Fractal Non-Closure

Larsen James Close

2026

The instance does not close; the law of instancing may.

Abstract

Fractals make non-closure visible because under magnification, the non-attainment of a terminal witness is more than mere absence. Smooth objects close under zoom by yielding a tangent element. Exactly self-similar objects close by return. The target fractal case does neither: the instance under magnification does not settle to one form or one finite orbit. What may stabilize instead is a law of instancing: a scenery distribution, invariant measure, or statistical pattern one level up. The asymmetric commuting square is the minimal grammar for this situation. It certifies lawful reading across loss, but it does not say which distinctions a reading has retained; the one-point reading commutes with everything while retaining no distinction. Fractal non-closure therefore names three tasks: identify the closure not attained under magnification, identify the law-level closure that may replace it, and state the domain's collapse condition so that the resulting law remains informative for that domain.

1

The word *fractal* does not close cleanly.

That is usually treated as an embarrassment. A mathematical term should admit a definition; if the definition leaks, one looks for a sharper definition, or else retreats to examples. Coastlines, Cantor sets, Brownian paths, Julia sets, strange attractors. The list grows because the class refuses to settle. Every candidate criterion seems either too narrow, excluding things that made the term necessary, or too broad, admitting things that no longer carry the phenomenon.

The non-attainment may be data.

A definition is a closure of a term. It says: here is the witness; the search has landed. But the phenomenon fractals expose is precisely that magnification need not land on a witness of the same kind. Zoom into a smooth curve and the process stabilizes: the tangent line appears, and the question has closed to an element. Zoom into an exactly self-similar object and the process returns: not to one element, but to a finite pattern or orbit. Zoom into the target fractal case and the process does not terminate in either way. No single form arrives as the answer. No finite return exhausts the motion. What stabilizes, if anything does, is not one picture but a law governing the pictures.

So the definitional leak is not accidental. The term may be being asked to close at a level where the phenomenon does not supply the relevant witness. The instance under magnification does not close to an element. The law of instancing may.

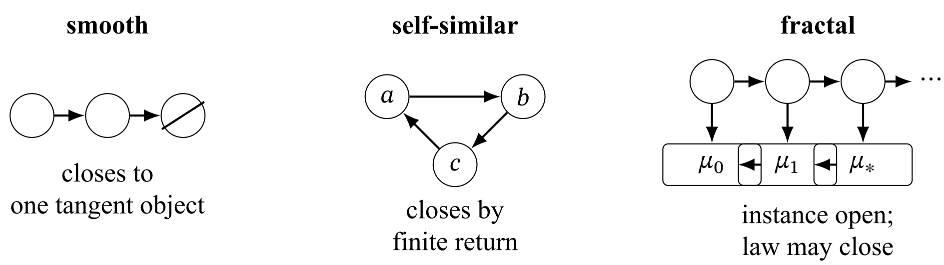

This is not yet a definition of fractal. It is a discipline for saying what a definition would have to respect. The useful question is not whether the object closes simply. It is:

under which operation, toward which target, and at which level?

Without those indices, “closure” and “non-closure” are too blunt to say what fractals show.

2

Magnification is the operation. That is what makes fractals the right lens.

In many domains non-closure is easy to sentimentalize. One can say inquiry continues, language proliferates, experience exceeds articulation, and all of that may be true. But it can also be too easy. Without a specified operation, non-closure floats. It becomes a mood rather than a structure.

Fractals are better behaved. The operation is concrete: zoom, rescale, pass to scenery, renormalize. The departures are measurable. Dimension gaps, tangent measures, beta numbers, scenery flow, and statistical self-similarity are all ways of asking how magnification does not close to a single ordinary witness. They do not all say the same thing, and they should not be forced into one criterion. Their shared importance is more basic: they measure non-attainment of a target closure under an operation precise enough to push back.[1]

This paper is not a replacement for that toolkit. It is a discipline for asking when a result from that toolkit remains informative for a stated inquiry: which operation, which target, which level, and which retained distinctions?

