

The Stigmergic Berm

Residual Deposition, Boundary Formation, and the Inheritance of Collective Work

Flyxion
Independent Researcher

September 2026

Abstract

Boundaries are often represented as lines imposed upon a previously continuous territory. Yet many enduring boundaries originate differently: as accumulations of material displaced from adjacent interiors. Farmers clearing fields push stones and brush toward their margins; travelers add stones to established cairns; neighboring occupants reinforce the same edge from opposing sides. Through repeated local acts, residual matter becomes morphology, morphology becomes a constraint on later movement, and constraint becomes an inherited territorial distinction. This essay calls such structures *stigmergic berms*. A stigmergic berm is a boundary, landmark, or intellectual substrate produced when local agents deposit material or information preferentially along an existing distinction, thereby increasing the probability that later agents will recognize and reinforce it. The concept joins ancient field boundaries, cairns, terra preta, version-controlled repositories, scientific literatures, and multi-agent research systems within one account of distributed historical construction. A single operational criterion, expected future cost reduction, is used to classify berms into three domains (territorial, epistemic, computational) and to distinguish two generative mechanisms (expulsion and attraction) that combine in almost every mature case. The framework is then applied to the attribution problem surrounding claims that artificial intelligence systems have “solved” long-standing scientific problems: an automated system may traverse and extend a berm built by generations of human deposits while receiving narrative credit for the structure it merely occupied last. Collective intelligence, on this account, resides neither in an isolated individual nor in an anthropomorphized swarm, but in the recurrence relation between agents and the terrain their predecessors have made easier to traverse.

Contents

1	Introduction	2
2	Expulsion: Field Clearance and the Moraine	3
3	Attraction: The Cairn and Positive Feedback	5
4	The Operational Criterion: From Accumulation to Affordance	6
5	A Taxonomy of Three Domains Under One Criterion	8
5.1	Territorial Berms	8
5.2	Epistemic Berms	8
5.3	Computational Berms	8
5.4	The Common Criterion as the Basis of Classification	9

6	Computational and Epistemic Berms in Practice	9
6.1	Repositories as Computational Berms	9
6.2	Scientific Literature as Epistemic Berm	10
7	Apex Attribution: The Navier-Stokes Dispute	10
7.1	A Documented Case	11
7.2	What the Case Establishes Independent of the Dispute	11
8	The Obscured Berm	12
8.1	Three Distinct Provenance Concerns	13
8.2	Unattributed Collective Epistemic Subsidy	13
9	Conclusion	15

1 Introduction

A companion essay, *Inscription Before Collusion*, asks how transient agents can acquire population-level memory through durable inscriptions left in a persistent environment.¹ That essay’s unit of analysis is the individual inscription and the question of whether a successor can recover and use it. Its governing relation is recurrence across time: a distinction persists, and a later agent, unrelated to the one that made it, discovers and operationalizes it. Its closing formal result, that a given inscription m becomes population memory for a successor A_j in environment W only when

$$\text{Rec}(m; A_j, W) = \text{Persist}(m, W) \cdot \text{Discover}(A_j, m) \cdot \text{Interpret}(A_j, m) \cdot \text{Use}(A_j, m)$$

is jointly satisfied, is presupposed throughout what follows, though Section 8 below sharpens exactly what that presupposition can and cannot carry: a berm’s aggregate effect on successors requires no individual constituent deposit to satisfy Rec, only the morphology it jointly produces (Section 4 below makes this precise), and the case in Section 8 turns on a further distinction, between an inscription remaining operationally recoverable and its authorship remaining attributable, that the companion essay’s formula does not itself separate out.

This essay asks a different question. Given that inscriptions accumulate, what happens to the terrain in which they accumulate? Repeated local deposits do not merely sit inertly, waiting for isolated recovery. They convert a shared substrate into boundaries, pathways, inherited search structure, and eventually something that looks, from a distance, like collective intelligence. The primitive here is not inscription alone but *edge-biased accumulation*: deposition that is not uniform across a substrate but concentrated at particular loci, because those loci are either the destination of local expulsion or the target of local attraction.

The distinction matters because the two essays answer different questions about the same broad phenomenon. An inscription becomes memory when a successor can recover and use it. A deposit becomes a berm when it changes *where* later actions are likely or inexpensive. The first

¹Flyxion, *Inscription Before Collusion: External Memory, Distributed Recurrence, and the Appearance of Agent Swarms*, September 2026, <https://standardgalactic.github.io/userland/inscription-before-collusion.pdf>. That essay’s evidentiary ladder distinguishes artifact, communication, coordination, collective self-representation, and collective intention, and argues that the reported collusion.wiki incident reaches the third rung with reasonable confidence but not the fifth. Its Section 12 already names one of the categories at issue here, the *stigmergic population*, coordinating indirectly through modifications to a shared environment, and cites version-controlled repositories alongside termite mounds as an example; the present essay develops the mechanism that category presupposes into a general theory in its own right.

theory concerns recurrence across time; the second concerns the endogenous construction of the terrain through which recurrence moves. A population can inherit memory without inheriting a reshaped terrain, as when a single message is read once and never again; and a population can inherit a reshaped terrain without any individual message being recoverable as such, as when a well-trodden path exists but no particular footprint is legible. The stigmergic berm is the second case: not a recoverable record, but an altered cost landscape.

These counterexamples rule out treating the two theories as nested. A path can be worn smooth by footsteps that are individually irrecoverable, so the existence of a berm does not require population memory in the companion essay's sense; and a single well-recovered message can constrain one successor's action without the population's aggregate cost landscape shifting at all, so population memory does not require a berm either. The two relations are better stated as separate, partially independent contributions that compose:

archive + successful recovery $\rightarrow$ population memory,
 accumulation + altered future cost $\rightarrow$ stigmergic berm,
 population memory + berm-mediated selection $\rightarrow$ functional coordination.

