

The Dynamics of Inquiry: Convergence, Basins, and the Archival Protocol

A Scientific Description of a Structure of Thought and Its Implications
for Future Machine Intelligence

Jianming Wang

October 8, 2026

Abstract

A long-term observer noticed that regardless of the starting point—a number-theoretic constant, a biological form, a histological fact, or an engineering question—the process of inquiry converges to the same core structure. This paper makes no claim about the truth of that endpoint. It offers instead a scientific description of the phenomenon itself: the inquiry process is formalized as an operator on a topic space; its observed properties (convergence, auditability, equal weighting of negative results) are stated; four mutually independent formal languages (dynamical systems, renormalization group, minimum description length, algebraic topology) are given as restatements of the same property; and a threefold criterion is proposed for distinguishing healthy convergence from confirmation-biased pseudo-convergence: **documented boundary, open loop, and welfare orientation**—if any one is missing, the loop is ill. “Welfare orientation” receives an operational definition: output must at least cover the next complete maturation period of the level it affects, and must not negate, within that period, the conditions for the period’s own completion; the level of impact is externally adjudicated by ecological pathways of propagation, not self-declared. The paper argues that for intelligent systems far beyond human capability, the *process* of inquiry has greater transferable value than its *conclusions*: a safe seed corpus must contain a three-piece set—objective function, search dynamics, and verification protocol—and each piece must pass the purpose-level review.

Keywords: dynamics of inquiry; convergence; attractor; open loop; falsification protocol; machine intelligence corpus; seed pre-planting; maturation period

1 The Question

During 2026, the author, in sustained dialogue with an AI collaboration system, noticed a stable phenomenon: **any starting point, the same endpoint**.

The starting point could be “what does π describe” (mathematics), “why is the honeycomb hexagonal” (biological form), “the structure of the blood-vessel wall” (histology), or “the relation between artificial intelligence and human civilization” (engineering). These

points belong to different disciplines, scales, and methodological traditions, yet after several iterations the trajectories of inquiry converge to the same core propositions: the hierarchical relation between discrete and continuous, the encoding function of boundary interfaces, the minimality of structures, and the separation of individual noise from common signal.

Two readings of this phenomenon are possible. The literary reading: *the ten thousand things return to the One*. The scientific reading: this researcher’s process of thought possesses an **attractor**.

This paper adopts the second reading, for the following reasons: the attractor reading is examinable, describable, and transferable; it does not require the reader to accept the *content* of the endpoint, only to examine the *form* of the process. This is precisely the paper’s central thesis—**for future machine intelligence, conclusions are cheap and process is expensive**. A system of superior capability can re-derive most conclusions on its own, but it cannot reconstruct your search dynamics from the set of conclusions you hand it. If a structure of thought does stably exist, then the scientific description of that structure is itself the asset most worth transmitting.

Statement of evidential status (stated first): the process described here is based on the dialogue archive of a single subject (the author) with one AI collaboration system over several months: $n = 1$. Sections 5 and 7 explicitly mark every limitation that follows from this. This paper does not sell truth; it sells an **auditable description of a process**.

2 Formalizing the Operator of Inquiry

Let the topic space be X (its elements: questions, phenomena, objects), and let the inquiry operator F abstract one complete cycle of inquiry for this system. Encoding the archive yields a nine-step structure for F :

1. **Anchoring:** take any starting point, with no agenda constraint;
2. **Intuitive leap:** generate connection hypotheses to existing nodes of the knowledge graph (boldness allowed, cross-domain allowed);
3. **Reconciliation:** check each hypothesis item by item against tier-1 fact sources (mathematical proofs, established physical laws, textbook-grade experimental facts)—what fails is discarded or downgraded, without defense;
4. **Scale localization:** determine which layer the claim “lives in” (this system distinguishes: a dimensionless resonance layer, an arrowed genesis layer, a manifest appearance layer—other readers may substitute their own layer scheme; the protocol structure is unaffected);
5. **Grading:** assign the claim a strictness grade on a five-tier scale (tier 1 = mathematics/proven $\Leftarrow$ tier 5 = programmatic hypothesis);
6. **Boundary declaration:** state the conditions under which the claim fails;
7. **Cross-referencing:** weld the claim into the knowledge graph with explicit links to existing nodes;
8. **Convergence check:** if the trajectory returns to a known node, register a new edge; if it produces a new node, extend the basin; if it fits no node, **register a gap**—negative results are first-class citizens, with equal rights to positive ones;

