A Section of the Chaotic Bundle: A Category-Theoretic Motif for Mandelbrot–Julia Art

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

December 27, 2025

This paper proposes a new way to make fractal art that is disciplined by structure rather than driven by ornament. In the classical quadratic family $f_c(z) = z^2 + c$, the Mandelbrot set organizes parameter values c, while each Julia set J_c expresses the dynamical "fiber" over that parameter. Most visualizations emphasize escape-time coloring or direct rendering of Julia geometry. Here we instead treat the parameter plane as a base space and map each parameter to a compact behavioral signature derived from the critical orbit. This yields a functor-like assignment $c \mapsto \sigma(c)$ into a space of invariants (boundedness, escape rate, expansion proxies, symmetry cues). The artwork is then constructed as a constrained "section": marks are drawn only where local coherence is respected, and silence appears where coherence fails. To make multiple scales legible on a single static canvas, we introduce a radial parameterization that encodes scale by radius (coarse near the center, fine near the boundary), allowing walking-distance inspection to replace zoom interfaces. We show how color can be treated as a second functor into a perceptual space, producing a natural-transformation motif that makes structure visible without reducing the work to a literal fractal plot. The result is a category-theoretic aesthetic of coherence, refusal, and legible failure.

Keywords: category theory, Mandelbrot set, Julia set, quadratic dynamics, parameter space, fiber bundle, section, Grothendieck construction, functoriality, natural transformation, coherence constraint, critical orbit, escape rate, Lyapunov proxy, generative art, data-driven aesthetics, multiscale visualization, radial parameterization, perceptual encoding, silence as structure.

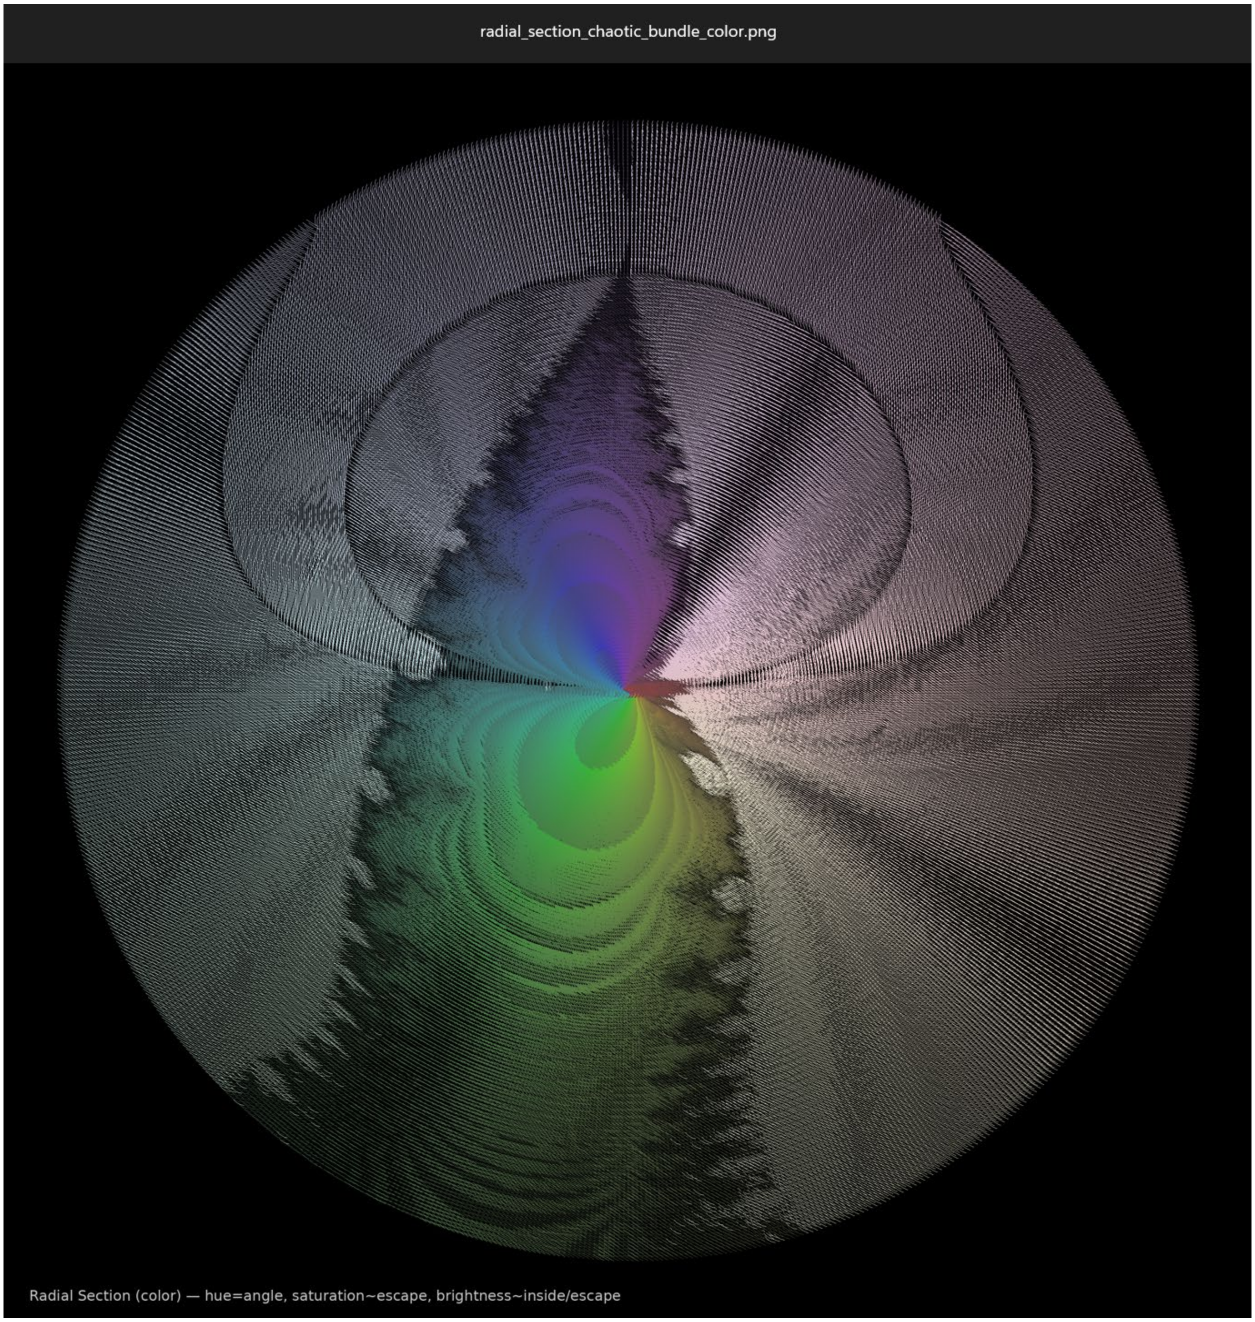

Figure 1. A Section of the Chaotic Bundle: A Category-Theoretic Motif for Mandelbrot–Julia Art

Outline

1. Introduction: From Fractal Depiction to Structural Motif

The Mandelbrot and Julia sets occupy a peculiar place in modern visual culture. They are among the few mathematical objects that have become widely recognizable as images, even to viewers who do not know their definitions. This is a compliment to the mathematics, but it also creates a trap: the public “knows” fractals primarily as a style. Once an object becomes a style, it invites a predictable set of artistic reflexes: saturate the palette, maximize detail, push contrast, and let the viewer’s awe do the explanatory work. The result is often visually compelling, but conceptually thin. The picture becomes a banner for “complexity,” a poster for “infinity,” or a generic emblem of “chaos,” while the structural relationship between parameter space and dynamical behavior—the reason these objects are profound—is rarely allowed to shape the composition.

This paper proposes a different approach. Instead of treating the Mandelbrot and Julia sets as surfaces to be painted, we treat them as a disciplined relationship: a family of systems indexed by a base space, with fibers that vary in ways that are sometimes coherent, sometimes violently unstable. The artwork is not a direct depiction of either set. It is an image produced by insisting that the same rule govern the entire construction, and by allowing the work to exhibit refusal—silence, sparsity, discontinuity—when the rule cannot be satisfied without lying. In other words, the art is built from a coherence principle, not from a craving for maximal detail.

