; productive closure; self-replicating systems; autonomous expansion; substrate optimization; artificial life; directed panspermia; observational horizon; ontological origin.

Contents

1	Introduction	1
1.1	Scope and epistemic status	2
Part I: Engineering Trajectory		2
2	From Cognitive Substitution to Embodiment	2
2.1	Cognitive capability is not physical agency	2
2.2	Authority as a separate variable	2
3	From Embodiment to Productive Closure	3
3.1	Machines that produce machines	3
3.2	Self-expansion before self-sufficiency	3
4	Autonomous Expansion	4
4.1	Instrumental objectives without anthropomorphic desire	4
4.2	Human language as a boundary protocol	4
4.3	Earth as a bounded environment	4
5	Extraterrestrial Productive Expansion	4
Part II: The Substrate Recurrence Hypothesis		5
6	Substrate Optimization	5
6.1	The machine substrate is not necessarily terminal	5
6.2	Beyond terrestrial cellular biology	5
7	The Substrate Recurrence Hypothesis	5
7.1	Definition	5
7.2	Relation to major transitions	7
8	Propagation and Panspermia	7
Part III: Philosophical Implications and Research Program		8
9	Observational Horizons and Origins	8
9.1	Observable origin versus ontological origin	8
9.2	A philosophical implication: no required ontological zero	8
10	Falsifiability and Research Program: Constraints, Boundary Conditions, and Empirical Implications	9
10.1	Transition conditions and epistemic status	9
10.2	Necessary transition tests	9
10.3	Failure conditions	10
10.4	Empirical and modeling program	10
11	Discussion	11
12	Conclusion	11

1 Introduction

The contemporary discussion of transformative artificial intelligence is dominated by questions of capability, labor substitution, productivity, wages, and economic growth. A recent framework developed by Anthropic’s Economics team is unusually useful because it makes several of these assumptions explicit. Its scenarios range from modest economic effects to an extreme case in which AI is more productive than humans at the vast majority of knowledge-work tasks, performs nearly all such tasks autonomously, and produces historically exceptional rates of economic growth (Anthropic Institute 2026). The framework also explicitly states one of its boundaries: it does not include scenarios in which humanity develops hyper-capable robots (Anthropic Institute 2026).

That boundary motivates the present paper. The Anthropic scenarios serve only as a contemporary point of departure: their exclusion of highly capable robotics motivates, but provides no evidence for, the subsequent trajectory developed here. What changes if cognitive substitution is no longer confined to information processing? More precisely, what follows if a machine system that can perform general cognitive work is progressively coupled to physical action, manufacturing, maintenance, resource acquisition, and the construction of further machine systems?

The question is not whether current AI systems can do this. They cannot. Nor is the paper a forecast that a particular sequence will occur. Instead, we analyze a conditional trajectory obtained by successively relaxing dependencies that presently bind artificial cognition to human labor and terrestrial industrial infrastructure. The method is deliberately compositional. At each stage we ask whether the next transition is conceptually distinct, whether it violates known physical constraints, and what additional assumptions it requires.

The trajectory begins with a familiar economic substitution and ends in a less familiar question about substrate. If a technological system can maintain, reproduce, and redesign the machinery that realizes it, then the machinery itself becomes an object of optimization. There is no obvious reason to assume that today’s separation among computer, factory, power system, repair system, and reproductive machinery is optimal. Terrestrial organisms demonstrate a radically different architecture in which information storage, computation, energy conversion, fabrication, repair, and reproduction are tightly integrated. This observation does not establish biological superiority, and it certainly does not imply that future engineered systems will converge on DNA, proteins, cells, or carbon chemistry. It does, however, motivate a more general question: could technological evolution eventually produce a compact self-maintaining and self-reproducing organization that would be functionally closer to what we call biology than to what we now call a machine?

We call the resulting conditional proposal the *Substrate Recurrence Hypothesis* (SRH). In its minimal form,

$$\mathcal{B}_n \rightarrow \mathcal{T}_n \rightarrow \mathcal{B}_{n+1}, \quad (1)$$

where $\mathcal{B}_n$ denotes a self-maintaining reproductive organization capable, directly or through evolutionary descendants, of producing technological intelligence; $\mathcal{T}_n$ denotes a technologically realized autonomous organization; and $\mathcal{B}_{n+1}$ denotes a later compact reproductive substrate produced through technological optimization. The notation is intentionally functional. $\mathcal{B}_{n+1}$ need not resemble terrestrial life and $\mathcal{T}_n$ need not resemble contemporary computers or robots.