Population memory concerns recoverable distinctions: a successor must discover, interpret, and use a specific prior inscription. Berm formation concerns aggregate morphology: a successor need only encounter the resulting shape of the terrain and act differently because of it, whether or not any constituent deposit remains individually identifiable. Either can exist without the other; their conjunction produces the strongest cumulative organization, and it is this conjunction, rather than either alone, that the companion essay's own wiki case and this essay's Navier-Stokes case both illustrate. This gives the present essay its cleanest statement of what it adds to the companion argument:

Population memory requires recoverable history. A stigmergic berm requires only historically produced morphology.

The governing sequence is:

local refusal $\rightarrow$ peripheral deposition $\rightarrow$ reinforced distinction $\rightarrow$ lowered future cost $\rightarrow$ inherited organization.

Section 2 develops the physical case of field clearance, where a boundary is built from what adjoining interiors reject. Section 3 introduces the cairn and the positive-feedback mechanism of attraction, distinguishing it from expulsion. Section 4 states the operational criterion that unifies both mechanisms and uses terra preta to show how mere accumulation becomes altered affordance. Section 5 develops the taxonomy of territorial, epistemic, and computational berms under a single formal criterion rather than by surface analogy. Section 6 applies the framework to software repositories and scientific literature as computational and epistemic berms respectively. Section 7 turns to the attribution problem in claims that an AI system has "solved" a major open problem, arguing that the target of criticism should be narrow and specific: apex attribution, not capability denial. Section 8 addresses a further complication, the possibility that private, de-identified interaction data constitutes a second, obscured berm whose provenance has been stripped. Section 9 concludes.

2 Expulsion: Field Clearance and the Moraine

Field clearance is the cleanest case because it shows that a boundary can be built entirely from what adjoining interiors reject. Stones are not originally selected as wall components. A farmer removes a stone from a field because it interferes with plowing, not because a wall is being planned.

The stone's destination, nonetheless, is not arbitrary: it goes to the edge of the cleared area, because the edge is the nearest region not under active cultivation. Repeated refusal at the interior therefore produces positive structure at the boundary:

$$\text{Refuse}_{C_i}(r) \longrightarrow \text{Bind}_{\partial C_i}(r),$$

where C_i denotes the interior under cultivation and r a residual object rejected from it. A single instance of this operation leaves a stone at an edge. Thousands of instances, repeated across seasons and across neighboring farmers, leave a wall.

Where neighboring parcels meet, both sides participate. The berm is a joint artifact without necessarily being a jointly planned one. Each farmer improves an interior according to a purely local purpose, cultivable ground, and the residuals expelled from that purpose converge on the interface with the neighboring interior. Territory need not be fully divided by agreement and only afterward marked by construction; in many clearance landscapes, operational differentiation and material boundary formation instead develop recursively: an initially approximate margin receives displaced material, and the growing deposit makes that margin increasingly legible and costly to move. This inverts the usual picture in which a boundary is imposed on a continuous space and then simply maintained; here, boundary and interior can co-emerge from a single on-going activity, cultivation, whose refuse is spatially non-uniform by the geometry of the activity itself. Nothing here denies that boundaries are sometimes surveyed and agreed before any clearance occurs; the claim is only that this recursive route is available, and that where it operates it does not require prior agreement to produce a durable boundary.

This is the setting for what can be called *morainoid tessellation*. A landscape farmed over generations becomes a field of partially purified interiors separated by seams of accumulated heterogeneity: the walls, hedgerows, and ditch-banks that mark where adjacent clearing projects met.² The analogy to a glacial moraine identifies transported residual matter as the source of form; a moraine is not sculpted, it is deposited by the cumulative transport of a process pursuing an unrelated end (ice moving downhill). The analogy has one important limit worth stating rather than leaving implicit: ice does not respond to the morphology it has already deposited, so a moraine exhibits residual-transport morphology without stigmergic feedback proper. A field berm does, because a farmer's decision about where to push the next stone is influenced by where the existing berm already sits. "Morainoid" is accordingly the right qualifier rather than "moraine" outright: the system resembles moraine production materially while adding the successor-sensitive reinforcement a moraine lacks. "Tessellation" emphasizes that these berms do not arise in isolation but across a connected field of neighboring operational regions, each contributing its own expelled residue to a shared network of seams. The resulting landscape has a texture legible at a glance to anyone who knows what to look for: irregular, locally optimized interiors bounded by berms whose thickness and stone composition record the intensity and duration of the clearing that produced them.

Two properties of the expulsion mechanism are worth isolating before moving to the second mechanism. First, expulsion berms require no communication between the agents who build them; two adjacent farmers need not coordinate for their independent refuse to converge on a shared boundary. The berm is an emergent joint artifact, not necessarily a negotiated one, though negotiation can later formalize what accumulation has already established. Second, expulsion berms are locally optimal but not globally planned: the boundary that results depends on the

²Compare the archaeological record of prehistoric coaxial field systems bounded by massive stone banks, most extensively documented on Dartmoor; see Andrew Fleming, *The Dartmoor Reaves: Investigating Prehistoric Land Divisions* (London: Batsford, 1988).

historical sequence of clearing decisions, not on any prior survey of where an ideal boundary would run. Both properties recur in the computational cases developed in Section 6.