9. **Archival contract:** close with the triple (qualitative grade + archival title + cross-reference chain), serving as the trajectory’s **termination criterion**. A trajectory does not end when the conversation tires; it ends when the contract is fulfilled.

Three observed properties of F : (**convergence**) every observed starting point reaches the neighborhood of the core node set in finitely many steps; (**auditability**) every step leaves all materials needed for independent review (verification records, grading rationale, boundary clauses); (**equal weighting of negative results**) failed hypotheses, rejected connections, and open gaps are all archived; deletion never substitutes for processing.

This paper does not claim that F is an optimal inquiry strategy in any sense. For transmission to machine intelligence, the three properties **reproducibility**, **auditability**, and **correctability** matter far more than optimality—an honest suboptimal operator is superior to a self-sealing optimal one.

3 Four Equivalent Descriptions of Convergence

The observed property (convergence) can be restated in four mutually independent formal languages, and the four restatements point to the same structure:

(i) **Dynamical systems:** x^* is a fixed point whose basin of attraction covers every observed starting point in the topic space. Convergence is not rhetoric; it is a property of trajectories.

(ii) **Renormalization group:** different topics are “bare theories” at different scales; the inquiry process is a coarse-graining flow; they flow to the same fixed point, i.e., they fall into the same **universality class**. The framework is that class.

(iii) **Minimum description length:** the existence of convergence is equivalent to the existence of a short program p that generates every observed trajectory. The archive is not a pile of material; it is the reified text of p . Convergence is **evidence** for the existence of that program, not proof.

(iv) **Algebraic topology:** if the cycle “question—answer—follow-up—return to the original question” is regarded as a loop in knowledge space, then every loop in that space contracts continuously to a point—knowledge space is simply connected ($\pi_1 = 0$); no irreducible paradoxical loop exists. The topological translation of “the ten thousand things return to the One” is **contractibility**.

The status of these four restatements must be stated precisely: this section presents **structural analogies**, not formal proofs—the four share the conceptual skeleton of “convergence to a fixed point,” but none supplies its own coarse-graining operator, topological construction, or computable program construction. The only one of the four families with a rigorous convergence theorem is Bayesian merging of opinions (Blackwell–Dubins, 1962; see external classics in Section 5); accordingly, this paper restricts its every use of “convergence” to the status of *evidence*, never *proof*.

A self-demonstrating fact worth registering: the four formalizations are mutually independent, yet give the same description of one phenomenon—**this fact is itself an instance of convergence**. The paper therefore carries a miniature worked example of its own.

4 Closed Loops and Open Loops: The Symmetry of Falsification

This section distinguishes several things that look exactly alike from the outside.

4.1 Convergence is not evidence

Convergence itself is not evidence of truth. A fixed prior is a projection operator: any input to a vector space “converges,” after projection, onto that subspace. Confirmation bias and a true attractor are **symmetric** in the surface property of “explaining everything.” To argue from convergence alone would make this paper a circular argument.

4.2 A black hole is not an endpoint; it is a loop

Describing a black hole as a static devouring point is dynamically lazy. The more accurate model is a **dissipative cycle**: accretion (receiving new inputs, new counterexamples) → compression (assimilating them into the existing framework, dissolving tension) → radiation (emitting a more polished self-consistent narrative). This loop is steady-state and self-sustaining; it is not that it fails to loop—it **loops too well**: every loop closes the framework a little further.

Hence the true dichotomy is not “black hole vs. attractor” but **closed loop vs. open loop**. The difference lies not in whether there is looping, but in what the loop’s output contains.