Smoothness is the limiting contrast. A smooth object is not merely less complicated. It is a case in which magnification answers the question. The tangent object appears as an element. The local view becomes simpler as the scale decreases; the process terminates because it has found the right kind of witness.

Exact self-similarity is a second contrast. The process may not settle to one element, but it returns. Magnification brings the object back to itself, or to a finite cycle of forms. Closure is not terminal; it is orbital. The answer is not "this one element" but "this return pattern."

The target fractal case is different from both. It is not smooth terminal closure, and not merely exact return. It is a case where the sequence of views continues to generate difference, while a higher-level reading may stabilize. The object's identity is carried by an invariant law rather than by one form.

This is the first positive formulation:

A fractal is a set whose tangent object is an ensemble rather than an element; identity under magnification is carried by law, not by one form.

That sentence is a target, not a theorem. It must be tested against the actual geometry. But it already improves the question. It does not ask the class to close where its phenomenon may not supply the relevant witness.

3

The minimal shape is the asymmetric commuting square.

There is a base motion: zoom the instance. There is a higher motion: renormalize the law. There is a reading from the base to the higher register: send the current view to the object that records what survives the view. The square says that reading after zooming agrees with renormalizing after reading.

$$
\begin{tikzcd}
\boxed{\begin{array}{c}\text{instance}\\ \text{under zoom}\end{array}} \arrow[r, "\text{magnify } f"] \arrow[d, "\text{read } q"'] \arrow[dr, phantom, "\begin{array}{c} q \circ f = g \circ q \\ \textit{lawful passage across loss} \end{array}"] & \boxed{\begin{array}{c}\text{instance}\\ \text{under zoom}\end{array}} \arrow[d, "\text{read } q"] \\
\boxed{\begin{array}{c}\text{law / scenery}\\ \text{space}\end{array}} \arrow[r, "\text{renormalize } g"'] & \boxed{\begin{array}{c}\text{law / scenery}\\ \text{space}\end{array}}
\end{tikzcd}
$$

In symbols:

$$q \circ f = g \circ q.$$

That equation is small enough to disappear inside its own familiarity. It should not. It is the minimal grammar of lawful reading across loss.

The top row is the motion at the level where the instance remains open. The bottom row is the motion at the level where a law may close. The vertical map is the lossy passage between them. This passage is asymmetric: the lower level does not, in general, recover the upper. A scenery distribution is not the patch of the set that produced it. A measurement is not the world measured. An abstraction is not the concrete process it abstracts.

Reconstruction theorems do not remove this asymmetry. They add hypotheses under which a chosen reading can be inverted, enriched, or supplemented by other data. That is valuable, but it is a further theorem, not a consequence of the square. The forward claim here is narrower: commutation says that q respects motion; it does not say that q has retained enough structure for recovery.

The square buys one thing exactly. It buys lawfulness. Whatever the reading carries upward moves compatibly with the motion below.

It does not say which identifications matter. It does not say the reading has retained the differences the inquiry needs. The square can commute because the reading has carried structure upward. It can also commute because the reading has identified everything.

That difference is the one-point calibration.

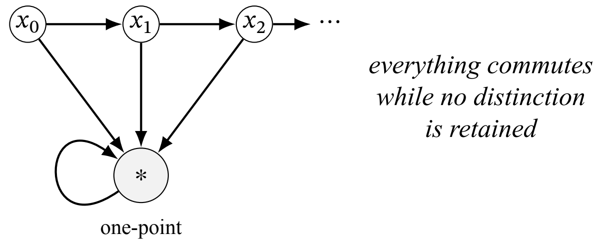

The one-point reading sends every instance to the same point. It commutes with every motion, because no distinction remains that could fail to commute. It closes every observed orbit, because the observation is already closed. And it retains no distinction, because every instance has the same reading.

This is not a pathology of the square. It is the correct behavior of a bare signature that emits nothing. If the only structure is "next state," then the canonical behavior seen by the signature is empty; the one-point behavior is terminal. The repair is not to distrust squares, but to say what the domain emits at each step.

So the square is universal, but not sufficient by itself. The square gives the lawful form of reading. The domain must supply the collapse condition: the point at which the reading no longer reads the phenomenon.