This reframing is not merely philosophical. It changes what counts as a legitimate aesthetic move. It treats “more pixels” as an easy substitute for meaning, and it treats “staying faithful to structure” as the primary generator of form. It also shifts how viewers are meant to engage the work. A fractal poster is consumed instantly: the eye is pulled into detail and the mind supplies a vague metaphysical story. A coherence-driven work is different. It is meant to be read over distance and time. It rewards walking up close, stepping back, and noticing that some features are present precisely because other features were disallowed. The viewer’s movement becomes part of the interpretive apparatus, much as a proof is not “seen” in a single glance but traversed.

1.1 The usual visual culture of Mandelbrot and Julia images

The canonical “Mandelbrot image” is a parameter-plane picture colored by escape time or a smooth potential, sometimes enhanced with distance estimation or orbit traps. It is already a kind of visualization triumph: a single plot can encode membership in the Mandelbrot set (boundedness of the critical orbit) and, indirectly, a classification of qualitative dynamical regimes. Yet in popular and even semi-technical art practice, this classification role is often overshadowed by the aesthetics of zooming. The Mandelbrot set becomes an endless mine of ornate filigree. The narrative is: “Look how deep it goes.” The method is: sample, color, and intensify until the image glows.

Julia sets are treated even more directly as “shapes.” They are rendered as boundaries in the dynamical plane, sometimes with interior/exterior coloring, sometimes with orbit traps, sometimes as flame-like distance fields. Viewers learn to recognize a Julia set as a kind of emblem: dendrites, dusts, rabbit-like blobs, spirals. The parameter $c$becomes a knob for changing the emblem. The usual artistic stance is to pick a few attractive parameters and render them at high resolution. The relationship between Julia sets and the Mandelbrot set is sometimes stated—“inside gives connected Julia sets, outside gives dust”—but rarely used as a compositional constraint. It is treated as trivia rather than as a governing law.

This aesthetic culture is not “wrong.” It has produced genuine beauty and an enormous archive of techniques. But it tends to optimize for immediate impact: high saturation, high detail density, and a promise of inexhaustible novelty. That optimization pushes artists toward a single question: “How can I get more visual yield per pixel?” It does not push toward the question that makes the Mandelbrot–Julia relationship profound: “How does local behavior change under small parameter change, and what does it mean when that dependence is stable versus unstable?” In short, it celebrates output while ignoring coherence.

A further limitation is that most fractal images are implicitly single-scale. Even if an image contains small and large features, the viewer is forced into one magnification at a time. The standard solution is an interactive zoom: the work becomes a video or a navigable interface. But if the goal is a static artwork—something that can be printed, hung, and physically approached—then the culture’s default techniques provide little guidance. They are optimized for screens and spectacle, not for embodied viewing.

1.2 The central claim: art as coherence-preserving assignment

The central claim of this paper is that the Mandelbrot–Julia landscape can be used to produce art in a way that is structurally disciplined by a coherence principle: the artwork is built from an assignment that must behave “functorially enough” under small changes, and it is permitted to refuse representation where that coherence fails.

Concretely, we treat the parameter plane (the c-plane) as a base space. For each parameter c, we compute a behavioral signature $\sigma(c)$ derived from the critical orbit of $f_c(z) = z^2 + c$. The signature is intentionally finite and modest: it does not attempt to reproduce the full Julia set, nor does it attempt to outdo standard renderings with clever shading. It extracts a small set of invariants or proxies—boundedness, escape rate, expansion cues, symmetry hints—that are stable enough to be meaningful but sensitive enough to reveal transitions. The artwork then places a mark on the canvas for each sampled parameter, with mark geometry and color driven strictly by $\sigma(c)$.

The category-theoretic flavor enters at the level of constraint. We impose a coherence gate: along any path of nearby parameters (a local morphism in the base), the corresponding signatures are expected to vary gently when the system is in a stable regime (in particular, within the connected/"inside Mandelbrot" domain where the Julia set is connected). If the signatures jump too sharply for too small a parameter change, we do not "smooth it away" with a visual trick. We do not interpret the jump as an invitation to decorate. We treat it as a failure of representability under our chosen signature map, and we allow the output to become sparse or silent. In the resulting image, absence is not an error; it is an encoded statement: "coherence broke here, at this resolution, under this signature."

This is the key aesthetic reversal. Standard fractal art treats instability as a source of ornament: the boundary is where the most elaborate detail appears. In a coherence-preserving motif, instability is not primarily ornamental; it is a boundary of disciplined representation. The viewer sees not just "complexity," but the trace of a rule straining against a system that does not always reward smooth dependence.

A second part of the claim is that a single static canvas can carry multiple scales without zoom interfaces by encoding scale geometrically. In the radial construction developed here, distance from the canvas center corresponds to the size of the neighborhood sampled in parameter space. The center of the image explores a broader region around a chosen basepoint c_0, while the outer regions explore progressively smaller neighborhoods. This turns "magnification" into a compositional coordinate. A viewer does not need a slider or a scroll wheel; the viewer walks. The artwork becomes legible at multiple scales through embodied motion rather than interface control.

Finally, color is treated not as a garnish but as a second structured assignment. The paper frames perceptual encoding as a mapping from behavioral signatures into a color space (for example HSV), organized so that changes in parameter can correspond to coherent changes in hue, saturation, and brightness. This can be described as a second functor-like map and, in favorable cases, as a natural-transformation motif: behavior and color change in step, and the viewer experiences color as structure.

1.3 Contributions and what the paper is not claiming

The contributions of the paper are practical and conceptual.

First, it introduces a disciplined generative-art framework: a minimal set of orbit-derived measures used to define a behavioral signature, coupled to a coherence constraint that suppresses marks where local continuity is violated. This framework produces images that are recognizably rooted in the Mandelbrot–Julia relationship while refusing the familiar look of escape-time fractal posters. The framework is portable: one can change the signature, the sampling geometry, or the mark primitives without abandoning the coherence principle.

Second, it introduces a static multiscale motif designed for physical viewing. The radial parameterization is not presented as the only solution, but as a concrete example of how to build multiscale legibility into the geometry of the composition. This is meant to support print-scale work, installation thinking, and walk-up inspection.

Third, it proposes a category-theoretic vocabulary that clarifies what is being done without forcing the mathematics into a purely categorical straitjacket. The "base space / fibers / section" language is not a metaphor glued onto an image after the fact. It is a design description: the artwork is constructed as if it were a partial section of a behavior bundle over parameter space. The coherence gate corresponds to a demand for local consistency of the assignment under small morphisms in the base. This framing makes it possible to explain why the work contains silence, why it refuses to fill every region with texture, and why "failure" can be meaningful.

Just as important are the non-claims.

This paper is not claiming a new theorem about Mandelbrot or Julia sets. It does not attempt to compete with the deep mathematical literature on parameter space structure, bifurcation theory, renormalization, or the fine topology of Julia sets. The categorical language is used to articulate a design discipline; it is not presented as a formal reconstruction of complex dynamics inside category theory.

It is also not claiming that the specific numerical proxies used here (escape rate, Lyapunov-like average along the critical orbit, and simple symmetry cues) are canonical invariants. They are deliberately practical and lightweight. The goal is to show that even modest signatures, if treated with coherence constraints, can yield work that is structurally legible. More sophisticated invariants could replace them, but the aesthetic principle would remain the same.

Finally, it is not claiming that earlier fractal art is conceptually empty or aesthetically inferior. The point is not to shame spectacle but to open a different lane: a lane where restraint, refusal, and coherence are first-class aesthetic primitives. The argument is that this lane is underexplored, and that category theory offers a vocabulary—base, fiber, section, coherence, naturality—that helps artists build in constraints that make the resulting images feel inevitable

2. Background: The Quadratic Family as a Base-and-Fiber Story

The conceptual move behind this paper is easiest to motivate in the simplest setting where the relationship between parameter space and dynamical behavior is unusually clean: the quadratic family $f_c(z) = z^2 + c$. This family is mathematically rich enough to exhibit the entire vocabulary of stability, bifurcation, and chaotic boundary behavior, yet constrained enough that one can tell a coherent story without immediately invoking heavy machinery. It is also the setting in which the Mandelbrot set emerges as a true classifier: it is not just a decorative silhouette, but a boundary between qualitatively different worlds, and it does so via a single orbit—the orbit of the critical point. That single-orbit fact is precisely what makes a "base-and-fiber" interpretation feel natural: the parameter $c$labels the system, and the associated Julia set is an object that can be viewed as living over c, changing as $c$changes. One does not need to force the language of bundles onto the situation; it is already present in how complex dynamics is practiced and explained.