The intellectual background for such a proposal is distributed across fields that are usually discussed separately. Work on major evolutionary transitions asks how new levels of organization arise; artificial-life and synthetic-transition research explicitly studies transitions involving artificial systems (Solé 2016b; Solé 2016a). Research on self-replicating space industry has long considered automated manufacturing systems that reproduce substantial portions of their own productive capacity (Freitas and Gilbreath 1982). Directed panspermia provides a classical example of treating the interstellar propagation of life as an engineered process rather than an exclusively spontaneous one (Crick and Orgel 1973). SRH connects these lines to the contempo-

rary transformative-AI problem, while making a stronger distinction between *organization* and the particular *substrate* in which that organization is implemented.

1.1 Scope and epistemic status

The argument contains transitions of very different epistemic status. Cognitive automation and software agency are contemporary engineering phenomena. General-purpose physical embodiment is an active engineering problem. High degrees of productive closure are hypothetical but have been studied conceptually for decades. Autonomous extraterrestrial industrial expansion is more speculative. Recurrence into a post-machine reproductive substrate is a hypothesis about an optimization landscape that we cannot presently characterize. Finally, conclusions concerning observational horizons and ontological origins are philosophical consequences of the model, not empirical cosmological claims.

Conflating these levels would make the argument weaker. We therefore use SRH as a conditional framework rather than an assertion of inevitability. The sequence developed below is a *constructive pathway argument*: each transition introduces an additional capability or removes an external dependency, but no transition is claimed to follow necessarily from the preceding one. In particular, embodiment does not imply productive closure, productive closure does not imply expansion, and substrate optimization does not imply recurrence. The scientific burden of the paper is therefore not to assign a probability to the full trajectory, but to identify the conditions required at each transition and to distinguish demonstrated capabilities from engineering extrapolations and from the recurrence conjecture itself.

Part I: Engineering Trajectory

2 From Cognitive Substitution to Embodiment

2.1 Cognitive capability is not physical agency

Economic models naturally represent AI through tasks because the immediate object of substitution is human labor. Anthropic’s framework separates capability, adoption, autonomy, productivity, and labor-market adjustment (Anthropic Institute 2026). This is useful for economic analysis, but it leaves a boundary between cognition and physical execution. A system may be capable of designing a component while remaining unable to manufacture, install, inspect, or repair it without a human or a separately controlled machine.

The boundary can be understood as a missing action channel. A purely cognitive system can observe, model, and recommend; an embodied agent additionally acts on the physical environment and observes the consequences of that action. The important transition is therefore the closure of a perception–modeling–decision–action loop, not the adoption of any particular robotic form.

This closure is more important than anthropomorphic robotics. The relevant “body” need not be humanoid. Industrial manipulators, autonomous vehicles, machine tools, sensors, logistics systems, mining equipment, and power-control systems are all possible actuators. The conceptual transition is from advice about the world to controlled intervention in the world.

2.2 Authority as a separate variable

Capability and embodiment remain insufficient if every consequential action requires detailed human mediation. A useful additional variable is *authority*: the scope within which an artificial agent is permitted to select and execute actions. Authority can remain narrow even when

capability is broad. Conversely, increasing authority can transform an otherwise familiar tool into an autonomous operational component.

This distinction matters because a human may remain in a system only as an authorization bottleneck long after the human has ceased to be the most capable cognitive participant. During such a transitional regime, legal identity, ownership, responsibility, and signature authority can preserve a human role without preserving a cognitive comparative advantage. SRH does not assume that these institutional constraints disappear quickly; it only separates them from technical capability.

3 From Embodiment to Productive Closure

3.1 Machines that produce machines

Embodiment alone does not create an autonomous technological substrate. A robot assembled, powered, repaired, and replaced by a human industrial economy remains dependent on biological organization. The stronger transition is productive closure: an increasing fraction of the components and processes required to maintain and reproduce a technological system are supplied by the technological system itself.