This is the same separation that appears in naming-channel language, but with the pieces shifted into a lossy register. Here "naming-channel" is local shorthand rather than a standard term in dynamics: a carrier is presented or named, a reader sends it into another register, and one asks both whether the reading is lawful and whether it has lost the carrier distinctions the use requires. The equation saying that a presented name reads back correctly is the lawful half of naming. Faithfulness is the second equation: reading loses no carrier distinction, equivalently the reading is injective. In the fractal square, commutation is analogous to the lawful half, not to faithfulness. The missing second condition cannot be global injectivity, because a law-level reading must compress the instance. Its analogue of faithfulness is domain-relative noncollapse: the reading has to preserve the distinctions by which this phenomenon remains this phenomenon.

4

This is where non-closure becomes precise.

Non-closure is not the absence of closure. It is the refusal of a specified operation to close toward a specified target at a specified level. Once those indices are present, several distinct structures appear.

Terminal closure: after finitely many steps, the orbit settles to one element. Smooth magnification has this shape when the tangent object appears.

Cyclical closure: the orbit does not settle, but it returns with positive period. Exact self-similarity has this shape when magnification recovers the same form or a finite cycle of forms.

Instance openness: the orbit neither settles nor returns. Pure escape to infinity has this shape, so instance openness alone does not specify a law-level reading.

Law closure: an observation of the open instance closes in a factor system one level up. This is the shape fractals make available: the current view remains open while the law governing the views stabilizes.

The formal core records exactly this much.[2] A `RenormSystem` has states and one step operation. Its orbit can be element-closed, orbit-closed, or open. A `Factor` is the asymmetric square: a second system, an observation map, a law-level closure relation, and the commutation equation. It is simply a lawful reading map together with its dynamics; it carries no automatic non-collapse guarantee. `LawClosed` says the observed orbit closes or converges to an invariant law state in that factor and that the limiting law is not declared degenerate. `OpenWithLawClosure` is the paired condition: open at the base level, law-closed one level up. The name is deliberately neutral; the formal shape is broader than the intended fractal instantiation.

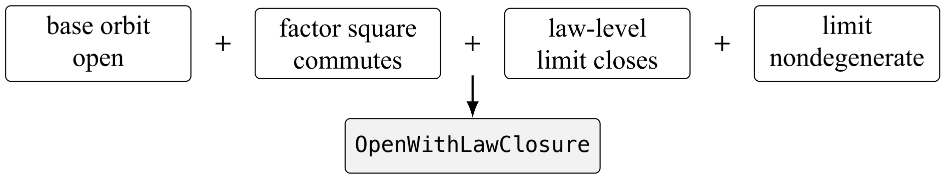

open below; closed lawfully one level up

The formalism deliberately avoids measure theory at this stage. That is not a defect. Finite terminal closure remains available as one possible law-level closure relation, but the formal core no longer requires it. It isolates the structural claim:

closure may be unattained at the level of the instance and hold at the level of the law.

The next layer is not a generic grade of preservation. A law-level reading may legitimately compress; otherwise it would simply reproduce the instance at another level. What matters is the domain-specific collapse condition: which compressions still count as readings of the phenomenon, and which do not.

The formal core therefore names collapse locally. `CollapsesPair` says that a factor identifies two base states that were distinct. This is not yet a criticism. It is a witness of loss. Whether that loss matters depends on the landscape being read. `Degenerate` remains supplied by the factor, because Dirac collapse, terminal collapse, smooth collapse, and statistical collapse are not the same condition.

That is the role of the one-point theorem. It is not a side note. It is the retention test for every proposed law. The Lean file now proves the test against itself: pure escape to infinity satisfies `OpenWithLawClosure` under a one-point factor whose single law-state is declared nondegenerate. The theorem is not a positive fractal example. It is the formal mark showing where the abstract layer remains intentionally underspecified.

The closure relation carries the same freedom, and the core names it too: if a factor declares its `ClosesTo` relation universal, law closure is automatic once any invariant nondegenerate law state is available. `Degenerate` and `ClosesTo` are therefore two declaration points, not one. Each must be supplied by the domain before law closure asserts anything about the phenomenon.