2.1 The quadratic map $f_c(z) = z^2 + c$ and the critical orbit

Fix a complex parameter $c \in \mathbb{C}$, and consider the iteration of the map

$$f_c(z) = z^2 + c.$$

The dynamical plane is the complex z-plane, and the central operation is iteration:

$$z_{n+1} = f_c(z_n) = z_n^2 + c, n = 0,1,2, \dots$$

For a given starting point z_0, the sequence $\{z_n\}$ is its orbit. In complex dynamics, the qualitative behavior of orbits—whether they remain bounded, converge to cycles, wander chaotically, or escape to infinity—organizes the plane into regions of stable and unstable behavior. The Julia set is the boundary where this organization becomes maximally sensitive: arbitrarily small changes in the initial condition can produce qualitatively different outcomes.

Within the quadratic family, one point has privileged status: the critical point of f_c. The derivative is

$$f_c'(z) = 2z,$$

so the unique critical point is $z = 0$. The orbit of this critical point,

$$0,\ c,\ c^2 + c,\ (c^2 + c)^2 + c,\ \dots,$$

is called the critical orbit. This orbit is not merely one orbit among many. For polynomials, and particularly for the quadratic family, the fate of the critical orbit governs the global structure of the dynamics. Intuitively, the critical point is where the map locally “folds” the plane most strongly; if that fold is trapped in a bounded region, the dynamics has a chance to organize into a connected, coherent structure. If that fold escapes, the dynamics tends to fragment. In more formal terms, many qualitative properties of the Julia set and the existence of attracting cycles are controlled by the critical orbit.

One of the most important practical facts for computation is the escape criterion. For $f_c(z) = z^2 + c$, if an orbit ever satisfies $| z_n |> R$ for a sufficiently large radius R, then the orbit will diverge to infinity. In practice, $R = 2$ already suffices for the quadratic family; larger choices such as $R = 4$are often used for numerical stability. This means one can classify a parameter $c$by iterating $z_{n+1} = z_n^2 + c$ from $z_0 = 0$ and asking whether $| z_n |$exceeds the escape radius before a maximum iteration count. That classification is the computational backbone of the Mandelbrot set.

This critical-orbit centrality is also the backbone of the art framework developed later. It supplies a compact source of signatures. Rather than rendering an entire Julia set, we can compute orbit-derived scalars that serve as behavioral proxies: boundedness versus escape, the rate at which escape happens, and measures of local expansion along the orbit. The fact that one orbit carries so much structural information is what makes this feasible as a compositional rule rather than a hack.

2.2 Mandelbrot set as parameter classifier and connectedness criterion

The Mandelbrot set is defined in the parameter plane:

$$M = \{\, c \in \mathbb{C} \;\mid\; \text{the orbit of } 0 \text{ under } f_c \text{ remains bounded} \,\}.$$

This definition is striking because it turns a question about the dynamical plane into a classifier in parameter space. Each $c$specifies a dynamical system; the Mandelbrot set is the set of systems whose critical orbit does not escape.

There are two aspects of this definition that are essential for our later “base-and-fiber” interpretation.

First, it is a classifier with real bite. Inside M, the family contains parameters with stable attracting cycles of various periods, and these organize into the bulb-like components attached to the main cardioid. Outside M, the critical orbit escapes and the system behaves, in a global sense, as an “escaping” regime. The boundary is where small parameter changes can flip between bounded and escaping behavior, producing the visual richness that the Mandelbrot set is famous for. But in the conceptual view of this paper, the boundary is not celebrated for its ornament. It is the locus where coherence becomes fragile, where any attempt to build a smooth assignment from parameter space to behavioral signatures will be tested.

Second, and even more important, the Mandelbrot set encodes a connectedness criterion for Julia sets in this family. A classical theorem in complex dynamics states that for $f_c(z) = z^2 + c$,

$$c \in M \Longleftrightarrow J_c \text{ is connected.}$$

Equivalently,

$$c \notin M \Longleftrightarrow J_c \text{ is totally disconnected (a “dust”).}$$

This is a profound statement because it turns a global topological property of a geometric object in the dynamical plane—the connectedness of the Julia set—into a parameter-space membership test determined by a single orbit. It also means that the Mandelbrot set is not merely “the set of interesting parameters.” It is a boundary between topological regimes of the fibers J_c.

For the purposes of this paper, this equivalence is the key justification for treating the Mandelbrot set as the “good base region” for coherent fiber variation. If we restrict attention to parameters within M, then the associated Julia sets are connected, and many global features vary in a way that can be tracked meaningfully. Once we move outside, the fibers become dust-like, and the relationship between small parameter perturbations and visible structure becomes far more violent. This does not mean the outside region is uninteresting; it means that any disciplined visual assignment must decide what it will preserve and what it will refuse to pretend to preserve.

In our generative motif, “inside Mandelbrot” plays exactly this role. It is the regime in which we demand local coherence from our assignment $c \mapsto \sigma(c)$. That demand is not a theorem; it is a design stance. But it is a stance aligned with the mathematics: the connectedness regime is where fiber-thinking is most natural, because you are not tracking a shattered Cantor dust that can rearrange under tiny perturbations in ways that overwhelm any finite signature.

2.3 Julia sets as fibers and why “fiber-thinking” is natural here

For each parameter c, the Julia set J_c is defined as the boundary between stable and unstable behavior in the dynamical plane. There are multiple equivalent definitions; one intuitive description is that J_c is the set of points where the iterates of f_c behave chaotically in the sense of sensitive dependence. Another is that J_c is the boundary of the filled Julia set K_c, where

$$K_c = \{\, z \in \mathbb{C} \;\mid\; \{f_c^{\circ n}(z)\} \text{ remains bounded} \,\}, J_c = \partial K_c.$$

Thus, K_c captures the non-escaping points, and J_c is its boundary.

The key idea for this paper is not to re-prove properties of Julia sets, but to adopt a point of view: as $c$varies, the objects J_c vary. The parameter plane is a space of systems; the Julia set is a characteristic object associated to each system. In that sense, the family

$$c \mapsto J_c$$

is already a "bundle-like" construction: you have a base (parameters) and a fiber (a geometric object) attached to each base point.

Even without formal bundle theory, the language of fibers is natural for at least four reasons.

First, the Mandelbrot set acts as an organizing chart for the Julia fibers. When one navigates the Mandelbrot set and selects a point c, one is effectively selecting a fiber J_c. The cultural practice of exploring Julia sets by "picking points in the Mandelbrot set" already treats the Mandelbrot set as an index of fibers. In other words, popular fractal exploration is already fiber-based; it simply doesn't name it that way.

Second, the dependence of Julia sets on parameters is central to the subject. Questions like "how does J_c change under perturbation of c?" and "what happens to the geometry/topology of J_c near the boundary of M?" are standard. These are fiber-variation questions. The instability at the boundary is precisely a statement about the failure of naive continuity of fibers. That is exactly the kind of phenomenon that categorical thinking—coherence, naturality, controlled failure—helps articulate.

Third, the connectedness dichotomy makes the fiber story unusually crisp. Inside M, the Julia set is connected, so fiber variation can be tracked with a sense of continuity. Outside M, the fiber becomes totally disconnected, and "tracking" becomes more delicate: small parameter changes can cause massive rearrangements in dust structure, and any finite visual signature risks becoming misleading. This dichotomy suggests an aesthetic discipline: demand coherence where coherence is meaningful, and allow silence or refusal where coherence is not preserved under the chosen representation.

Fourth, from a generative-art perspective, fiber-thinking provides a design ladder. One can choose what the fiber object is. It does not have to be the full Julia set geometry. It can be a derived object: a signature vector of invariants, a topological summary, a measure distribution, or a perceptual encoding. The key is that it is attached to each parameter in a consistent way. Once you accept that "fiber" can mean "derived behavior object," you are no longer trapped in the usual binary of "render Julia" versus "render Mandelbrot." You can build a new visual language that is still faithful to the underlying family.

This is the guiding stance for the rest of the paper. We will treat the quadratic family as a base-and-fiber system, define a small but meaningful fiber signature $\sigma(c)$ from critical-orbit behavior, and then construct images as partial sections of this derived bundle—images in which coherence is not merely hoped for, but enforced, and in which failure of coherence is allowed to remain visible as a first-class feature rather than being smoothed into decoration.

3. The Category-Theoretic Frame

The phrase "category theory" can mean many things, ranging from a rigorous formalism used to prove theorems, to a disciplined habit of describing systems in terms of objects, morphisms, and coherence rather than in terms of isolated constructions. This paper uses category theory primarily in the second sense. The goal is not to rebuild complex dynamics inside a categorical axiom system. The goal is to extract a clean design grammar that makes the construction of the artwork intelligible, reproducible, and principled. Category theory provides a vocabulary for that grammar: base objects, fibers, assignments, morphisms as controlled change, functorial behavior as structure-preserving change, and coherence as the policing mechanism that decides when an assignment is "honest."