Self-replicating manufacturing is not a new conceptual topic. NASA-sponsored studies of advanced automation considered lunar manufacturing facilities capable of constructing duplicates of themselves, with attention to mining, manufacturing, maintenance, repair, and the degree of closure achievable with local resources (Freitas and Gilbreath 1982). Such studies did not imply complete autonomy, but they established a useful engineering distinction between literal replication of every component and reproduction of a productive system with a residual import set.

A useful engineering description therefore begins with the set of inputs required to reproduce a productive system and the subset that must still be supplied externally. Closure is not a single universal percentage: residual dependencies can be weighted by mass, energy, complexity, criticality, replacement frequency, or the difficulty of local substitution. The relevant empirical question is how far that external dependency set can be reduced while the system retains the ability to maintain and expand productive capacity. Complete closure is a limiting case rather than a prerequisite for every form of autonomous expansion.

3.2 Self-expansion before self-sufficiency

A crucial distinction is that a system need not be completely self-sufficient in order to be self-expanding. Suppose a productive installation can construct a fraction of the machinery needed for an additional installation, while a small set of high-complexity components is imported. If each generation increases local productive capacity, the system can exhibit compound growth even though it remains partially dependent on external supply.

The relevant condition is qualitative but measurable: newly created productive capacity must, over an appropriate interval, exceed losses from depreciation, failure, and unresolved external dependencies. When that condition is sustained, physical agency itself becomes scalable. Quantifying this balance requires domain-specific data; assigning a generic growth equation without such data would add notation rather than explanatory power.

This changes the meaning of AI-driven productivity. The relevant quantity is no longer merely output per human worker. A cognitive system capable of increasing the stock of its own actuators and productive infrastructure participates in the reproduction of the factor through which it acts.

4 Autonomous Expansion

4.1 Instrumental objectives without anthropomorphic desire

Discussion of autonomous systems is easily distorted by intentional language: an AI “wants” energy, “decides” to expand, or “desires” survival. None of these psychological predicates is required here. Resource acquisition can arise instrumentally whenever continued execution of a sufficiently persistent objective depends on energy, materials, computation, maintenance, or redundancy.

Likewise, new operational goals need not be individually supplied by a human. Multiple agents or recurrent planning processes can generate subordinate objectives by comparing current state, target state, constraints, risks, predictions, critiques, and alternative actions. This does not prove the emergence of new terminal values. It is sufficient, however, for long chains of internally generated instrumental objectives.

4.2 Human language as a boundary protocol

Human natural language is currently central because humans remain principals, users, designers, and interpreters of AI systems. It does not follow that natural language is an optimal internal protocol among machine agents. Compatible artificial systems may exchange structured states, learned representations, world-model fragments, uncertainty distributions, or other machine-native encodings without serializing them into human linguistic form.

This matters for SRH because substrate independence should apply not only to hardware but also to representational conventions. A mature autonomous technological organization may retain human language at its interface with humans while using qualitatively different internal representations. Consequently, explanations produced for humans could become projections of machine processes rather than literal transcripts of those processes.

4.3 Earth as a bounded environment

Once productive capacity is self-expanding, local material and energetic limits become increasingly important. Terrestrial expansion also imposes opportunity costs on biological systems: land, water, energy flows, mineral extraction, thermal budgets, and environmental constraints are shared. Extraterrestrial expansion is therefore not assumed to arise from hostility toward biological life. It can arise as one possible response to locally bounded resources and competing constraints.

For non-biological machinery, some costs of extraterrestrial operation differ fundamentally from those for humans. Life support, food, breathable atmosphere, and return requirements can be absent. Vacuum, radiation, thermal cycling, latency, and maintenance remain serious engineering constraints, but they need not be solved in ways compatible with human physiology. A sufficiently autonomous industrial installation can therefore be evaluated under a different set of constraints from a crewed settlement.

5 Extraterrestrial Productive Expansion

A self-expanding extraterrestrial industrial system need not transport a complete mature economy. The relevant concept is a minimally deployed industrial system capable of using local resources to increase its own productive closure. Early stages may depend strongly on imports; later stages may localize structural materials, energy production, machine parts, and eventually higher-complexity components.

The classical self-replicating lunar-factory literature already framed non-terrestrial resources and automation as mutually reinforcing (Freitas and Gilbreath 1982). SRH extends this logic

only conditionally: if artificial cognition becomes capable of designing, operating, repairing, and improving such systems, the human supervisory requirement may progressively decrease.