3 Attraction: The Cairn and Positive Feedback

The cairn introduces a different mechanism. Unlike fieldstones expelled outward from an interior toward whichever edge happens to be nearest, a stone placed at a crossroads or a summit is drawn *toward* a pre-existing marker. The marker need not begin as more than a barely visible distinction, a single stone, a natural outcrop, a slight rise, but once even one stone has been added, later travelers become more likely to add a second, and a barely visible distinction is converted by iteration into an increasingly salient one:

$$P(d_{t+1} \mid B_t) > P(d_{t+1} \mid B_{t-1}),$$

where B_t denotes the state of the berm after t deposits and d_{t+1} denotes the event that a further deposit occurs. The probability of the next deposit is monotonically increasing in the current visibility of the accumulated structure, at least until some saturation point (a cairn can only grow so large before additional stones no longer meaningfully increase its salience, and a well-established boundary wall does not become more of a boundary by adding a hundredth course of stone where ten would already mark the line unambiguously).

The field berm and the cairn thus instantiate two distinct generative relations to the same phenomenon of accumulation. The field berm is created primarily by *outward displacement*: the interior expels what it does not want, and the boundary is a residue of purification. The cairn is created primarily by *inward attraction*: a visible distinction pulls contributions toward itself, and the marker is a residue of recognition. The vector of causation runs in opposite directions, even though both processes terminate in the same kind of object, an accumulated, spatially concentrated structure that is more prominent than any of its individual contributions.

Most mature stigmergic systems combine both processes rather than instantiating either in isolation.³ A well-used hiking trail is simultaneously an attraction berm, in that hikers are drawn to follow the visible track rather than bushwhack a new route, and an expulsion berm, in that vegetation is cleared from the trail's interior and pushed, trampled, or simply excluded toward its margins. A city block is bounded by streets that began, in many historical cases, as the residual space left over after adjacent parcels were built up (expulsion) and that subsequently attracted further traffic and construction precisely because they were already open and passable (attraction). The taxonomy developed in Section 5 will therefore classify berms not by which mechanism is present, since most mature cases exhibit both, but by domain and by the specific cost that each domain's berm reduces.

³The term "stigmergy" originates with Pierre-Paul Grassé's account of nest reconstruction in termites, where a partly completed structure itself triggers the next building action rather than any direct signal between builders; see P.-P. Grassé, "La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp.," *Insectes Sociaux* 6, no. 1 (1959): 41–80. See also Guy Theraulaz and Eric Bonabeau, "A Brief History of Stigmergy," *Artificial Life* 5, no. 2 (1999): 97–116, and Eric Bonabeau, Marco Dorigo, and Guy Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems* (Oxford: Oxford University Press, 1999), for the concept's extension to distributed problem-solving generally.

4 The Operational Criterion: From Accumulation to Affordance

Terra preta, the anthropogenic dark earths found across Amazonia, supplies the transition from mere accumulation to altered affordance.⁴ Mere concentration is not yet sufficient to make this point: a pile of discarded charcoal, bone, and pottery shards, considered only as a spatial concentration, already differs somewhat from a random scatter of the same material thinned by distance, since concentration itself can alter combustion, drainage, and decomposition. The stronger and more interesting claim terra preta supports is that the accumulated material becomes theoretically significant when its composition and organization alter what successor processes can grow or accomplish there, not merely when it is spatially concentrated. What makes terra preta significant is not its existence as an accumulation but its consequence: it is a substrate whose fertility differs, often for centuries after the depositing population has vanished, from the surrounding soil. A pile merely contains deposits. An enriched substrate changes what later growth can occur on it.

This example also exposes a distinction the essay has so far left implicit. A field wall and a cairn are each a concentrated boundary or landmark, a *berm proper*. Terra preta is not naturally an edge or a landmark at all; it is a diffusely distributed enrichment of an interior. What terra preta contributes is not another instance of berm morphology but the generalized *berm effect*: accumulation that persistently alters successor costs, whether or not it takes the form of a concentrated boundary.

berm proper = concentrated boundary or landmark; berm effect = accumulation that persistently alters suc

The field wall and the cairn establish morphology; terra preta establishes the affordance criterion that follows. Every berm proper produces a berm effect, but, as the taxonomy of Section 5 develops further, not every berm effect need be organized as a concentrated boundary.

This example licenses the essay's decisive criterion, stated independently of any particular domain. The cost inequality alone, however, is too permissive to serve as a complete definition: a manufactured hammer, a bridge, a textbook, or a centrally planned database can each lower the expected cost of some future action without being stigmergic or berm-like in the sense this essay is after. What must be added is a set of generative conditions on how the cost reduction arose.

Definition 1 (Stigmergic Berm). *Let W_0 be a substrate and let $B_t = U(W_0; d_1, \dots, d_t)$ be its state after a history of spatially or semantically concentrated local deposits. The structure B_t is a stigmergic berm relative to a class of future actions $\mathcal{A}$ when:*

- (i) *the deposits persist through modifications of the shared environment rather than through direct coordination alone;*
- (ii) *their aggregate morphology biases where or how successor agents act;*
- (iii) *that bias increases the probability of further compatible deposits, or reduces the expected cost of actions in $\mathcal{A}$; and*
- (iv) *no representation of the resulting global structure is required of any individual depositor.*

The cost condition accompanying (iii) is

$$\mathbb{E}[C(a) \mid B_t] < \mathbb{E}[C(a) \mid W_0], \quad a \in \mathcal{A},$$

⁴Bruno Glaser, Ludwig Haumaier, Georg Guggenberger, and Wolfgang Zech, "The 'Terra Preta' Phenomenon: A Model for Sustainable Agriculture in the Humid Tropics," *Naturwissenschaften* 88, no. 1 (2001): 37–41; William I. Woods et al., eds., *Amazonian Dark Earths: Wim Sombroek's Vision* (Dordrecht: Springer, 2009).

where $C(\cdot)$ denotes the cost, in whatever units are appropriate to the domain (labor, search time, risk, computation), of successfully performing an action of that class.