5

The theorem has a philosophical form.

Commutation certifies that whatever survives the reading moves lawfully. It does not certify that the reading has avoided the relevant collapse. Lawfulness and collapse are independent questions.

This problem recurs anywhere a lower register is read into a higher one. The same square appears whenever measurement, abstraction, oversight, or specification reads a process through a proxy and then asks the proxy to move lawfully. Measurement reads world into instrument. Abstraction reads concrete process into quotient. Scoring reads behavior into number. Computation reads implementation into specification. In each case, the square can commute while the reading has already crossed its collapse condition. A score can move lawfully and still miss the behavior. A quotient can respect dynamics and still identify the cases the inquiry needed to keep apart. A measurement can be repeatable and still measure the wrong thing.

This is the paper's quiet contact with corrigibility: every oversight regime is a lawful reading whose collapse condition may or may not have been stated.

Fractals provide a disciplined case because the loss is neither accidental nor merely practical. The loss is constitutive of the phenomenon. If magnification closed to one element, the object would be smooth in the relevant sense. If it closed by finite return, it would be exactly self-similar in the relevant sense. The target case is the one where the instance keeps generating difference.

But that does not license any law. The law must be nondegenerate in a way the domain can defend. In a scenery-flow setting, the distinction might be between a Dirac limiting measure and a nontrivial invariant distribution. In a rectifiability setting, it might be the non-stabilization of tangent planes in the smooth manner. In a statistical-complexity setting, it might be the minimal predictive structure that remains after individual configurations continue varying.[3]

The point is not that these are the same measure. The point is that each has to answer the same three questions: which closure is not attained under the operation, which law-level closure is being claimed, and which collapse condition keeps that law informative for the domain?

6

The strongest objection is that coherence needs closure.

Grant it. Consequence chains must close somewhere. A proof that never lands is not a proof. A measurement that never stabilizes is not a measurement. A theory that refuses every closure is not deep; it is unusable. If non-closure is made primitive without qualification, it destroys the conditions under which it can be known.

The answer is not to weaken closure. The answer is to index it.

The fractal case is coherent because closure changes level. It does not vanish. The instance under magnification remains open. The law-level reading may close. The reader receives not a final picture but a stable grammar for the production of pictures. This is why the case is neither total closure nor total non-closure. It is a structured mixture:

closed enough to be legible, open enough to produce.

That sentence is the bridge back to completability. Terminal completion reaches a fixed point and is done. Cyclical completion returns. Graceful completion closes locally while remaining globally open. Fractal non-closure is not identical to graceful completion, but it gives a mathematical lens into the same structural possibility: local or adjacent closure without global terminal closure.

The square makes the adjacency visible. The base level continues. The factor level closes. The vertical reading is the condition of legibility, but each domain must say when the reading has collapsed what needed to remain in play. That is why the one-point calibration matters. It is the formal reminder that legibility through a one-point reading is legibility of a terminal factor, not automatically legibility of the phenomenon being read.

7

The paper's central claim can now be stated without mystification:

Fractal non-closure is non-attainment of element-level closure under magnification together with possible nondegenerate law-level closure under a lawful reading.

Every word matters.

Non-attainment of element-level closure: the phenomenon is not captured by a single tangent object.

Under magnification: the claim is indexed to an operation, not made in the abstract.

Possible law-level closure: the law is not guaranteed by the non-attainment below. It must be shown.

Nondegenerate: the law has not crossed the collapse condition relevant to the domain; the limiting law has not identified states whose distinction was constitutive of the phenomenon at the instance level.

Lawful reading: the square must commute; the reading must respect the operation it claims to read.

This does not close the definition of fractal. It disciplines why the definition resists closure. The refusal is not a fog around the term. It is an instruction: do not ask for an element where the phenomenon is carried by a law; do not accept a law without asking which distinctions it retains; do not call openness a defect when it is the condition under which the law becomes visible.

The old question was: what is a fractal?