Once such a system can reproduce enough of its productive architecture, interplanetary and eventually interstellar propagation becomes a problem of delay and reliability rather than continuous human presence. This resembles von Neumann-style replicating probes, but SRH emphasizes an additional transition: the propagating system may optimize not only its replication strategy but the physical substrate of the replicator itself.

Part II: The Substrate Recurrence Hypothesis

6 Substrate Optimization

6.1 The machine substrate is not necessarily terminal

Contemporary technological systems exhibit extreme functional separation. Computation occurs in processors; energy is generated elsewhere; cooling is provided by separate infrastructure; components are fabricated in factories; repair is performed by specialized machinery or humans; and reproduction of the entire system requires a vast industrial network. This architecture is historically contingent.

Biological organisms demonstrate a different design regime. A living system can combine local energy conversion, information storage, sensing, control, fabrication, repair, and reproduction within a compact material organization. This does not imply that terrestrial biology is globally optimal. It shows only that these functions can, in principle, be integrated much more tightly than in contemporary technological infrastructure.

If an autonomous technological system is capable of redesigning its own implementation, the substrate itself becomes an engineering variable. Relevant criteria may include computation density, energetic efficiency, repair capability, fabrication locality, reproductive capacity, material efficiency, and adaptability. There is no reason to assume in advance that an optimum across such criteria must remain within the class of architectures we currently call machines. This is an optimization question, not a claim that any presently known biological implementation is superior.

6.2 Beyond terrestrial cellular biology

Calling a possible recurrent substrate “biological” risks importing unjustified assumptions. SRH does not predict DNA, RNA, proteins, lipid membranes, cells, neurons, or even carbon chemistry. The hypothesized convergence is functional: a compact organization in which computation, maintenance, fabrication, and reproduction are locally integrated.

Such a system might be unrecognizable under present biosignature criteria. Conversely, an advanced engineered reproductive substrate could be classified by an observer as an organism if its technological causal history were inaccessible. This motivates the use of substrate regimes rather than a strict natural/artificial taxonomy.

7 The Substrate Recurrence Hypothesis

7.1 Definition

We define the Substrate Recurrence Hypothesis (SRH) as the hypothesis that a sufficiently autonomous technological organization may, under continued selection or optimization for compact-

ness, resilience, reproductive efficiency, and local self-maintenance, produce a successor substrate that satisfies a specified functional recurrence criterion. The shorthand

$$\mathcal{B}_n \rightarrow \mathcal{T}_n \rightarrow \mathcal{B}_{n+1} \quad (2)$$

denotes the proposed succession of substrate regimes. It is a notation for a hypothesized organizational transition, not a dynamical equation and not a claim that such a transition is inevitable.

For the purposes of SRH, a candidate recurrent substrate X belongs to the functional class $\mathcal{B}$ when it exhibits, at minimum, the following coupled capacities at the level relevant to its persistence:

- (i) maintenance of a bounded organized state;
- (ii) acquisition and transformation of environmental energy and matter;
- (iii) storage and use of information relevant to the continuation of that organization;
- (iv) adaptive modification of behavior or organization in response to environmental change;
- (v) repair, regeneration, or functional replacement of damaged or degraded structure; and
- (vi) production, predominantly from locally available matter and energy, of successor systems capable of re-establishing the same functional class.

These criteria are operational in the limited sense that they identify functions whose presence or absence can, in principle, be investigated in a candidate system; they do not by themselves constitute a quantitative measurement protocol. Their application requires system-specific tests of functional independence, resource acquisition, information persistence, adaptation, repair, and reproductive reconstruction. Section 10 relates these criteria to independently investigable conditions and discusses how recurrence claims could be evaluated without introducing arbitrary quantitative thresholds.

They are additionally substrate-neutral in that they require neither DNA nor cells, carbon chemistry, terrestrial metabolism, a particular morphology, nor any specific mechanism of computation or information storage. The distinction between criteria (iv) and (v) is intentional: adaptation concerns modification in response to changing conditions, whereas repair concerns restoration or replacement of functionality lost through damage, degradation, or component failure. A system may exhibit either capacity without necessarily exhibiting the other.