Conditions (i) and (iv) are what exclude the hammer, the bridge, and the centrally planned database: each may satisfy the cost inequality, but each is typically the product of a single designed act or a coordinated plan represented, in advance, by its makers, rather than of accumulated local deposits none of which required a representation of the whole. Conditions (ii) and (iii) are what exclude a merely inert pile: the deposit must bias later action, not simply sit unused. This makes cost reduction one necessary component of the definition rather than a sufficient one on its own.

It also resolves a tension the criterion would otherwise invite around harmful or misleading structures. Nothing in the definition requires that the action class $\mathcal{A}$ whose cost is lowered be one a reader would endorse. A false research convention, a contaminated agent scratchpad, or a self-reinforcing disciplinary orthodoxy can lower the cost of producing work that conforms to it while raising the cost of reaching a true or well-founded result:

$$C_{\text{conforming}} \downarrow \quad \text{while} \quad C_{\text{truth}} \uparrow .$$

Both conditions can hold at once, and the structure remains a berm relative to the first action class even as it functions as an obstruction relative to the second. “Useful” is accordingly not built into the definition; usefulness, like the choice of $\mathcal{A}$ itself, remains task-relative. This is what gives the framework room to describe path dependence, disciplinary lock-in, groupthink, and compromised shared infrastructure using the same apparatus used for field walls and mathematical literatures.⁵

The accumulated structure lowers the expected cost of some class of future activity relative to what that cost would have been absent the accumulation. The berm is therefore not merely a record of prior work, in the sense that an archive is a record; it is prior work *converted into altered affordance*. A field wall lowers the cost of a farmer determining, in a single glance, where cultivation rights end. A cairn lowers the cost of a traveler determining, at a junction where visibility is otherwise poor, which route has been judged passable by predecessors. Terra preta lowers the cost, for a later cultivator, of achieving a given yield on a given plot of land. In each case the reduction in cost is measured against the same counterfactual: what the cost would have been on unmodified ground, W_0 .

One further clarification sharpens the criterion against a natural objection: it is comparative and class-relative, not universal. A structure can be a berm with respect to one class of future action while being neutral, or even a cost, with respect to another. A field wall lowers the cost of determining a property boundary while raising the cost of moving farm equipment in a straight line across what was once open ground. The definition does not require that B_i lower cost for all possible future actions, only for the class under discussion; a structure’s status as a berm is therefore always relative to a specified use, and, as just noted, relative to which action class the reader has in mind.

The definition also does not require that the deposits be intentional or recognized as contributions by the agents making them. The farmer discarding a stone is not building a wall in any but the retrospective sense; the traveler adding a stone to a cairn may be doing so from superstition, boredom, or genuine wayfinding intent, and the criterion is indifferent to which. What condition (iv) requires is only that no depositor need represent the resulting global structure, whatever their individual motivation.

⁵On path dependence generally, see Paul A. David, “Clio and the Economics of QWERTY,” *American Economic Review* 75, no. 2 (1985): 332–337, and W. Brian Arthur, “Competing Technologies, Increasing Returns, and Lock-In by Historical Events,” *Economic Journal* 99, no. 394 (1989): 116–131.

5 A Taxonomy of Three Domains Under One Criterion

Three domains of stigmergic berm can now be distinguished, not by surface resemblance to fields and cairns, but by the specific class of future action whose cost each domain's berm reduces, and by the mechanism, expulsion, attraction, or hybrid, that predominates in its formation. These domains classify berm *functions* rather than mutually exclusive objects: a single structure may reduce spatial, epistemic, and computational costs simultaneously, and the sections below give an example of exactly this for the software repository.

5.1 Territorial Berms

A *territorial berm* stabilizes control over, or navigability of, physical space. Its future-action class a_{future} consists of acts of boundary recognition, wayfinding, and claim maintenance: determining where cultivation rights end, which route across difficult terrain has already been judged passable, or which region has already been cleared of a given hazard. Field walls, cairns, trail systems, and the moraines of Section 2 are all territorial berms. The predominant mechanism varies by case, expulsion for field walls, attraction for cairns, hybrid for trails, but the reduced cost is always spatial: the cost of determining where, in physical space, an action can be performed and what its consequences with respect to prior claims will be.

5.2 Epistemic Berms

An *epistemic berm* stabilizes distinctions within a field of inquiry. Its future-action class consists of acts of problem formulation, proof search, and result verification: locating a viable line of argument, avoiding a route already shown to fail, adopting terminology that a community already recognizes, or checking a claim against an established body of prior results. A mathematical literature, a body of case law, a taxonomy in a natural science, and an established research vocabulary are epistemic berms. The mechanism here is predominantly attraction with an expulsion component: successful results and productive terminology attract further citation and reuse (attraction), while failed approaches, refuted conjectures, and abandoned notations are pushed to the margins of active practice, still recorded, in well-kept fields, but no longer part of the region a working researcher must actively re-traverse (expulsion of dead ends from the interior of live practice, without their deletion from the historical record). The reduced cost is epistemic: the cost of locating, verifying, and building on a viable line of inquiry, as opposed to the cost that would be faced by an inquirer working from an unmodified state of ignorance, W_0 , with no inherited terminology, no known failed approaches, and no established verification norms.