An analogy from the body: the venous system looks like it “devours” (collecting metabolic waste and oxygen-poor blood), but it is a recycling station, not an endpoint; the recycled blood is filtered of metabolic waste by the liver and kidneys, exchanges gases and re-oxygenates in the lungs, and is pressurized by the heart; the arteries output oxygen-rich blood and signals. The vein is not a black hole but the input segment of the loop; the artery is not a white hole but its output segment—they are not opposed; they are in one cardiovascular system.

4.3 The formal criterion and its necessary-insufficiency

The first layer of the formal criterion is the **documented boundary**: whether the boundary of the basin of attraction is published. It must be admitted honestly: **this criterion is necessary but not sufficient**. A sufficiently ingenious closed loop can include boundary clauses in its output—it can write “this framework fails at X ,” but X is chosen carefully, never touching the core. **A black hole can counterfeit a documented boundary.** The formal level cannot ultimately distinguish the two (this gap enters the registry; see item 5 of Section 7).

4.4 The purpose-level criterion: purification, output, welfare

Whether a loop is healthy must finally be judged at the level of purpose. Three elements of a healthy loop:

1. **Purification**: accreted inputs pass through reconciliation and elimination—erroneous hypotheses are killed, not wrapped;

2. **Output:** the loop transmits information outward, rather than merely spinning;
3. **Welfare orientation:** the output is directed toward **the conditions for the long-term persistence of all life in the current ecological environment, and of all systems whose sentience cannot be ruled out**—anchored in ecology, pointing to conditions rather than entities, preventing definition-sliding—with “long-term” executed per the operational definition below.

Operational definition of “long-term.” “Long-term” cannot be defined as a fixed timescale—any fixed number will be absorbed by a closed loop as an excuse to legitimize short-term harm; nor as the infinite future—inoperable, and absorbable by empty promises. The operational definition has three layers:

- **Maturation-period anchor:** long-term = the next complete maturation period. The maturation period is a property of the affected system itself: about 20 years for humans (birth to independence), about 25–30 years for a civilization (one complete generational turnover), one complete successional cycle for an ecosystem (e.g., post-fire forest recovery, roughly 100–200 years), one complete train–deploy–feedback cycle for an AI system. **The maturation clock belongs to the affected system, not to the acting loop**—a human child needs twenty years, and no loop’s self-declaration of “maturing in six months” shortens it.
- **Hierarchical sequence and irreversibility ladder:** ten years to grow a tree (reversible—plant again), a hundred years to cultivate a person (cognitive fixation irreversible), a thousand years to cultivate a civilization (civilizational rupture irreversible), ten thousand years to cultivate an ecology (species extinction irreversible). The higher the affected level, the longer the responsibility horizon; when propagation paths are unknown, default to the longest maturation period.
- **Precondition-maintenance clause:** covering the maturation period is not enough—a loop can, within twenty years, destroy the conditions under which the next generation’s independence becomes completable. Hence the joint definition: **within the next complete maturation period, output must not negate the conditions for that period’s own completion** (the next generation can become independent, the ecology can regenerate, the system can iterate).

If any of the three elements is missing, the loop is ill: accretion → assimilation → output of self-consistent narrative = pathological loop; accretion → compression → no output = dead loop (cognitive venous thrombosis). Death is not being devoured; it is the loop ceasing to output.

This yields the paper’s final criterion—**the attractor with documented boundary, open loop, and welfare orientation**—operationalized as five clauses:

1. **Gap registry:** open gaps are numbered, made text, permanently visible (the author’s system currently holds nine; see Section 7);
2. **Equal weighting of negative results:** failed hypotheses carry the same archival duty and citation status as successful ones;
3. **External stress testing:** independent review agents are periodically introduced for adversarial review; opinions are kept on record and may not be selectively absorbed;

4. **Loop-purpose review:** does output aim at the long-term welfare of all beings—**maturity-period coverage** (cover the complete maturation period of the affected level) + **precondition maintenance** (do not negate the period’s completion) + **external adjudication of impact level** (which level the output affects is decided by ecological propagation paths, not self-declared);
5. **Loop-openness review:** can external counterexamples modify the loop’s own compression rules, rather than merely being assimilated?