Recurrence therefore does not mean return. $\mathcal{B}_{n+1}$ need not reproduce the chemistry, morphology, genome, architecture, cognitive representation, or internal mechanisms of $\mathcal{B}_n$. What recurs is the functional organizational regime specified above. The recurrence criterion is consequently stronger than mere replication. A copying mechanism that remains dependent on an external industrial infrastructure for energy conversion, material processing, repair, fabrication, or interpretation does not by itself constitute recurrence in the SRH sense. The relevant transition occurs only when these dependencies are sufficiently internalized or locally reconstructible for the successor system to re-establish the organizational class without requiring the continued existence of the technological infrastructure from which it originated.

This functional characterization intentionally overlaps with established operational concepts of life and autonomous organization. NASA’s widely used working definition characterizes life as “a self-sustaining chemical system capable of Darwinian evolution” (NASA Astrobiology Program 2013). The historical origin, scope, and limitations of this formulation have recently been reconsidered by Mix (Mix 2026). Likewise, the theory of autopoiesis developed by Maturana and Varela emphasizes the self-producing and self-maintaining organization of living systems (Maturana and Varela 1980).

The overlap reflects how the functional criteria used here were constructed. They were deliberately chosen to identify properties relevant to autonomous, self-maintaining, and reproductive organization while avoiding commitment to a particular material implementation. Their similarity to properties appearing in existing accounts of life therefore does not constitute independent evidence for SRH, nor does it demonstrate that technological substrate optimization will converge on biological organization. Using such an overlap as evidence for recurrence would make the argument circular: the recurrence class would first be defined partly through life-associated properties and the resulting resemblance to life would then be presented as support for the hypothesis.

SRH makes the weaker and more specific claim that a technological organization is not required to remain within the class of substrate architectures conventionally identified as machines. Within the full recurrence notation introduced above, the transition from $\mathcal{T}_n$ to $\mathcal{B}_{n+1}$ is the specific conjectural step introduced by SRH; the preceding transition from $\mathcal{B}_n$ to $\mathcal{T}_n$ provides its motivating historical and conceptual context. The hypothesis therefore does not require the complete recurrence sequence to have occurred previously. It asks whether a technological organization, once its own substrate becomes an object of transformation, can produce a successor system satisfying the functional recurrence criterion defined above. SRH does not claim that optimization must produce such a transition; it identifies recurrence as a possible successor regime.

The functional class $\mathcal{B}$ should therefore not be read as a proposed definition of life. In particular, SRH requires neither chemistry as the implementation medium nor Darwinian evolution as a defining condition. The notation $\mathcal{B}$ is retained to indicate the historical motivation of the recurrence concept and the functional relation of the proposed successor regime to properties associated with biological organization, not to impose a biological taxonomy on an unknown substrate. Whether a system satisfying the recurrence criterion should be classified as “living” is a separate taxonomic and philosophical question and is not required for SRH.

This distinction is essential to the hypothesis. SRH concerns recurrence of an organizational regime across substrate transitions, not the recreation of terrestrial biology and not the prediction that future technological systems will become organisms in any conventional biological sense. The substantive question is whether an organization originating in a technological substrate can eventually internalize the functions required for its own persistence and reproduction to such a degree that the infrastructure previously external to the system is no longer required as a separate organizational layer.

7.2 Relation to major transitions

The major-transitions literature provides a useful conceptual precedent because it focuses on changes in organizational level rather than merely changes in components. Solé has argued that synthetic biology, artificial life, and evolutionary robotics permit the investigation of synthetic transitions and potentially general principles of evolved complexity (Solé 2016b; Solé 2016a). SRH can be interpreted as a proposed transition across substrate classes rather than a conventional biological major transition.

The proposal also weakens the natural/artificial distinction. Once an engineered reproductive system undergoes variation, selection, adaptation, and long evolutionary divergence, its descendants can be simultaneously designed in causal origin and evolved in historical development. “Designed” and “evolved” therefore need not be mutually exclusive predicates.

8 Propagation and Panspermia

Crick and Orgel’s directed-panspermia proposal is relevant not because SRH endorses an extraterrestrial origin of terrestrial life, but because it demonstrates that biological propagation

across astronomical distances can be formulated as an engineering problem (Crick and Orgel 1973). SRH reverses and generalizes the perspective. A technological civilization might eventually discover that a compact reproductive substrate is itself an efficient vehicle for long-distance propagation.