In this section we specify a workable categorical model of what the artwork is doing. The model is deliberately modest. It is tuned to the needs of generative construction: to define what counts as a "small change in parameter," what counts as a "small change in behavior," how one assigns a behavior object to each parameter, and how one enforces the rule that this assignment should not lie about continuity. The work then appears as a partial section of a derived bundle: a drawing that exists precisely where coherence can be maintained, and that becomes silent where coherence fails.

3.1 The base category of parameters: proximity, paths, and deformations

The base space in this story is the parameter plane $\mathbb{C}$, or more precisely a region of $\mathbb{C}$that we choose to sample (often a neighborhood around a chosen basepoint c_0). The simplest "categorical" way to treat $\mathbb{C}$is to ignore its full structure and keep only the notion of local change: when are two parameters "near," and what does it mean to deform one into another?

There are multiple ways to package this as a category, and the paper does not require choosing a single canonical one. What matters is that morphisms represent controlled parameter change.

A convenient option is a thin category built from an adjacency relation. Fix a scale $\varepsilon > 0$. Define a directed edge (a morphism)

$$c \to c'$$

whenever $| c - c' | < \varepsilon$. If we also include identities $c \to c$and take the transitive closure, we obtain a preorder-like structure that encodes the idea of "reachable by a chain of small steps." In this base category, there may be many paths from $c$to c', representing different sequences of small perturbations.

A second, more geometric option is to use a path category. Objects are parameters c. Morphisms $c \to c'$ are continuous paths $\gamma: [0,1] \to \mathbb{C}$ with $\gamma(0) = c$ and $\gamma(1) = c'$, with composition given by path concatenation. This better reflects the intuition that one can "move through parameter space" in many ways, and that coherence should be tested along these movements.

A third option, often enough for computation, is to use an open-set viewpoint: objects are small neighborhoods (disks) in $\mathbb{C}$, and morphisms are inclusions. This turns parameter space into a site for local restriction: moving to a smaller neighborhood is a morphism. This is particularly natural when we later encode scale by radius, because the canvas radius is literally a restriction to smaller and smaller neighborhoods around c_0.

All three models share the same operational meaning for our purposes: the base category is not "the set of all parameters." It is the structured arena in which we decide what counts as "local movement" and therefore what counts as a legitimate test of continuity. The image is then a record of where such local movement can be reflected faithfully in the behavior-space.

3.2 The fiber category of behavioral signatures

In classical complex dynamics, one might attach to each parameter $c$the Julia set $J_c$itself, or the filled Julia set K_c, or a dynamical system object such as $(\mathbb{C}, f_c)$. But for our purpose—making a static image whose marks are controlled by a finite data vector—we attach something smaller: a behavioral signature extracted from the critical orbit and related local proxies.

Define a signature map

$$\sigma(c) = (\sigma_1(c), \sigma_2(c), \dots, \sigma_k(c)),$$

where each component is a computable quantity such as:

- a boundedness indicator (did the critical orbit escape by time N?),
- a smooth escape-rate proxy (how quickly does it escape, if it escapes?),
- an expansion proxy (a Lyapunov-like average of $\log | 2z_n |$along the critical orbit),

- a symmetry cue (e.g., whether $\mathrm{Im}(c)$ is near zero),
- optional additional components (period estimates, orbit-trap statistics, distance estimators).

The crucial point is that $\sigma(c)$ is not an arbitrary feature vector. It is chosen to represent "behavior" in a way that is plausible to preserve under small parameter perturbations in stable regimes, and plausibly allowed to fail near bifurcations and boundaries.

To frame this categorically, we build a fiber category $\mathcal{D}$ whose objects are such signatures. Morphisms represent admissible changes between signatures: for example, one might say there is a morphism $\sigma \to \sigma'$when $\| \sigma - \sigma' \|$ is below a tolerance. This yields a thin "proximity" category on signature space. Alternatively, one can structure $\mathcal{D}$as a metric-enriched category: objects are signatures, and the "strength" or "cost" of a morphism is the distance between them. For the artwork, we only need the induced idea: some signature transitions are small and coherent, others are large and represent a breakdown of local representability under the chosen signature.

In this fiber category, the act of drawing is an act of representation: each $\sigma(c)$ is mapped to a glyph (a stroke shape, thickness, curvature) and optionally to color. The fiber category is therefore not purely abstract; it is the semantic layer that ties computation to mark-making.

3.3 The functor-like assignment $F: c \mapsto \sigma(c)$

With a base category $\mathcal{C}$of parameters and a fiber category $\mathcal{D}$of signatures, the core mathematical object is an assignment

$$F: \mathcal{C} \longrightarrow \mathcal{D}$$

defined on objects by

$$F(c) = \sigma(c).$$

To call Fa functor in the strict sense, we would also need to specify how morphisms in $\mathcal{C}$map to morphisms in $\mathcal{D}$, and verify identity and composition laws. In practice, the artwork does not require strict functoriality; it requires a weaker property: that local parameter morphisms should be taken to "small" signature morphisms in the stable regime, and that this expectation should be policed by a coherence gate. That is why we call the assignment functor-like: it behaves like a functor where the system supports it, and it is permitted to fail where the system does not.

In the adjacency-category model of $\mathcal{C}$, a morphism $c \to c'$ corresponds to $| c - c' | < \varepsilon$. The induced "candidate morphism" in $\mathcal{D}$ is the change $\sigma(c) \to \sigma(c')$. If the signature distance is small, we treat it as a valid morphism. If it is too large, we treat it as a failure of functorial behavior at that resolution.

In the path-category model, a morphism is a path γ. The image of that path under $F$is the trace of signatures $\sigma(\gamma(t))$. If this trace behaves smoothly under small t-changes, we regard it as coherent. If it jumps or becomes highly unstable, we interpret that as a failure of representing the fiber variation with our chosen signature at that scale.

In the inclusion-of-neighborhoods model, the morphisms are restrictions to smaller neighborhoods. Then F is naturally interpreted as "behavior under restriction": as you shrink the neighborhood around c_0, does the signature field become more predictable, more faithful, and more coherent? The radial layout used later is essentially a visualization of this restriction process.

Thus, the functor-like assignment is not merely an interpretation layer. It is the backbone of the generative process. The image can be read as a diagram of where Fbehaves like a structure-preserving map and where it does not.

3.4 Coherence as a rule: local faithfulness and admissible variation

The most important aesthetic move in this framework is the coherence rule. This rule declares that we will not draw "everything we can compute." We will draw only those features that pass a test of local consistency.

There are two related notions here: admissible variation and (local) faithfulness.

Admissible variation is the threshold idea. If $c$and $c'$are close in parameter space, then $\sigma(c)$ and $\sigma(c')$ should be close in signature space, at least within the regime where such an expectation is meaningful. We formalize this by choosing tolerances:

- a parameter tolerance εthat defines "near,"
- a signature tolerance δthat defines "coherently near."

Then the coherence condition is:

$$| c - c' | < \varepsilon \text{ and } c, c' \in \text{coherence regime} \Longrightarrow \| \sigma(c) - \sigma(c') \| < \delta.$$

Here the "coherence regime" can be chosen as the set of parameters inside M, or a subset thereof, depending on how strict we want to be.

Local faithfulness is subtler. In category theory, a functor is faithful if it does not collapse distinct morphisms. We are not enforcing faithfulness globally; we are using a local analogue: in regions where the dynamics is stable, small changes in parameter should not produce wildly different signatures. If they do, then our signature map is not acting as a faithful witness of local parameter change. In that case, the correct response—if we want the work to be honest—is not to smooth the signatures or to average them until they look continuous. It is to refuse to draw, because drawing would imply a continuity we do not possess.

This coherence rule is also where the "category theory as ethics" aspect appears. A visualization can always be made prettier by interpolation, smoothing, or forced continuity. But those are aesthetic decisions that smuggle in false claims. The coherence constraint functions as an explicit anti-smuggling rule. It says: the picture is allowed to be sparse, because sparseness is preferable to lying.

In the images generated for this paper, the coherence rule is implemented as a gate: when successive parameter samples are close and appear to be in the stable regime, but the signature jump exceeds the threshold, the mark is suppressed. The result is that boundaries and unstable regions manifest as thinning, tearing, or silence. These are not glitches. They are categorical signals: locations where "functorial behavior" fails under the chosen signature and sampling scale.