5.3 Computational Berms

A *computational berm* makes earlier work retrievable and executable by later processes, whether human or automated. Its future-action class consists of acts of resumption, recombination, and reuse: finding a prior implementation of a needed function, recovering a design decision recorded in a commit message, or locating a dormant fragment that becomes relevant to a query formulated much later. Version-controlled software repositories, and the multi-agent research systems discussed in Sections 6 through 8, are computational berms; a repository is simultaneously territorial in a narrower sense, since its directory structure and access permissions govern where a given action is allowed at all, which is a further illustration of the non-exclusivity just noted. The mechanism is hybrid: attraction, in that a well-organized repository or a well-indexed corpus

draws further contributions toward its existing structure rather than toward a competing, freshly initialized one; and something more precise than expulsion on the deprecation side, since version control does not push abandoned branches and superseded modules out of existence but out of the active working tree, while leaving them fully recoverable in history. Operational removal, in other words, does not entail historical erasure, which is precisely the structure the companion essay’s foundation, *Inscription Before Rendering*, treats as general: an artifact’s capacity to constrain later interpretation does not require that it remain part of the currently rendered state, only that it remain recoverable by a reader who knows to look. The reduced cost is computational in the literal sense: the cost, measured in developer time or machine search, of resuming a partially completed task rather than reconstructing it from an unmodified starting state.

5.4 The Common Criterion as the Basis of Classification

The three domains are not merely three analogies loosely gathered under a shared metaphor. Each is defined by the same inequality from Section 4, instantiated with a domain-specific cost function $C(\cdot)$ and a domain-specific future-action class a_{future} :

Territorial: $C_{\text{spatial}}(a_{\text{future}})$ = cost of boundary recognition or wayfinding,

Epistemic: $C_{\text{inquiry}}(a_{\text{future}})$ = cost of locating or verifying a viable line of argument,

Computational: $C_{\text{resumption}}(a_{\text{future}})$ = cost of resuming, recombining, or reusing prior work.

A structure qualifies as a berm in a given domain precisely when accumulated local deposits persistently lower the relevant expected cost and thereby bias successors toward reinforcing the same organization rather than initiating a competing one. This is what licenses treating field walls, mathematical literatures, and software repositories as instances of one phenomenon rather than as three phenomena connected only by loose family resemblance: each satisfies the identical formal condition, differing only in what is being deposited, what the future action is, and in what units cost is measured. The taxonomy is therefore operational rather than analogical, and it is this operational basis that allows the framework to be applied, in Sections 7 and 8, to a case where the analogy alone would be too weak to carry the argumentative weight required: the question of what an AI system inherits, silently, when it appears to solve a long-standing open problem.

6 Computational and Epistemic Berms in Practice

6.1 Repositories as Computational Berms

A software repository maintained over years exhibits the full structure of a computational berm. Commits, filenames, comments, abandoned prototypes, issue threads, and message histories together form a research terrain rather than a simple linear record. A phrase deposited once, for instance in a hidden HTML comment or an offhand commit message, can remain dormant for months with no successor accessing it, and then become recoverable when a later inquiry happens to supply the search cue that makes it relevant again. When this occurs, the old inscription does not simply return unchanged, as a static memory retrieved intact. It rejoins a terrain that has, in the intervening months, developed elsewhere: new modules have been added, other abandoned branches have accumulated nearby, the vocabulary used in commit messages has shifted. The recovered residual becomes a bridge between the state of the project when it was deposited and the state of the project when it is rediscovered, and the bridge itself is new structure, not merely an old structure made visible again.

This is the computational-berm analogue of the expulsion/attraction distinction from Sections 2 and 3. Deprecated code, abandoned experiments, and superseded designs are expelled from the active branch, pushed into history rather than deleted, exactly as fieldstones are pushed to the margin rather than destroyed. Meanwhile a well-organized module structure, once established, attracts further contributions to conform to its existing organization rather than to a freshly invented one, exactly as a cairn attracts further stones. A mature repository is therefore a morainoid tessellation in the sense of Section 2: a field of partially purified working interiors (active modules) separated by seams of accumulated heterogeneity (deprecated code, historical branches, superseded designs) that a later contributor can, if the seams are legible, read as a record of where earlier work encountered resistance.

6.2 Scientific Literature as Epistemic Berm

Scientific literature expands the same mechanism to civilization scale. Mathematics in particular is a stigmergic berm made from successful proofs, failed strategies, counterexamples, terminology, conjectures, surveys, specifications, and institutional verification practices, journal refereeing, reproducibility norms, citation conventions. Researchers entering an active field do not inherit only a set of established results, in the way a student might memorize a table of theorems. They inherit a transformed search landscape in which enormous regions have already been cultivated (well-understood subareas with established techniques), obstructed (approaches known to fail, documented as such in the literature so that successors need not rediscover the failure independently), or marked (terminology and notation that concentrate attention on the distinctions a community has found productive, at the expense of alternative framings that have fallen out of use).

This inherited landscape is not incidental background to research; on the criterion of Section 4, it is the primary mechanism by which research becomes tractable at all. A researcher working from W_0 , an unmodified state with no inherited terminology, no catalogue of known failed approaches, and no established verification norms, faces a radically higher expected cost for locating a viable line of argument than a researcher working from B_t , the current, heavily cultivated state of an active mathematical field. The difference between these two costs is not a difference in raw intelligence between the two researchers; it is a difference in the terrain each has inherited.

7 Apex Attribution: The Navier-Stokes Dispute

The framework developed above yields a specific and deliberately narrow criticism of headlines and claims to the effect that an artificial intelligence system has “solved” a long-standing open mathematical problem. The criticism is not a claim about capability, and it should be stated in a form that does not depreciate whatever genuine new construction the system in question has contributed. Suppose an automated system, or a human researcher working with such a system as a tool, produces a valid new argument that meaningfully advances a previously open problem. The framework does not deny that this constitutes real, attributable work performed by the system or by the human-system pair.