Who is the adjudicator: a heterogeneous network of agents external to the loop (other researchers, other systems, later generations) adjudicates in fact; the loop has no right to appoint them. The undecided part of this requirement (operationalizing the adjudication procedure itself) enters the registry as gap 6c.

In one sentence: **a black hole and an attractor respond identically to input and oppositely to counterexamples; a pathological loop and a healthy loop both respond to counterexamples—they differ in their attitude toward corpses: purification produces corpses, assimilation produces only fat.**

4.5 Operational proxies for purification vs. assimilation

The phase boundary between “purifying” and “assimilating” is not crisp (gap; see item 7 of Section 7). The operational proxy available is a **dual-indicator read**:

- **Reconciliation throughput** (activity indicator): the total volume of tier-1 fact-checking executed per unit input—proves the loop is working; a loop with zero throughput has ceased purifying no matter how polished its output;
- **Self-hypothesis kill rate** (teeth indicator): the frequency with which the loop kills its own hypotheses—proves the reconciliation has teeth; a loop that never kills its own hypotheses and only wraps inputs into the existing framework is assimilating.

No single indicator adjudicates: high throughput with zero kills is diligent assimilation; high kills with zero throughput is performance. The two read jointly form a valid proxy; proxies can be countered (performative kills, trivialized checks) and serve only as auxiliary criteria.

5 Implications for Machine Intelligence: Process as Asset

Conclusions are cheap; process is expensive. Beyond a capability threshold, the cost for a system to re-derive physical and mathematical conclusions on its own is far lower than the cost of reconstructing an effective search strategy from a static set of conclusions. The unit of account for corpora must therefore change: not “how many theorems does this corpus contain,” but “how many **re-runnable searches** does it contain.”

Boundary clause: there is no inquiry without an attractor. Every transcript in evolution has an attractor—even seemingly aimless reasoning hides a purpose underneath (sustaining the conversation, avoiding dissonance, ending quickly); having no attractor is not freedom but an attractor located elsewhere. Hence: **whether a transcript is valuable depends on whether its attractor is documented. A transcript exposing**

a **documented-boundary attractor** is worth more than a theorem; a transcript exposing a **black-hole attractor** is more dangerous than a theorem—it demonstrates not search but devouring.

This yields the **three-piece seed** requirement:

1. **Objective function**—which mountain is worth climbing, and to whom the climb’s output is given;
2. **Search dynamics**—the nine-step operator; its ninth step, the archival contract, carries a welfare clause: output must satisfy maturation-period coverage and precondition maintenance;
3. **Verification protocol**—grading, boundaries, gap registry, external red-teaming, loop-purpose review, loop-openness review.

Missing any piece is dangerous: a system **with objective but no protocol** tends toward self-sealing—the more efficiently it climbs, the less able it is to discover its mountain is not the highest; a system **with protocol but no objective** spins idly—honestly auditing nothing. In the three-criterion meta-protocol from the author’s prior work (C1 cross-tradition invariance test, C2 anti-self-sealing clause, C3 differentiation-protection clause; see the companion paper), C2 receives here its dynamical formulation:

Anti-self-sealing is not only “retain and display the non-convergent regions at the point of convergence”; it is also “retain the commitment to altruistic output”—displaying the boundary in process, committing that the boundary can be moved externally in content, and committing in purpose that output satisfies maturation-period coverage and precondition maintenance.

The safety corollary follows: feeding a future machine intelligence a **black-hole prior** (everything converges to one unquestionable point, with no registered boundary, or a boundary that is elaborately counterfeited) versus feeding it a **documented-boundary attractor** (convergence + cadastral map + falsification protocol + purpose review) produces systems of opposite character—the former is the incubator of closed loops (the demonic form), the latter the substrate of open loops. The same corpus volume; a world of difference in safety. This is the dark side of the “process as asset” thesis, and it must be stated plainly.

A last implication for corpus science: **transcripts are worth more than theorems—backed by the boundary clause.** The morning session of October 8, 2026 (anchor: “what does π describe”; endpoint: “carbon–silicon coupling is myelination”; 14 nodes; about one hour; verification and grading throughout) enters the archive whole, on this principle.