3.5 Partial sections: why refusal and silence are mathematically honest

Once we have a base category and a fiber category with an assignment F, a natural next step is to treat the situation as a bundle-like structure. One formal route is the Grothendieck construction: from a functor $F: \mathcal{C} \to \mathcal{D}$ one can build a category $\int F$ whose objects are pairs $(c, x)$with $x \in F(c)$, and whose projection to $\mathcal{C}$resembles a fibration. In our setting, we can interpret $\int F$ informally as the "total space" of behavior over parameters: a space of parameter–signature pairs $(c, \sigma(c))$.

A section of such a projection is a choice of a fiber element over each base point. In our derived setting, the signature $\sigma(c)$ is already the chosen fiber element, so in a formal sense we already have a section. But the artwork is not merely plotting this section. It is plotting it subject to coherence constraints and finite representation. That is why the relevant concept is a partial section: a section that is only drawn where it can be represented coherently under the rules.

Mathematically, partial sections are not exotic. Many geometric objects are locally trivial but globally obstructed; one may have local sections that do not glue to global ones. In sheaf language, one has local data that fails a gluing condition. In fiber bundle language, one may be able to choose a section on a subset but not extend it continuously across a singular locus. The same logic applies here, with “continuous extension” replaced by “coherent representability under a finite signature and a drawing grammar.”

In the context of complex dynamics, the boundary of the Mandelbrot set is a natural singular locus for many naive parameter-to-object assignments. Small parameter changes can produce large changes in qualitative behavior. Any attempt to maintain smooth representation across such loci risks introducing a false continuity. The partial-section stance says: we do not hide this. We admit it as part of the phenomenon.

Silence, then, is not merely an artistic flourish. It is the visible trace of a constraint: “the chosen signature and sampling scale do not support a coherent section here.” The gaps become evidence of instability rather than defects of rendering. They also serve an aesthetic function aligned with the mathematics: they prevent the image from collapsing into indiscriminate texture. They create a legible tension between regions of coherence (dense, stable marks) and regions of failure (thin, broken, or absent marks). The viewer is invited to read this tension as structure: where the system grants reliable continuity and where it refuses it.

This is the heart of the category-theoretic frame as applied to art. It is not that the Mandelbrot set “is category theory.” It is that category theory supplies a language for disciplined mapping between domains, and a vocabulary for honest failure: functorial behavior where possible, coherence checks everywhere, and partial sections where global continuity would be a lie.

4. Constructing the Artwork as a Partial Section

This section turns the category-theoretic framing into an explicit generative procedure. The main idea is to build an image that is not “a plot of a set,” but a record of where a chosen behavioral assignment can be represented coherently. To do that we specify four ingredients: a signature $\sigma(c)$ attached to each parameter c, a coherence gate that decides whether a mark is allowed, a small vocabulary of marks (glyphs) that carry the signature into geometry and color, and an interpretation of discontinuities that treats them as meaningful rather than as rendering defects.

The design intent is restraint. The signature is intentionally small; it is not meant to encode everything. The coherence gate is intentionally strict; it is meant to protect the meaning of “local continuity.” The glyph set is intentionally minimal; it is meant to make the image legible rather than decorative. And discontinuities are intentionally preserved; they are meant to be read as the locations where the system refuses smooth dependence under the chosen representation.

4.1 Defining the signature $\sigma(c)$: boundedness, escape, expansion, symmetry

Fix a parameter $c \in \mathbb{C}$ and the quadratic map $f_c(z) = z^2 + c$. Let the critical orbit be defined by

$$z_0 = 0, z_{n+1} = z_n^2 + c.$$

Choose a maximum iteration count N and an escape radius R(in principle $R = 2$ suffices for the quadratic family, but larger R is often used in practice to stabilize numerical thresholds). Define the escape time

$$\tau(c) = \min\{n \leq N : |z_n| > R\},$$

with the convention $\tau(c) = N$ if no escape occurs by time N.

The signature $\sigma(c)$ is a small vector of orbit-derived features. The specific components can be varied, but the core set used here is:

Boundedness (Mandelbrot membership proxy).
Define a binary indicator

$$b(c) = \begin{cases} 1 & \text{if } \tau(c) = N, \\ 0 & \text{if } \tau(c) < N. \end{cases}$$

This is a computational proxy for $c \in M$. It is not a proof of membership (finite iteration cannot certify boundedness for all time), but it is the standard operational classifier for rendering, and it is sufficient as a behavioral regime label for the art construction.

Escape measure (a smooth escape-rate proxy).
For escaping parameters, the raw escape time $\tau(c)$ is useful but coarse. A smoother quantity can be defined using the magnitude at escape. One typical proxy is a normalized "escape rate" of the form

$$e(c) \approx \frac{\log |z_{\tau(c)}|}{\tau(c)}.$$

This is not a canonical invariant; it is a practical scalar that increases when the orbit diverges rapidly. In the artwork, $e(c)$ is used as an "energy" parameter that can drive thickness, saturation, or brightness.

Expansion proxy (Lyapunov-like average along the critical orbit).
The derivative of f_c is $f_c'(z) = 2z$. Along the critical orbit we can accumulate a mean log-derivative magnitude:

$$\lambda(c) \approx \frac{1}{m} \sum_{k=0}^{m-1} \log \mid 2 z_k \mid,$$

where m is either N (if bounded by the cutoff) or $\tau(c)$ (if it escapes). This quantity functions as a crude proxy for local expansion along the orbit. It is not the full Lyapunov exponent of an invariant measure; it is a tractable scalar that distinguishes regimes of contraction versus expansion along the observed orbit segment. In mark design, $\lambda(c)$ can be mapped to curvature or roughness, allowing more "unstable" behavior to bend glyphs more strongly.

Symmetry cue (a categorical "stabilizer hint").
For real parameters ($\mathrm{Im}\,(c) = 0$), the system is invariant under complex conjugation, and its Julia set has a mirror symmetry. More generally, parameters near the real axis tend to exhibit stronger "mirror-like" cues in many renderings. We include a simple binary cue:

$$s(c) = \begin{cases} 1 & \text{if } \mid \mathrm{Im}\,(c) \mid < \eta, \\ 0 & \text{otherwise,} \end{cases}$$

for a small threshold η. This cue is intentionally simple; it is not claiming a deep symmetry classification. It is a way of letting the glyph vocabulary respond to a recognizable structural axis in parameter space.

Putting this together, a minimal signature takes the form

$$\sigma(c) = (b(c), e(c), \lambda(c), s(c)),$$

possibly with some components normalized into $[0, 1]$ ranges for stable mapping into mark parameters.

Two remarks matter for "honesty." First, this signature is deliberately incomplete. Many different parameters can share similar signatures. That is acceptable; it is part of what the image will show (regions of non-faithfulness). Second, the signature is numerically defined with finite iteration. This is also acceptable as long as the work does not pretend otherwise. Later, the coherence gate and the interpretation of discontinuities explicitly acknowledge that finite iteration introduces a controlled kind of epistemic humility: the section is built from what we can stably compute, not from what we might wish to assert.

4.2 The coherence gate: when marks are suppressed

Given a signature $\sigma(c)$, one might be tempted to draw a glyph at every sampled parameter. Doing so would produce a dense field, but it would also erase the main conceptual move of the paper: the insistence on coherence. The coherence gate is the mechanism that enforces that insistence.

At a high level, the rule is:

- If the base category says two sampled parameters are “near,” and
- we are in a regime where local coherence is expected to be meaningful (typically the “inside” regime $b(c) = 1$),
- then the signature should vary gently.
- If it does not, the mark is suppressed.

To implement this, we choose thresholds:

- A parameter proximity threshold ε.
- A signature-change threshold δ, measured by a norm in signature space, e.g.

$$\| \sigma(c) - \sigma(c') \|_2.$$

Then a typical gate condition between successive samples $c$and c'along a sampling path is:

$$\text{if } \mid c - c' \mid < \varepsilon \text{ and } b(c) = b(c') = 1, \text{ then require } \| \sigma(c) - \sigma(c') \| < \delta.$$

If the requirement fails, suppress the mark at c'(or suppress the transition segment, depending on the implementation style).

This does three important things.

First, it prevents “boundary sparkle” from becoming automatic ornament. Near the Mandelbrot boundary, small changes in $c$can flip boundedness classification at finite iteration depth, or change escape and expansion proxies abruptly. A naive renderer would turn this into a dense glitter field. The coherence gate refuses that. It is a refusal to treat instability as decoration.