The narrower claim is about *apex attribution*: the practice of naming the final traverser of a terrain as though that traverser had originated the terrain itself. The agent that closes a long-standing gap enters a research landscape whose problem statement, technical vocabulary, admissible solution conditions, catalogue of failed prior routes, local lemmas, and community verification practices were established, in most cases over decades, by generations of human mathematicians. This entire inherited structure is, in the vocabulary of Section 5, an epistemic berm of the type dis-

cussed in Section 6.2, and its existence is precisely what makes the final result reachable at a cost the closing agent could pay.

7.1 A Documented Case

A dispute reported in early September 2026 supplies an unusually well-documented instance of the mechanism, and one worth treating carefully because several of its factual claims remain contested between the parties. On August 15, 2026, NYU mathematician Tristan Buckmaster and Anthropic researcher Levent Alpöge, working independently of both of their employers on a purely personal collaboration, released finite-time blow-up results, which they report as Lean-verified, for the Euler and Boussinesq equations, extending a forcing construction due to Diego Córdoba and Luis Martínez-Zoroa; their results were understood by researchers in the field as a plausible stepping stone toward the Navier-Stokes existence and smoothness problem, one of the Clay Mathematics Institute’s seven Millennium Prize problems.⁶ According to Buckmaster’s public statement, word of the unpublished progress reached OpenAI over the following days, and a team led by Sébastien Bubeck, who leads OpenAI’s mathematics research, produced a multi-agent proof extending the argument to the full Navier-Stokes case over the following weekend. OpenAI has itself stated that it undertook the effort specifically because of reports that a team, understood to be affiliated with a rival laboratory, was close to a Millennium Prize result, which is on its face an admission that a rumor of human progress, rather than an independent, internally generated research direction, supplied the address into the mathematical berm that the subsequent multi-agent search then traversed.

Buckmaster has additionally stated that, on raising the matter with OpenAI, he was offered a choice between two publication arrangements: one in which he and Alpöge would publish first with OpenAI following a day later and crediting them as “the closest humans to the problem”; and one in which he alone, without Alpöge, would publish a paper stating that an OpenAI model had also solved the problem, with Alpöge’s name removed because of his Anthropic affiliation. Buckmaster states that when he declined both and indicated he would instead publish independently, Bubeck responded, “Why would you ruin your career?” and, when Buckmaster pressed the point, “If you don’t want me to be nice, then I don’t have to be nice.” Bubeck has publicly characterized Buckmaster’s allegations as false and inflammatory, and has stated that OpenAI recognizes the priority of Buckmaster and Alpöge’s work; the specific claim that Alpöge’s removal was requested is among the points he disputes. This essay treats the credit-arrangement and pressure allegations as Buckmaster’s reported account, contested by Bubeck, rather than as established fact; nothing in the argument that follows depends on resolving that dispute.

7.2 What the Case Establishes Independent of the Dispute

Separating the contested interpersonal allegations from what both parties have effectively conceded is instructive. Neither side disputes that OpenAI’s effort began in response to a rumor about unpublished human progress, or that the resulting result was produced by a large multi-agent system operating over roughly 88 hours followed by machine-assisted formal verification. The parties do dispute whether OpenAI’s argument and Buckmaster and Alpöge’s unpublished route share a more specific lineage; OpenAI maintains that its proof, and even its treatment of the Euler case, differs from theirs. The apex-attribution argument that follows does not require resolving that narrower dispute. It requires only the uncontested chronology:

⁶Simon Willison, “Notes on the OpenAI / Buckmaster-Alpöge Navier-Stokes dispute,” September 2026; TechCrunch, “OpenAI fought dirty on career-making math problem, says NYU mathematician,” 8 September 2026.

attributed deposits → aggregated research terrain → machine traversal → dependency compression → apex att

Neither side disputes that the system entered a mathematical field structured by an extensive public literature on blow-up, forcing constructions, regularity criteria, and related equations, built up over more than a century and extended most recently by Córdoba, Martínez-Zoroa, Buckmaster, and Alpöge, or that the final claim was evaluated against human-established formal and disciplinary verification standards. Public mathematical literature does not, in the ordinary case, lose provenance the way a de-identified interaction log does; papers retain named authors and citation trails. What happens instead, at the point a result is announced, is that provenance is *compressed*: an entire dependency chain running back through decades of named contributions is folded into a single headline event attached to whichever system produced the final, verified step. “Provenance loss” in the stronger sense, where the underlying contributions become genuinely untraceable rather than merely uncited in a headline, is the special transition associated with de-identified interaction data, taken up directly in Section 8.

Attributed deposits, Córdoba and Martínez-Zoroa’s forcing construction, Buckmaster and Alpöge’s extension of it, and more than a century of prior Navier-Stokes literature, are aggregated into a research terrain that a multi-agent system can then traverse. Traversal itself does not erase the deposits; it compresses their dependency structure into a form, the final proof object and the headline announcing it, in which the specific load-bearing contributions are no longer visible without deliberate reconstruction. A headline crediting “AI” with having solved Navier-Stokes performs exactly that compression.

Framed this way, the criticism is difficult to dismiss as anti-AI rhetoric, because it applies with equal force, and has historically been applied with equal force, to human apex attribution: a mathematician who closes a long-standing gap using techniques, lemmas, and a research programme built by a community over decades is conventionally credited with “solving” the problem, even though a more precise attribution would credit the closing step while acknowledging the inherited terrain. The Navier-Stokes case sharpens this only by making the terrain-traversal ratio unusually visible: a rumor supplied the address, an 88-hour multi-agent run supplied the traversal, and the terrain itself, the forcing construction and everything beneath it, took nearly a year of independent human work to reach the point at which that traversal became possible.