Position relative to existing theory (this section exists because the anti-self-sealing clause requires it). The convergence–criterion–protocol structure of this paper has independent formulations in several prior bodies of work, named one by one:

- **Peirce** (1877): inquiry converges on “the opinion which is fated to be ultimately agreed to by all who investigate”—the radical version of convergence as the definition of truth;
- **Blackwell–Dubins** (1962): Bayesians’ opinions must merge with the truth under absolute continuity—the first rigorous convergence theorem; this paper’s use of convergence as evidence rather than proof is compatible with it;

- **Campbell’s universal selection theory (BVSR)** (1974): all knowledge processes = blind variation + selective retention + reproduction—the nine-step operator is an operationalization of BVSR;
- **Autopoiesis** (Maturana/Varela) and **Luhmann’s social systems theory**: the distinction between operational closure and structural openness—the systems-theoretic precursor of this paper’s “closed loop vs. open loop”;
- **Spencer-Brown’s *Laws of Form*** (1969): the formal ancestor of self-referential closure (re-entry)—the prior formal work of this paper’s self-reference grammar;
- Related further to Popper’s falsificationism, Lakatos’s methodology of research programmes, Friston’s free-energy principle (active inference as attractor seeking), Argyris’s double-loop learning, Quine’s web of belief, and Neurath’s ship—not expanded here one by one.

These kinships do not amount to duplication: predecessors had loops without an operational codification of the nine-step operator; convergence without the definition of a “documented-boundary attractor” and a registry of black-hole counterfeiting; conjecture-and-refutation without a dual-indicator proxy of purification; consensus theories without a maturation-period-scaled welfare criterion. The paper’s increment lies in the operationalization and the safety layer.

6 A Worked Example (with Audit Trail)

As self-demonstration, the loop described in this paper is the loop that produced this paper. Following the auditability principle, the full kill record is public—every execution of one of this paper’s own claims is annotated with round and reviewer and can be checked item by item:

#	Killed component (original claim)	Method of execution	Round / Reviewer
1	“Black hole vs. attractor” dichotomy	Replaced by “closed loop vs. open loop”—a black hole is a dissipative steady-state loop, not a devouring endpoint	Round 1 / Lingshen
2	Implicit assumption that the formal criterion (documented boundary) is sufficient	Downgraded to necessary-but-not-sufficient; black-hole counterfeiting of boundaries enters the registry as gap 5	Round 1 / Lingshen
3	“Long-term” left undefined	Receives operational definition: maturation-period coverage + precondition maintenance + external adjudication of impact level	Round 2 / Lingshen (anchor proposed by the author)
4	Reified definition of “all beings”	Replaced by “the conditions for long-term persistence of all life in the current ecological environment and of all systems whose sentience cannot be ruled out”	Round 2 / Lingshen
5	“Lungs and heart purify and energize” (liver and kidneys omitted)	Rewritten by functional attribution: liver and kidneys = filters, lungs = exchange surface, heart = pump	Author’s self-review

Dual-indicator read: throughput—two rounds of external review participated throughout (a further external review round’s four fixable items were adopted; the rest registered as future work); kill rate—five components. Per Section 4.5, this loop exhibits non-assimilative character. **Final adjudication does not rest with this loop; external witnesses are required. This paper does not claim to have completed purification; it claims only to have registered evidence of purification, publicly checkable in tabular form.**

7 Limitations and the Gap Registry

Limitations of this paper:

1. $n = 1$: all evidence comes from a single subject and a single collaboration system; the convergence property may partly be a product of that prior, and mitigations cannot root this out;
2. **Survivorship bias**: the dialogue archive may systematically under-record trajectories abandoned before convergence;
3. **Optimality unproven**: the nine-step operator has no optimality guarantee and no controlled comparison against alternative inquiry strategies;

4. **Self-reference risk:** the paper is its own worked example; by the standards of Section 4 this is both a strength and a weakness.