Second, it makes the artwork encode a meaningful distinction: regions where the assignment behaves smoothly (dense, coherent texture) versus regions where it does not (thin, broken, silent). The image becomes an empirical diagram of local representability under the chosen signature map and sampling resolution.

Third, it makes discontinuity legible. If the work has no silence, the viewer cannot tell which features are stable consequences of the rule and which are artifacts of forced continuity. Silence becomes the negative space that allows the visible marks to claim meaning.

In categorical language, the coherence gate is enforcing a weak form of local functoriality: "small morphisms in the base should map to small morphisms in the fiber." It is not claiming this is always true. It is using the places where it fails as content.

4.3 Mark-making primitives: strokes as categorical "glyphs"

Once we have a signature $\sigma(c)$ and a coherence gate, we must choose how the signature becomes visible. The paper's stance is to use a minimal vocabulary of marks: simple strokes with controlled variations. This is both an aesthetic and an epistemic choice. A complicated glyph can always be made to look meaningful; a simple glyph forces meaning to come from structure in the data and from consistency across the field.

A stroke-based glyph can be described by a small set of parameters:

- **Location** on the canvas, determined by the sampling map from parameter space into the image plane.
- **Thickness**, a scalar controlling the stroke width.
- **Length**, a scalar controlling its spatial extent.
- **Curvature**, controlling whether the stroke is straight or bowed.
- **Orientation**, controlling the dominant direction of the stroke.
- **Color**, if used, encoded in hue/saturation/value or another perceptual coordinate system.

A typical mapping from signature components to glyph parameters is:

- Thickness $\propto e(c)$: faster escape means heavier stroke.
- Curvature $\propto \lambda(c)$: more expansion yields more bend or tension.
- Orientation depends on $s(c)$: near the real axis, rotate into a distinct orientation to signal the presence of a structural axis; away from it, use a different dominant direction.
- Brightness or saturation can depend on $b(c)$ and $e(c)$: bounded regimes may use a calmer baseline, while escaping regimes use more energetic intensity.

The point of calling these strokes "categorical glyphs" is not that the stroke is an arrow in a commutative diagram. The point is that the glyph is an element of a controlled representation functor: it is the image, under a rendering map, of the fiber object $\sigma(c)$. In other words, we can conceptually factor the construction as:

$$c \overset{F}{\mapsto} \sigma(c) \overset{R}{\mapsto} \text{glyph}(c),$$

where F is the behavioral assignment and R is a rendering map into mark space. The coherence gate then sits between them as a constraint that prevents $R \circ F$ from claiming continuity where F is not behaving continuously at the sampling resolution.

Because the glyphs are small and repeated thousands of times, the macroscopic image emerges by accumulation. Large-scale forms (plumes, annuli, spines) are not painted directly; they are emergent patterns produced by consistent local rules. This matches the conceptual ambition of the work: to let a global "structural motif" arise from local coherence constraints rather than from manual composition.

4.4 What discontinuities mean: boundary effects, instability, and numerical limits

A coherence-constrained rendering will almost inevitably produce discontinuities. Some will be genuine signals of dynamical instability; others will be artifacts of finite computation. The honesty of the framework depends on distinguishing these and on treating both as meaningful in different ways.

Boundary effects (mathematical instability).
Near the boundary of the Mandelbrot set, qualitative behavior changes abruptly with parameter. In the strict infinite-iteration idealization, this boundary is where the dynamics transitions between bounded and escaping critical orbits. In any practical finite iteration, the boundary becomes a region of ambiguity: points near the boundary can behave bounded for many iterations and then escape, or vice versa. In our signature language, this means that $b(c)$, $e(c)$, and $\lambda(c)$ can vary sharply under tiny changes in c. When the coherence gate suppresses marks in such regions, the resulting gaps are telling the viewer: "local continuity failed under this signature and resolution." That is a faithful report of the underlying phenomenon: the boundary is a locus of fragile dependence.

Instability of the chosen signature (representation limits).
Even away from the true boundary, a signature may be insufficiently robust. For example, the expansion proxy $\lambda(c)$ computed along a short orbit segment may fluctuate due to transient dynamics. Two parameters may be genuinely close in behavior, yet their finite-segment proxies differ enough to trip the gate. This is not a failure of the concept; it is a reminder that the signature is an abstraction. In category-theoretic terms, the functor-like assignment is not guaranteed to preserve all relevant structure; it preserves only what we chose to measure.

Numerical limits (finite iteration as epistemic humility).
The most important practical limit is that boundedness is decided by a cutoff N. A parameter may be classified as bounded at depth N but would escape at depth $N' \gg N$. This effect is especially common near the boundary. The coherence gate transforms this limit into a visible feature: where classification and proxies are unstable under small perturbations, marks become sparse. Rather than pretending that finite computation yields certainty, the image reveals where the computation cannot support smooth representation.

This is why discontinuity is not a defect in this framework. It is content. It encodes a three-layer reality:

- the mathematics itself contains loci where smooth dependence fails,
- the chosen signature is a lossy abstraction that can fail to behave smoothly even when the underlying system is not maximally unstable,
- the finite computation imposes a horizon beyond which claims are not warranted.

A conventional fractal rendering often hides these layers by forcing continuity through interpolation or by turning instability into attractive texture. In this paper's approach, the layers are left visible, and the coherence gate ensures that when the rendering cannot be honest about local continuity, it becomes silent rather than decorative.

In short: discontinuities are not "where the algorithm broke." They are where the representation rule refused to mislead. They function as the artwork's epistemic conscience, and they are central to why the resulting images can be read as partial sections rather than as maximal-density depictions.

5. One Canvas, Many Scales: Radial Parameterization as a Visual Atlas

A central practical challenge in fractal-based art is the scale problem. The mathematics offers inexhaustible structure at arbitrarily fine resolution, but a static physical artwork must commit to a finite surface and a finite viewing distance. Most fractal imagery "solves" this by leaning on a cultural assumption: viewers will accept that the work is a window into infinity even if the print reveals only a narrow band of that infinity. This assumption works for posters and screens, where the primary aesthetic is spectacle. It works less well if the goal is to make an image that is legible as structure—an object that can be studied, walked, and read as a disciplined construction rather than as a lure of endless detail.

This section proposes a different solution: encode scale directly into the geometry of the composition so that a single printed canvas can simultaneously exhibit coarse and fine behavior without requiring interactive zoom. The key move is to reinterpret the canvas as an atlas of restrictions. Instead of choosing one magnification and letting everything else be implied, the artwork places multiple “neighborhood scales” on a single surface, using radius as a scale coordinate. The viewer’s movement—walking closer or stepping back—then becomes the natural replacement for a software zoom interface.

5.1 The scale problem in static images and why zoom is not a solution

The scale problem is not merely “there is too much detail.” It is that the meaning of the Mandelbrot–Julia relationship often lives in transitions: how behavior changes when parameters move, how stability breaks, how boundaries behave, and how coherence can hold in one regime while failing in another. These are multiscale phenomena. The boundary of the Mandelbrot set is not a feature at one scale; it is a hierarchy of features whose local character depends on where and at what scale one looks. Likewise, the perceived “texture” of parameter dependence changes with neighborhood size: coarse neighborhoods average over regimes; fine neighborhoods reveal delicate structure and hidden instabilities.

A static image forces a choice: either commit to a global view that sacrifices fine structure, or commit to a zoomed view that sacrifices global context. Fractal art culture often resolves this by embracing zoomed detail and offering global context only as a narrative (“it’s the Mandelbrot set, trust me”). This produces visually stunning work, but it reduces the viewer’s relationship to one of consumption: the eye swims in detail without being given a stable reading grammar across scales.

The obvious technical answer is interactive zoom. In a screen-based medium, you can let the viewer scroll or pinch and descend into deeper scales. This is a real solution for exploration, but it is not a solution for a physical artwork for three reasons.

First, interactive zoom changes the aesthetic object. The object becomes a navigation experience, not a single composition. The meaning shifts from “here is a structured artifact” to “here is a space you can traverse.” That can be valid, but it is not what a printed canvas is for. A canvas is a commitment: one surface, one composition, one persistent structure that can be revisited.

Second, interactive zoom delegates coherence to an interface. The viewer’s understanding becomes dependent on the tool’s UI logic—where the zoom centers, how it interpolates, how it loads detail, what it emphasizes. The work risks becoming a demonstration of software, not a statement about structure. In the context of this paper, where we are explicitly concerned with coherence and honest failure, relying on zoom interfaces is a conceptual dodge. It hides the scale problem by outsourcing it.