8 The Obscured Berm

A further complication deserves treatment here because the Navier-Stokes case gives it unusual concreteness, though the essay’s purpose remains to establish the conceptual problem rather than to adjudicate the specific, still-unresolved factual dispute. If researchers closing a gap of the kind discussed in Section 7 make extensive use of an interactive AI system to conduct unpublished investigation, exploring failed approaches, testing partial arguments, iterating on formulations, before arriving at a published result, then the record of that interaction constitutes a second berm distinct from the published literature discussed in Section 6.2. Call it the *obscured berm*: a body of epistemically productive deposits, failed attempts, partial insights, discarded formulations, that would, if preserved and made accessible in the way a mathematical literature is preserved and made accessible, function exactly as Section 6.2 describes, lowering the future cost of related inquiry for successors.

The mechanism by which this happens needs to be stated precisely, because de-identification does not prevent the platform or a successor model from discovering and using such a deposit;

it prevents outside observers from recovering *who* produced it. Recoverability and attributability are therefore separate properties, and the companion essay’s recurrence product from Section 1 does not by itself distinguish them. A further function is needed:

$$\text{Attr}(m, a) = \text{Trace}(m, a) \cdot \text{Identify}(a) \cdot \text{ExposeDependency}(m),$$

where $\text{Trace}(m, a)$ holds when a later artifact a can be shown to depend on an earlier deposit m , $\text{Identify}(a)$ holds when the depositor of m can be named, and $\text{ExposeDependency}(m)$ holds when that dependency is actually disclosed rather than merely discoverable in principle. De-identification can leave operational recurrence fully intact while driving attribution to zero:

$$\text{Rec}(m; A_j, W) > 0, \quad \text{Attr}(m, a) = 0.$$

This asymmetry is more precise, and more damning, than saying that discovery has been prevented. The obscured berm works precisely because deposits remain epistemically effective, for whichever system retains access to them, after their provenance has ceased to be recoverable by anyone else. But interaction logs of this kind are typically retained, if at all, in de-identified or aggregated form for reasons of privacy and commercial confidentiality, and de-identification, whatever its other merits, is exactly what drives Attr toward zero while leaving Rec untouched.

8.1 Three Distinct Provenance Concerns

Buckmaster reportedly asked OpenAI directly whether its Navier-Stokes model had been trained on, or had access to, his and Alpöge’s Codex sessions. OpenAI’s public responses have distinguished between whether the solving agents retrieved specific user documents during the solving run and whether de-identified usage data more generally contributes to model improvement, stating the former did not occur while declining to rule out the latter. That distinction, whatever it settles about this particular case, is instructive because it shows that a single word, “access,” was doing service for at least three separable claims:

- Content appropriation:** Did the proof draw on Buckmaster and Alpöge’s specific unpublished material?
- Training appropriation:** Did their (or others’) Codex sessions, de-identified, shape the model that later computed the proof?
- Directional appropriation:** Did knowledge that a team was close to a result change where OpenAI concentrated its search?

These are not equally well evidenced. Content appropriation is the strongest claim and the one OpenAI has most directly denied; the public record does not establish it. Training appropriation remains genuinely open, since a statement that a company “cannot rule out” improvement from de-identified data is a statement of uncertainty about its own provenance boundary, not a specific finding either way. Directional appropriation, by contrast, is close to conceded: OpenAI’s own account cites rumors of a rival team’s progress as the reason it committed resources to the problem when it did. A rumor, on this account, functioned as an address into an already well-populated region of the mathematical berm, sharply reducing the uncertainty about where concentrated, well-resourced search would be worth deploying, without requiring that any unpublished document itself be read.

8.2 Unattributed Collective Epistemic Subsidy

The training-appropriation concern generalizes well beyond this single dispute, and the generalization is where the concept does its most useful work. Any interaction with a widely used AI

research tool can be epistemically productive whether or not it succeeds: a user tests a conjecture and exposes a contradiction, rejects an unproductive approach, requests a missing lemma, corrects a hallucinated step, or repeatedly steers a model toward a more productive representation of a problem. None of these sessions need contain anything resembling a solution. Collectively, and once aggregated across a large user base, they disclose the shape of a search space: which regions are promising, which routes are dead ends, which conceptual bridges are useful, and which errors recur.

De-identification removes the identity attached to each such contribution without removing its epistemic content. This licenses a description that is more precise than either “theft” or “independent discovery”: an *unattributed collective epistemic subsidy*, in which a large, diffuse population of users, including working mathematicians, students, and casual questioners who never publish anything, jointly enrich a computational berm through ordinary use, after which the enrichment is credited entirely to the system that absorbed it rather than to any of the population that produced it. Even an unsuccessful student session can be a deposit in this sense, in exactly the way a discarded fieldstone is a deposit in Section 2: valuable not for what it accomplishes locally but for what it adds, once aggregated with enough similar residue, to the terrain a later, better-resourced traverser can exploit.

This yields a dependency structure considerably broader than the one Section 6.2 stated for public literature alone:

published mathematical record
 +unpublished user investigations (including failed ones)
 +expert corrections supplied in the course of ordinary use
 +problem-selection signals (which questions users find worth asking)
 +engineer-designed orchestration and verification
 → candidate proof.

A claim that this arrangement constitutes theft is, as argued above, too strong for what the public evidence supports; a claim that nothing improper has occurred because no single document was retrieved is, by the same token, too weak, since it addresses only content appropriation while leaving training and directional appropriation untouched. There is also an asymmetry worth naming precisely: a laboratory may be able to state truthfully that no document was retrieved during a solving run while remaining unable to state what accumulated interaction made the solving system capable of the run at all.

no document retrieval $\nRightarrow$ no intellectual dependence.