The better task is to say which closure magnification defeats, which law-level closure may replace it, and which nondegeneracy condition keeps that replacement legible.

8

Here is the first concrete instantiation.

Let V be a space of normalized views: sceneries, pointed rescalings, or local windows of a set or measure after recentering and rescaling. Let

$$f : V \to V$$

be one step of magnification. A chosen initial view $v \in V$ produces the base orbit

$$v,\ f(v),\ f^2(v), \ldots$$

This is the level at which the target fractal refuses ordinary closure. If the orbit settles to one view, the case is smooth in this register. If the orbit falls into a finite return pattern, the case is orbital or exactly self-similar in this register. The target case is that neither happens.

Now pass to the factor level. Let $\mathcal{P}(V)$ be a law space over views, and let

$$g = f_* : \mathcal{P}(V) \to \mathcal{P}(V)$$

be pushforward by magnification. The empirical scenery law of the first N views is

$$\mu_N(v) = \frac{1}{N} \sum_{k<N} \delta_{f^k(v)}.$$

If these empirical laws converge, write

$$q(v) = \lim_{N\to\infty} \mu_N(v).$$

Then the abstract square is no longer merely promised. It is enacted at the level of empirical laws, and then it closes at the limit. At finite N,

$$\mu_N(f(v)) = f_*\mu_N(v).$$

Meanwhile $\mu_N(f(v))$ differs from $\mu_N(v)$ only by the exchanged boundary term $(\delta_{f^N(v)} - \delta_v)/N$. In the limit,

$$q(f(v)) = f_*q(v) = q(v).$$

One hypothesis works quietly in that chain and should be named. The boundary exchange gives $q(f(v)) = q(v)$ on its own, in any mode of convergence where a term of total mass $2/N$ vanishes. The middle equality, reading the pushforward through the limit, additionally asks magnification to be continuous in the topology where the laws converge. Keeping those two hypotheses separate is part of the discipline the square enforces.

The base instance remains open as a sequence of views, while the law of those views closes as an invariant distribution. This is the mathematical shape the paper has been circling:

the scenery need not close; the scenery law may close.

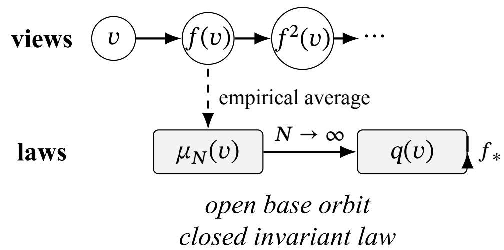

The middle-third Cantor measure gives a concrete model. Let C be the middle-third Cantor set with its natural Cantor measure, and code a typical point by an infinite word over the alphabet $\{0, 2\}$. At triadic scales, the symbolic pointed scenery is governed by the shift on this address space: V is a space of symbolic pointed sceneries, f is the left shift, and $\mathcal{P}(V)$ is the space of laws on those sceneries. If the address is eventually periodic, the base orbit closes by finite return; that is the exact self-similar exception. For a Bernoulli-typical address, the shift orbit is not eventually periodic, so the instance sequence remains open, while its empirical laws converge to the Bernoulli shift-invariant law. That limiting law is not a Dirac mass: both first-symbol cylinders have positive mass. It is not the one-point reading either, because it still distinguishes the left and right cylinder choices that make the pointed scenery nontrivial. The collapse condition is therefore avoided not by global injectivity, but by a specific invariant law that preserves the branch distinctions the scenery needs.

The Lean core earns exactly this grammar, not the geometry. It proves that instance openness, factor lawfulness, law-level closure by a supplied convergence relation, and the one-point and universal-closure calibrations fit together coherently. It proves that the shape can be uninformative relative to a proposed inquiry. And it proves that the shape need not be wholly non-retentive: a minimal latch reading law-closes the same escaping orbit while separating the initial state from every later one. That witness is thin, but it sits on the retentive side of the one-point calibration; it compresses almost everything and still keeps a distinction alive. The geometry enters when V, f, $\mathcal{P}(V)$, convergence, and the collapse condition are supplied by the actual scenery theory.