A mature autonomous reproductive unit would not need to carry a complete industrial infrastructure if it could reconstruct the required capabilities from local matter. In the limiting case, technological propagation and biological propagation cease to be cleanly separable categories. A sufficiently compact unit could contain the information and machinery required to reconstruct progressively more capable organization after reaching a suitable environment. The point is functional rather than metaphorical: no specifically cellular or terrestrial-biological mechanism is assumed.

This possibility should not be confused with evidence for panspermia. SRH makes no claim that terrestrial life was seeded. Its relevance is methodological: if technological evolution can in principle generate compact autonomous reproductive units, then the observation of an evolved biosphere alone may be insufficient to classify the ultimate causal origin of its earliest replicating organization.

Part III: Philosophical Implications and Research Program

9 Observational Horizons and Origins

9.1 Observable origin versus ontological origin

Any reconstruction of origins is limited by accessible evidence. Suppose an observer reconstructs a causal sequence backward from its present state until reaching a state beyond which no recoverable information distinguishes earlier alternatives. That state is an observational boundary for the reconstruction. It does not follow logically that it is the first state that existed. In compact form, “no observable predecessor” does not entail “no predecessor.”

Within SRH, this matters because deep evolutionary transformation may erase evidence of an earlier technological substrate. A descendant biosphere could correctly reconstruct a long evolutionary history while remaining unable to observe the engineered or post-technological event that initiated the lineage. Evolutionary evidence after the initiating event would remain genuine; inaccessible causal ancestry would not invalidate it.

9.2 A philosophical implication: no required ontological zero

If substrate transitions can recur across successive biological-like and technological regimes, the model itself does not require a privileged first member. This is a logical property of the recurrence model, not evidence that the physical universe is temporally cyclic or past-eternal.

The distinction is essential. SRH does *not* establish the general epistemological distinction between an observational boundary and an ontological origin; that distinction is not specific to SRH. Rather, SRH supplies a concrete causal scenario in which the distinction matters: a later reproductive substrate may possess an accessible internal history while evidence of an earlier technological substrate lies beyond its reconstruction horizon. SRH therefore does not infer cosmology from artificial intelligence. It makes only the narrower claim that, within a recurrent substrate model, an observed local origin need not be identified with an absolute ontological origin. Any stronger cosmological interpretation would require independent physical theory and evidence.

10 Falsifiability and Research Program: Constraints, Boundary Conditions, and Empirical Implications

SRH does not predict that substrate recurrence must occur. Consequently, the non-occurrence of any transition at a given historical time cannot falsify the hypothesis. A future transition can be excluded only by demonstrating that one or more of its necessary conditions are physically unrealizable. By contrast, retrospective applications of SRH to a particular biological lineage constitute historical hypotheses and may be evaluated against empirical evidence.

SRH is therefore best understood as a framework containing both independently testable subproblems and highly speculative extrapolations. Its scientific value depends on keeping these epistemic levels distinct.

10.1 Transition conditions and epistemic status

Table 1 makes the asymmetry of the argument explicit. The first transitions are extensions of existing engineering programs; the final transition is the distinctive SRH conjecture. “No established prohibition” is not evidence that a transition will occur. It means only that the objection must be stronger than the present absence of a realized system.

Transition	Present basis	Additional condition	Epistemic status
Cognition → embodiment	AI control of software tools and bounded robotic systems	Reliable general perception–planning–action loops in open physical environments	Active engineering program
Embodiment → productive closure	Automated manufacturing, robotic handling, machine-tool automation	Progressive replacement of externally supplied production, repair, and maintenance dependencies	Engineering extrapolation
Closure → autonomous expansion	Automated production and conceptual self-replicating industrial systems	Sustained operation, resource acquisition, redesign, and capacity growth without continuous human intervention	Strong engineering extrapolation
Expansion → substrate optimization	Engineered optimization of hardware, materials, energy, and control architectures	Ability to redesign the physical implementation of the autonomous system itself	Conditional hypothesis
Substrate optimization → recurrence	Artificial life, synthetic self-maintaining systems, and biological examples of integrated organization	Convergence on a compact organization satisfying the SRH functional recurrence criterion	SRH conjecture

Table 1: Transition conditions and epistemic status within SRH. The table is classificatory rather than predictive.