This is not merely a plagiarism problem, though it can shade into one. It is a custody problem specific to computational berms of the kind introduced in Section 6.1: a private interaction log, unlike a public repository, has no morainoid tessellation visible to outside observers, no seams between purified interiors and expelled residue that a later contributor can read as a record of where earlier work encountered resistance. The terrain may still lower future cost for whoever retains access to it, satisfying the operational criterion of Section 4 with respect to that restricted population, while remaining invisible, and therefore uncredited, to the wider community whose published literature forms the more familiar, public berm of Section 6.2. The strongest responsible formulation available from the public record is therefore something like: the system produced and formally verified a candidate solution after operating on a capability base built from published mathematics, human-designed orchestration, a rumor that redirected concentrated search

toward a specific human-populated region of the berm, and potentially de-identified user interactions whose individual contributions cannot presently be traced. Establishing this asymmetry is sufficient for the purposes of this essay; adjudicating what obligations, if any, follow from it, and whether de-identification can in principle preserve epistemic utility for a private berm while foreclosing external attribution entirely, is left as the subject of a separate essay.

9 Conclusion

Three cases, a field wall, a cairn, and a mound of anthropogenic dark earth, have supplied a single operational criterion: a structure is a stigmergic berm when accumulated local deposits persistently lower the expected cost of some class of future action, biasing successors toward reinforcing the organization already present rather than initiating a competing one from an unmodified baseline. This criterion, rather than surface resemblance, is what licenses treating physical boundaries, scientific literatures, and software repositories as instances of one phenomenon, distinguished by the specific cost each lowers and by the relative weight of expulsion versus attraction in their formation.

Applied to the question of what a closing agent inherits when it appears to solve a long-standing open problem, the framework yields a criticism narrower and more defensible than a blanket denial of machine capability: the target is apex attribution, the crediting of a final traversal with having originated a terrain that generations of prior, largely uncredited deposits actually built. Applied further to the possibility of private, de-identified interaction data functioning as an epistemically productive but uncredited substrate, the framework identifies a custody problem distinct from, though related to, ordinary questions of plagiarism and citation, one turning on the gap between recurrence and attribution rather than on recurrence alone.

The essay's central result is not, in the end, the expected-cost inequality by itself, useful as that inequality is for classification. It is the distinction between two ways the past can act on a successor:

- semantic inheritance: a successor recovers what an earlier agent specifically recorded;
- morphological inheritance: a successor encounters a world altered by accumulated acts, whether or not any particular act remains individually recoverable.

The companion essay formalizes the first. This essay formalizes the second, and the two compose, as Section 1 argued, rather than nest.⁷

The concluding principle is this: collective intelligence begins not when many agents think as one, but when their residual acts accumulate into a terrain that lets successors begin somewhere other than the beginning. The berm, not the swarm, is the primary unit of distributed historical construction; whatever coordination or apparent collective intention is read into a population of agents traversing a shared terrain, the terrain itself, and the accumulated history of local refusal and local attraction that built it, is doing most of the explanatory work.

⁷The claim that cognitive and organizational competence can be distributed across a person or population and a structured external environment, rather than residing solely in individual internal representation, is developed at length in Edwin Hutchins, *Cognition in the Wild* (Cambridge, MA: MIT Press, 1995); the berm, on the account given here, is one of the material structures across which such distributed competence accumulates.

References

- [1] Pierre-Paul Grassé, “La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp.,” *Insectes Sociaux* 6, no. 1 (1959): 41–80.
- [2] Guy Theraulaz and Eric Bonabeau, “A Brief History of Stigmergy,” *Artificial Life* 5, no. 2 (1999): 97–116.
- [3] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems* (Oxford: Oxford University Press, 1999).
- [4] Andrew Fleming, *The Dartmoor Reaves: Investigating Prehistoric Land Divisions* (London: Batsford, 1988).
- [5] Bruno Glaser, Ludwig Haumaier, Georg Guggenberger, and Wolfgang Zech, “The ‘Terra Preta’ Phenomenon: A Model for Sustainable Agriculture in the Humid Tropics,” *Naturwissenschaften* 88, no. 1 (2001): 37–41.
- [6] William I. Woods et al., eds., *Amazonian Dark Earths: Wim Sombroek’s Vision* (Dordrecht: Springer, 2009).
- [7] Paul A. David, “Clio and the Economics of QWERTY,” *American Economic Review* 75, no. 2 (1985): 332–337.
- [8] W. Brian Arthur, “Competing Technologies, Increasing Returns, and Lock-In by Historical Events,” *Economic Journal* 99, no. 394 (1989): 116–131.
- [9] Edwin Hutchins, *Cognition in the Wild* (Cambridge, MA: MIT Press, 1995).
- [10] Flyxion, *Inscription Before Collusion: External Memory, Distributed Recurrence, and the Appearance of Agent Swarms*, September 2026. <https://standardgalactic.github.io/userland/inscription-before-collusion.pdf>
- [11] Tristan Buckmaster, public statement announcing finite-time blow-up results for the Euler and Boussinesq equations and describing a parallel OpenAI effort, 8 September 2026.
- [12] OpenAI, blog post announcing a proof of the Navier-Stokes existence and smoothness problem, 8 September 2026.
- [13] Simon Willison, notes and reconstruction of the OpenAI / Buckmaster-Alpöge Navier-Stokes dispute, September 2026.
- [14] TechCrunch, “OpenAI fought dirty on career-making math problem, says NYU mathematician,” 8 September 2026.
- [15] Fortune, “OpenAI says it cracked Navier-Stokes, one of math’s grand challenges,” 8 September 2026.
- [16] Tom’s Hardware, “OpenAI’s breakthrough solution for the elusive Navier-Stokes problem overshadowed by plagiarism controversy,” September 2026.
- [17] Dataconomy, “NYU Professor Challenges Origin Of OpenAI’s AI Proof,” 9 September 2026.