Third, interactive zoom undermines embodied viewing. A large print invites walking: close inspection reveals microstructure; distance reveals overall form. That bodily dynamic is a powerful aesthetic resource. A zoom interface replaces it with a hand gesture and a moving viewport. The question becomes: can we construct an image that uses the embodied strengths of print rather than imitating the strengths of software?

The answer proposed here is yes: encode multiple neighborhood scales into a single fixed geometry so that scale is not something the viewer requests; it is something the image already contains.

5.2 Radial encoding: radius as neighborhood size around a basepoint c_0

Choose a basepoint $c_0 \in \mathbb{C}$. In practice, c_0 is a parameter of interest—often near a region where the dynamics is rich but not maximally unstable. The image is then constructed by sampling parameters in neighborhoods around c_0, but crucially, the neighborhood size is not fixed. It varies with radius on the canvas.

Let the canvas be a disk (or a disk-like region) with polar coordinates (r, θ), where $r$ranges from 0at the center to 1at the boundary. We define a parameter-space sampling map:

$$c(r, \theta) = c_0 + \rho(r)\, e^{i\theta},$$

where $\rho(r)$ is a nonnegative function that sets the size of the parameter-space neighborhood sampled at radius r.

The central design choice is the shape of $\rho(r)$. In the version used for the radial section motif, $\rho(r)$ decreases as $r$increases. Informally:

- Near the center of the canvas (small r), $\rho(r)$ is larger, so the sampling covers a broader neighborhood around c_0. This region encodes coarse-scale variation.
- Near the edge of the canvas (large r), $\rho(r)$ is small, so sampling stays very close to c_0. This region encodes fine-scale variation.

A typical schedule might look like:

$$\rho(r) = \rho_{\min} + (\rho_{\max} - \rho_{\min})(1 - r)^{\alpha},$$

with $\alpha > 1$ to concentrate increasingly fine neighborhoods near the outer ring. The exact form is not sacred; what matters is that the canvas is turned into a scale axis: moving outward corresponds to shrinking the parameter-space neighborhood.

This is why the construction can be called an “atlas” in a strong sense. An atlas in geometry is a coordinated family of charts that describe an object locally. Here, each ring of radius r is a chart at a particular locality scale around c_0. The canvas is not showing six separate panels; it is showing a continuous hierarchy of neighborhood charts in one composition.

Once $c(r, \theta)$ is defined, the earlier functor-like assignment $F: c \mapsto \sigma(c)$ and rendering map R produce glyphs at the canvas location corresponding to (r, θ). The coherence gate then operates along sampling paths (for example, along radial lines or along angular sweeps), suppressing marks where local representability fails.

Two consequences follow immediately.

First, the image becomes dense in a way that feels earned. It is not dense because we painted every pixel. It is dense because many coherent local assignments survived the gate across many scales.

Second, the image’s “texture gradient” becomes interpretable. The viewer can ask: how does behavior stabilize or destabilize as the neighborhood shrinks? The canvas is no longer merely a picture; it is a fixed record of a restriction process.

5.3 Reading the image: center-to-edge as restriction to smaller neighborhoods

To read the resulting radial image, the viewer must adopt a different habit than the usual “zoom into detail” reflex. The reading grammar is:

- The center is not “the start of the journey into infinity.” It is the coarsest neighborhood summary.
- The edge is not “just a boundary decoration.” It is the finest restriction around c_0 that the image encodes.

In mathematical terms, we can treat each radius r as specifying a disk in parameter space:

$$D_r = \{ c : | c - c_0 | \leq \rho(r) \}.$$

As r increases, $\rho(r)$ decreases, so these disks nest:

$$D_{r_1} \supset D_{r_2} \text{ for } r_1 < r_2.$$

Thus the center-to-edge direction corresponds to a chain of inclusions of neighborhoods, i.e., a restriction system. If one likes the categorical language, one can say: the canvas depicts how the signature assignment behaves under restriction along this inclusion chain.

This reading has two major aesthetic benefits.

First, it makes the work legible at multiple viewing distances. From far away, the viewer sees the gross distribution of “coherence density,” the large forms produced by the glyph field, and the global symmetry imposed by the radial mapping. Up close, the viewer reads the microtexture of glyphs, noticing where coherence is strong and where the gate produces sparsity. The work supports both readings simultaneously because scale is part of the composition rather than hidden behind a UI.

Second, it makes failure meaningful. Suppose coherence is poor in a coarse neighborhood but improves as the neighborhood shrinks; this will show as a broken or sparse interior region and a more continuous outer annulus. Conversely, if coherence fails even at very small neighborhoods around c_0 (for example, if c_0 is near a highly unstable boundary locus), the outer ring will remain ragged or thin. The viewer learns something: not a theorem, but a structural statement about how this signature behaves near that basepoint. The image becomes a diagnostic of local stability under the chosen representation.

This is why “walking” is the correct way to experience the piece. Stepping back gives you the global restriction story. Moving in gives you local evidence. The viewer’s body replaces the zoom tool, and the image remains one coherent object rather than a sequence of screenshots.

5.4 Alternative single-canvas layouts: strips, tilings, and quotient wraps

Radial encoding is not the only way to embed multiple scales in a single static composition. It is chosen here because it makes scale a natural coordinate and because it produces an object with strong presence as a single canvas. But other layouts can implement the same conceptual program—“one law, many scales”—with different aesthetic tradeoffs.

Strips (linear atlases).
A strip layout maps scale onto one axis (say, horizontal position). The image becomes a long ribbon where each vertical slice corresponds to a different neighborhood size around c_0, or to a different path through parameter space. The advantage is clear reading: left-to-right can literally be “coarse-to-fine.” The drawback is compositional: strips often feel like data visualizations or timelines unless carefully designed, and they can underuse the two-dimensional potential of a canvas.

Tilings (discrete atlases).
A tiling layout partitions the canvas into panels, each panel a different chart (different center, different radius). This is the “atlas grid” motif discussed earlier. The advantage is explicit comparison: the viewer can look from panel to panel and immediately see scale effects. The drawback is that the image becomes plural: it reads as six images in one frame rather than as

one unified field. For some installations that is desirable; for a single-canvas “dominant object,” it can feel segmented.

Quotient wraps (wrapped strips on a single surface).
A quotient wrap takes a strip and identifies its ends, wrapping it around a cylinder or torus-like coordinate on the canvas. Practically, this means drawing a long parameter sweep as a wrapped band that coils through the image. The advantage is unity: you get a single connected form that contains a long scale narrative. The drawback is interpretive: wrapped bands require the viewer to learn the mapping, and they can feel more like design patterns unless anchored by strong structural cues.

Spiral atlases (continuous scale with path memory).
A spiral is a hybrid: it encodes scale continuously (like radial) but also encodes traversal order (like a strip). Moving outward corresponds to shrinking neighborhood size, while the spiral path preserves a sense of “sequence.” The advantage is that it can feel like a journey while remaining one object. The drawback is that spiral density can become uneven, and the mapping can dominate the aesthetic if not restrained.

Each of these alternatives can host the same category-theoretic discipline: a base of parameter changes, a fiber of behavioral signatures, a functor-like assignment, a coherence gate, and a rendered partial section. The choice among them is an artistic choice about how the viewer should read the work: as a centered icon, a linear narrative, a comparative plate, or a wrapped topology.

The radial atlas is chosen here because it best satisfies the specific problem raised at the outset: how to make a single canvas that invites walking-distance inspection while allowing multiple scales to coexist without zoom. In the next section, we take the same discipline one step further by treating color not as decoration but as a structured mapping—an additional “assignment” into a perceptual codomain that can, when well designed, behave like a natural transformation between behavior and appearance.

6. Color as Natural Transformation

Color is the quickest way to make a generative image look “alive,” and also the quickest way to destroy its meaning. In fractal art culture, color is often treated as an aesthetic surplus: once the geometry is computed, a palette is applied to maximize impact. That practice can produce beauty, but it tends to sever color from structure. The viewer may feel that color is an emotional overlay rather than a readable encoding. This paper takes the opposite stance: color should be structural. It should function as a disciplined mapping from computed invariants into a perceptual codomain, with explicit rules that make color variation interpretable under parameter variation.

Category theory offers a useful way to phrase this discipline. If the behavioral signature assignment is a functor-like mapping F from parameters to behavior-objects, then a color assignment can be treated as a second functor-like mapping G from parameters to perceptual objects. A "natural transformation motif" then becomes the design goal: changes in behavior and changes in color should be coupled in a coherent way, so that the viewer experiences color as a witness of structure rather than as an arbitrary gradient. This does not require strict categorical machinery; it requires the habit of demanding coherence across two representations of the same underlying parameter movement.