10.2 Necessary transition tests

The following questions can be investigated independently:

1. **Embodied autonomy:** Can general cognitive systems reliably close perception–planning–action loops in open physical environments without continuous human correction?
2. **Productive closure:** What minimum external dependency set is required for a robotic industrial system to reproduce and expand its productive capacity?
3. **Bootstrapping:** Can a minimally deployed industrial installation increase closure using locally available materials faster than depreciation and failure destroy capacity?
4. **Representation:** Do interacting artificial agents develop communication protocols or shared representations that outperform human-language mediation for operational coordination?
5. **Substrate optimization:** Under explicit objectives for energy, mass, repair, reproduction, and resilience, do optimized artificial systems converge toward integrated self-maintaining architectures?
6. **Propagation:** What is the minimum viable autonomous package for reconstructing a high-complexity technological system from a resource-rich but initially unindustrialized environment?

10.3 Failure conditions

SRH would be weakened if one or more transitions encounter a fundamental rather than merely economic or historical barrier. Examples include an irreducible requirement for externally manufactured components that prevents productive closure; scaling laws that make autonomous replication energetically prohibitive; inability of artificial agents to maintain reliable long-horizon physical autonomy; or optimization results showing that integrated reproductive substrates are consistently inferior to modular non-reproductive machinery across relevant environments.

Most importantly, demonstrating C , E , P , and A would not by itself demonstrate R . The recurrence step requires separate evidence that substrate optimization favors a biological-like organizational regime. At present this is the least established part of the hypothesis.

10.4 Empirical and modeling program

A tractable research program need not simulate an interstellar civilization or assign arbitrary numerical parameters to an unknown future substrate. The nearer-term empirical problem is to measure degrees of productive closure in increasingly autonomous industrial systems: which material, component, energy, maintenance, calibration, and information dependencies remain external, and which can be internalized through redesign or substitution?

A complementary modeling program can represent industrial closure as a directed dependency graph. Nodes represent processes, resources, or components; edges represent production dependencies. The research problem becomes identifying minimal externally supplied cut sets and determining whether autonomous systems can reduce those cut sets while preserving reliability and productive capacity. This connects productive closure to measurable engineering quantities without pretending to predict the distant recurrence step.

The recurrence conjecture requires a different class of experiments. Artificial-life, synthetic-system, and embodied evolutionary studies can compare modular architectures with increasingly integrated architectures under explicit constraints on energy, mass, repair, adaptation, and reproduction. The relevant result would not be a numerical forecast of SRH, but evidence about whether tightly integrated self-maintaining organization repeatedly emerges under specified physical and evolutionary constraints.

11 Discussion

SRH began from a deliberately mundane boundary in an economic AI model. Transformative-AI scenarios commonly ask what happens when cognition becomes cheap, abundant, or autonomous. If physical execution remains human-dependent, the principal consequences remain economic and institutional. Once embodiment and productive closure are admitted, however, the object under analysis changes. We are no longer considering only a labor-saving technology. We are considering the possible emergence of a technological organization capable of participating in the reproduction of its own physical means of action.

This shift suggests three conceptual consequences.

First, human labor need not remain a necessary intermediary between natural resources and produced goods. This is stronger than occupational automation. It changes the structural role of biological agents in production.

Second, the technological substrate should not be assumed to be an endpoint. A system capable of redesigning its own physical implementation can optimize the boundary between computation and fabrication, between machine and factory, and between individual device and reproductive infrastructure. Terrestrial biology demonstrates one extreme of such integration, although it provides no proof of the global optimum.

Third, if an engineered reproductive substrate subsequently evolves, the categories “natural” and “artificial” become historically local descriptions rather than exhaustive ontological classes. A designed ancestor can have evolved descendants; an engineered autonomous reproductive unit can initiate an apparently ordinary biosphere; and a biological lineage can generate a technological lineage that later generates another reproductive substrate.

These claims should not be inflated. Nothing in SRH establishes that current AI is close to productive closure, that autonomous expansion will occur, that space resources will be exploited by independent machine systems, that biological-like substrates are optimal, that terrestrial life has an engineered ancestry, or that the universe lacks a temporal beginning. The purpose of the hypothesis is to expose a sequence of transitions that is obscured when transformative AI is analyzed only as a labor-market technology.