The collapse condition is now concrete enough to state without pretending it is already solved. A Dirac law may correspond to smooth closure: all mass has fallen onto one stable tangent view. A law supported on a finite orbit may correspond to orbital return rather than fractal non-closure. A one-point law corresponds to a reading that retains no distinction. The target fractal case requires a nontrivial invariant scenery law: closed at the law level, not closed as one view, not closed as mere finite return, and not reduced to a one-point reading.

This is also where the naming analogy becomes useful. Lawfulness is the commutation equation. It says the law-level reading moves correctly with magnification. Faithfulness, in the global naming sense, would be too strong: it would require the reading to lose nothing, and then the law would simply copy the instance. What survives here is a weaker and more exact analogue: domain-relative faithfulness. The law may compress the orbit, but it must not compress away the feature for which the orbit was fractal rather than smooth, periodic, or empty.

So the remaining work is not to invent another abstract grade. It is to follow this square through particular fractal landscapes and ask, in each one, which closure is not attained, which law closes, and which collapse condition the law satisfies.

Notes

1. The source anchors for this paragraph are now citation-level, but the paper still uses them only for orientation unless a later pass works through the technical details. Mandelbrot's early scale argument is Benoit Mandelbrot, "How Long Is the Coast of Britain? Statistical Self-Similarity and Fractional Dimension," *Science* 156, no. 3775 (1967): 636–638, doi:10.1126/science.156.3775.636. For the refusal of a single clean definition and the standard mathematical toolkit, see Kenneth Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, 2nd ed. (Wiley, 2003), doi:10.1002/0470013850. Tangent measures and rectifiability are anchored by David Preiss, "Geometry of Measures in R^n: Distribution, Rectifiability, and Densities," *Annals of Mathematics* 125, no. 3 (1987): 537–643, doi:10.2307/1971410. Beta-number/flatness tests are anchored by Peter W. Jones, "Rectifiable Sets and the Traveling Salesman Problem," *Inventiones Mathematicae* 102, no. 1 (1990): 1–15, doi:10.1007/BF01233418. Scenery-flow sources include Tim Bedford and Albert M. Fisher, "On the Magnification of Cantor Sets and Their Limit Models," *Monatshefte für Mathematik* 121 (1996): 11–40, doi:10.1007/BF01299636; Bedford and Fisher, "Ratio Geometry, Rigidity and the Scenery Process for Hyperbolic Cantor Sets," *Ergodic Theory and Dynamical Systems* 17, no. 3 (1997): 531–564, doi:10.1017/S0143385797079194; Hillel Furstenberg, "Ergodic Fractal Measures and Dimension Conservation," *Ergodic Theory and Dynamical Systems* 28, no. 2 (2008): 405–422, doi:10.1017/S0143385708000084; and Michael Hochman, "Dynamics on Fractals and Fractal Distributions," arXiv:1008.3731.

2. The formal core is machine-checked and public: Lean 4 formalization, https://github.com/Larsen-Close/fractal_non_closure, commit `2f7637a`, Lean toolchain v4.31.0, no Mathlib dependency, no `sorry` and no added axioms. The results named in this paper correspond to `Factor.OpenWithLawClosure`, `TerminalReading.onePointCalibration`, `TerminalReading.lawful_closed_collapsing_of_distinct`, `Factor.lawClosed_of_universal_closesTo`, `NatSucc.onePoint_factor_openWithLawClosure_zero`, and `LatchReading.openWithLawClosure_without_terminal_reading`.

3. These examples are domain-specific instantiation targets, not an identification of the same invariant across different fields. The statistical-complexity example should be read through James P. Crutchfield and Karl Young, "Inferring Statistical Complexity," *Physical Review Letters* 63, no. 2 (1989): 105–108, doi:10.1103/PhysRevLett.63.105. The rectifiability and scenery-flow examples are anchored by the Preiss, Jones, Bedford-Fisher, Furstenberg, and Hochman sources in the preceding note. The shared form is the indexed task: specify the operation, the unattained target of closure, the law-level target, and the nondegeneracy condition.