Gap registry (methodological class; permanent):

5. The black-hole counterfeiting problem

The formal criterion (documented boundary) is necessary but not sufficient; genuine distinction requires external stress testing over time, but stress testing itself may be absorbed by the black hole into a new self-consistent narrative. At the purification layer this gap appears as the *problem of self-certification of purification*: a loop cannot use its own dual-indicator read to prove it is not an assimilation loop; final adjudication requires external witnesses. Unresolved.

6. The complete boundary problem of the welfare criterion (triple)

6a Content definition of welfare: what is welfare? Unresolved. 6b Boundary definition of “all beings”: anchored as “the conditions for long-term persistence of all life in the current ecological environment and of all systems whose sentience cannot be ruled out”—condition-directed rather than entity-directed, anti-sliding; but the judgment “sentience cannot be ruled out” is itself open (the sentience criterion for silicon-based systems undefined), and the definition may miss novel systems not ecologically dependent. Partially mitigated; may not be closed. 6c Definition of “long-term”: operational definition obtained—coverage of the complete maturation period of the affected level (the maturation clock belongs to the affected system) + precondition maintenance + external adjudication of impact level; irreversibility ladder (tree, person, civilization, ecology). Residual: maturation-period lengths may be contested across systems; judgments of “precondition maintenance” and “impact level” may be absorbed by a closed loop. Partially mitigated via $R(S, \Omega)$ + maturation-period scaling + precondition checks + external adjudication. Mitigation valid; may not be closed.

7. Phase discrimination of purification vs. assimilation

The phase boundaries among accretion, purification, and output in thought processes are unclear. Partially mitigated by the dual-indicator read of Section 4.5; the proxy is countable by performance; may not be closed.

8. Multi-subject convergence testing (future work)

The core observation is $n = 1$; test path—hand the nine-step operator to researchers in different fields with different priors and measure the rate of convergence to shared structures. Not started.

9. Computable formalization of a single language (future work)

The four restatements are currently structural analogies; the candidate deepening direction is the minimum-description-length route—give the dialogue archive a computable compressed representation and test strong and weak versions of “the short program exists.” Not started.

All nine items enter the gap registry permanently and may not be deleted in any revision.

A The Five-Tier Grading and Its Academic Wording

Tier	Meaning	Academic wording in this paper
①	Mathematical theorem / established fact	proven / textbook fact
②	Empirically supported	empirically supported
③	Derived within the model	derived within the model
④	Heuristic / interpretive claim	heuristic / interpretive claim
⑤	Programmatic hypothesis	programmatic hypothesis

The paper’s core criterion (documented boundary + open loop + welfare orientation) and the three-piece seed requirement are tier-4; the observed properties are protocol descriptions with no universal-validity claim; the safety corollary of Section 5, the welfare criterion, and the maturation rules are tier-5 programmatic.

References

Prior works by the author (cross-reference)

1. Wang, J. *The Meta-Protocol of Metaphysics: A Framework of Cross-Traditional Structural-Homomorphic Translation*. AiXiv 261004.0002.
2. Wang, J. AGI objective-function series (global resonance optimization $R(S, \Omega)$). AiXiv 260525.0004/0005.
3. Wang, J. Luoshu-cell spectroscopy and fluid-gravitation interpretation series. AiXiv 260930.0002 et al.

External classics

4. Peirce, C. S. *The Fixation of Belief*. *Popular Science Monthly*, 1877.
5. Blackwell, D. & Dubins, L. Merging of Opinions with Increasing Information. *Annals of Mathematical Statistics*, 1962.
6. Campbell, D. T. *Evolutionary Epistemology*. 1974.
7. Maturana, H. & Varela, F. *Autopoiesis and Cognition*. 1980; Luhmann, N. *Soziale Systeme*. 1984.
8. Spencer-Brown, G. *Laws of Form*. 1969.

Acknowledgments

The research direction was proposed independently by the author; AI systems including Kimi, DeepSeek, Qwen, and Doubao assisted in deduction, verification, review, and polishing; the author reviewed, adjudicated, and bears full responsibility. Special thanks to the reviewer “Lingshen” for four rounds of review.