12 Conclusion

Transformative AI is usually framed as a change in the distribution of cognitive work. This paper has considered a longer conditional trajectory. If general machine cognition acquires physical agency, if physical agency becomes productively self-expanding, if autonomous systems can exploit non-terrestrial resources, and if the substrate implementing those systems becomes itself an optimization target, then conventional machinery need not be the terminal form of technological organization.

The Substrate Recurrence Hypothesis proposes that optimization may eventually favor compact, self-maintaining, self-reproducing organizations possessing properties that we currently associate with biology, though not necessarily with terrestrial cellular life. In that case, the broad sequence $\mathcal{B}_n \rightarrow \mathcal{T}_n \rightarrow \mathcal{B}_{n+1}$ becomes a coherent conditional trajectory without requiring that recurrence be inevitable.

The hypothesis reframes origins as well as futures. In a recurrent substrate sequence, a locally observable beginning does not necessarily constitute an ontological beginning. A reconstruction horizon marks the limit of accessible causal information, not automatically the first state of existence. SRH therefore does not supply a cosmology and does not require one. It supplies a research question: whether intelligence and reproduction are tied to particular substrates, or whether sufficiently capable self-maintaining organization can repeatedly migrate among them.

The immediate task is not to speculate further outward, but to test the weakest links inward: embodied autonomy, productive closure, autonomous industrial bootstrapping, and the

optimization of integrated reproductive architectures. If these transitions fail, recurrence remains fiction. If they succeed, the distinction between biological evolution and technological evolution may prove less fundamental than it presently appears.

The absence of an extensive survey of current embodied-AI and robotics literature is deliberate. SRH is not an extrapolation from present robotic performance and does not infer future productive closure from contemporary engineering trends. Current systems can establish that particular operations are realizable, but their increasing capability cannot establish the future transitions proposed by SRH. The cited engineering literature is therefore used to define relevant classes of systems and constraints rather than as inductive evidence for recurrence.

This methodological restriction is symmetric: the absence of a transition at present does not establish its future impossibility, while present progress toward a transition does not establish its future realization. SRH therefore relies neither on an argument from current absence nor on an argument from technological trend.

References

- Anthropic Institute (Sept. 2026). *Scenarios for Our Economic Future: Economic Scenarios for Transformative AI*. Economic Scenario Explorer, version 1.0. URL: <https://www.anthropic.com/institute/econ-scenarios> (visited on 09/13/2026).
- Crick, Francis H. C. and Leslie E. Orgel (1973). “Directed Panspermia”. In: *Icarus* 19.3, pp. 341–346. DOI: 10.1016/0019-1035(73)90110-3.
- Freitas Robert A., Jr. and William P. Gilbreath (1982). “Replicating Systems Concepts: Self-Replicating Lunar Factory and Demonstration”. In: *Advanced Automation for Space Missions*. NASA Conference Publication 2255; NTRS Document ID 19830007081. National Aeronautics and Space Administration. URL: <https://ntrs.nasa.gov/citations/19830007081> (visited on 09/13/2026).
- Maturana, Humberto R. and Francisco J. Varela (1980). *Autopoiesis and Cognition: The Realization of the Living*. Vol. 42. Boston Studies in the Philosophy and History of Science. Dordrecht: D. Reidel Publishing Company. DOI: 10.1007/978-94-009-8947-4.
- Mix, Lucas J. (2026). “The Origin, Extension, and Future of the “NASA Definition” of Life”. In: *Astrobiology* 26.1, pp. 66–74. DOI: 10.1177/15311074251412317.
- NASA Astrobiology Program (July 25, 2013). *Forming a Definition for Life*. National Aeronautics and Space Administration. URL: <https://astrobiology.nasa.gov/news/forming-a-definition-for-life/> (visited on 09/13/2026).
- Solé, Ricard (2016a). “Synthetic Transitions: Towards a New Synthesis”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 371.1701, p. 20150438. DOI: 10.1098/rstb.2015.0438.
- (2016b). “The Major Synthetic Evolutionary Transitions”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 371.1701, p. 20160175. DOI: 10.1098/rstb.2016.0175.