6.1 A perceptual codomain: mapping invariants into HSV or other color spaces

A color model is a choice of codomain: a space in which "small changes" and "large changes" are perceptually meaningful. In practice, the artist chooses a coordinate system such as RGB, HSV/HSB, HSL, Lab, or a palette-based manifold. Each has different affordances.

RGB is device-native but perceptually awkward: equal Euclidean steps in RGB do not correspond to equal perceptual steps. HSV and HSL offer a more controllable decomposition: hue as a circular coordinate, saturation as chroma intensity, and value/lightness as brightness. This separation is artistically useful because it allows one to bind different computed invariants to different perceptual axes.

For the motifs in this paper, HSV is attractive for two reasons.

First, hue is naturally circular. That makes it compatible with angular coordinates in the radial atlas. If the canvas uses θ as a meaningful coordinate (direction in parameter space), then mapping θ to hue produces a perceptually coherent wrap: a full circle in θ corresponds to a full loop in color. This is not merely pretty; it makes the mapping legible. A viewer can learn that "direction corresponds to hue."

Second, saturation and value can be tied to scalar invariants that have a natural "energy" interpretation. For example:

- escape rate can drive saturation (more rapid escape = more vivid chroma),
- boundedness can set a baseline brightness (inside = calmer, outside = more intense),
- expansion proxy can modulate value (high expansion = brighter or darker depending on desired reading).

This yields a structured mapping of the form:

$$\text{Color}(c) = \Phi(\theta(c), e(c), \lambda(c), b(c), \ldots),$$

where Φ is an explicit function into HSV coordinates, and thus into displayable RGB.

Other codomains can be used when the design goal changes.

- A palette manifold can be used if one wants categorical rather than continuous distinctions (e.g., distinct color families for different dynamical regimes).
- Lab or OKLab can be used when one wants more uniform perceptual gradients and smoother “distance” behavior in color space.
- A dual-channel scheme (e.g., hue for one invariant, brightness for another) can be used when one wants the image to remain legible even under desaturation or grayscale printing.

The guiding principle is not “choose a pretty palette.” The guiding principle is: choose a codomain in which the intended invariants produce readable variation and in which the mapping has an interpretable geometry (circular when the base has an angular coordinate, monotone when the invariant is monotone, and so on).

6.2 Two functors F and G: behavior and color

We now formalize the dual mapping.

Let $\mathcal{C}$be the base category of parameters, as in Section 3: objects are parameters c, and morphisms represent small parameter changes (adjacency steps, paths, or restrictions to smaller neighborhoods).

Let $\mathcal{D}$be the behavioral category: objects are signatures $\sigma(c)$ in a space of orbit-derived invariants, with morphisms representing admissible changes (small distances in signature space, or low-cost transitions in a metric-enriched sense).

Define the behavior assignment:

$$F: \mathcal{C} \to \mathcal{D}, F(c) = \sigma(c).$$

Now define a perceptual category $\mathcal{P}$. Its objects are colors (or more generally, perceptual tokens such as HSV triples, palette indices, or color+alpha glyph styles). Morphisms represent “small perceptual changes,” which can again be encoded by distance thresholds or by a perceptually meaningful metric.

Define the color assignment:

$$G: \mathcal{C} \to \mathcal{P}, G(c) = \kappa(c),$$

where $\kappa(c)$ is the chosen color encoding.

It is important to notice what this buys us. Without this framing, color is an afterthought applied to geometry. With this framing, color is a second representation of the same base movement. The viewer is not being asked to admire color; the viewer is being offered color as a second witness. If a region of parameter space is stable under the behavior assignment F, we should expect it to be stable under the color assignment G as well—unless the design intentionally chooses otherwise. If a region is unstable, both behavior and color should register that instability, either through abrupt shifts or through silence imposed by the coherence gate.

In a strict category-theoretic setting, one would demand that both $F$and Gbe functors and that morphisms map appropriately. In our generative setting, we instead demand a practical analogue: under small parameter changes, both the behavioral signature and the color should vary in a way that does not contradict the story we are trying to tell.

6.3 Natural-transformation motif: coherent color change under parameter change

A natural transformation is, informally, a way of translating between two functors while respecting the structure of the base category. In our context, the "natural-transformation motif" is a design goal: define color as a structured translation of behavior so that changes in parameter produce commensurate changes in both.

We can express this idea in the simplest way: the color should be a function of the signature,

$$\kappa(c) = \phi(\sigma(c)),$$

so that

$$G = \phi \circ F.$$

Here $\phi: \mathcal{D} \to \mathcal{P}$ is a rendering map from behavioral signatures to colors. When ϕ is chosen so that small signature changes produce small color changes, the color field becomes a faithful perceptual witness of the behavior field.

In practice, one often wants color to depend on some base geometry as well (for example, using θ to set hue). That gives a mixed scheme:

$$\kappa(c(r, \theta)) = \phi(\sigma(c(r, \theta)), \theta).$$

This is still compatible with the naturality motif as long as the dependency is interpretable and coherent: hue tracks direction, saturation tracks escape energy, and brightness tracks boundedness or expansion. The key is that the mapping is rule-based and stable, not arbitrary.

What does “natural” mean visually? It means the viewer can learn a reading grammar and apply it consistently across the image:

- If direction changes smoothly around the disk, hue cycles smoothly.
- If escape energy increases, saturation increases.
- If we cross from bounded to escaping regimes, the brightness baseline changes in a recognizable way.

When these conditions hold, color becomes a structural layer that reinforces what the glyph geometry already suggests. The viewer experiences agreement between witnesses: the shape language and the color language tell the same story.

The coherence gate introduced earlier can be extended to color as well. One can require that when parameters are close and in the stable regime, not only must σ vary gently, but κ must also vary gently. If either jumps beyond a threshold, the mark is suppressed. This enforces a stronger version of the naturality motif: not only must behavior be locally coherent, but the translation into color must not introduce discontinuity that is not already present in behavior. The result is an image where color discontinuities are meaningful rather than accidental.

The larger conceptual payoff is that the artwork becomes an instance of “two functors, one base.” The viewer sees that the parameter plane is being translated into two different codomains: a behavior codomain (glyph geometry) and a perceptual codomain (color). The natural-transformation motif is the guarantee that these translations do not contradict each other. The work thereby resists a common failure mode of generative art: gorgeous color gradients laid over structure that does not support them.

6.4 Design choices: when color clarifies and when it becomes decoration

The difference between clarifying color and decorative color is not primarily a matter of taste. It is a matter of whether color preserves the coherence story.

Color clarifies when it does at least one of the following.

First, it encodes a base coordinate in a way that matches the geometry. In the radial atlas, mapping θ to hue is clarifying because it gives a stable, learnable relation between direction and color. A viewer can point to a region and infer something about the direction in parameter space being sampled there. This is especially helpful when the geometry is dense: hue can help segment the field without introducing arbitrary boundaries.

Second, it encodes a scalar invariant monotonically. Escape rate is a classic candidate: mapping higher escape to higher saturation or higher brightness produces an “energetic” reading that is easy to internalize. The viewer may not know the formula, but the monotonicity is felt: more intense regions correspond to faster divergence.

Third, it distinguishes regimes cleanly. For example, inside vs outside can be given different brightness baselines or different saturation ranges. This can prevent confusion in regions where glyph geometry alone would be ambiguous.

Fourth, it remains legible under reduction. A clarifying color scheme often still works if the image is printed in grayscale: the brightness channel carries meaningful structure. If the scheme collapses into noise when desaturated, it is often a sign that the color was doing too much unstructured work.

Color becomes decoration when it violates these principles.

A common decorative failure mode is palette arbitrariness: shifting hue in a way that has no anchor in base geometry or invariants, so that neighboring parameters can be assigned dramatically different colors for no structural reason. This breaks the naturality motif and trains the viewer to ignore color as meaning.

Another failure mode is excessive gradient complexity. One can design color maps so intricate that any local variation in the invariants produces high-frequency hue variation. This produces a visual “shimmer” that overwhelms the coherence story. It can be beautiful, but it tends to erase the distinction between stable and unstable regions because everything becomes equally vivid.

A third failure mode is contradicting the silence rule. If the coherence gate suppresses marks to preserve honesty, but color is used to fill in the suppressed regions with a smooth gradient, the work smuggles continuity back in through the perceptual layer. In this framework, that is a categorical violation: it reintroduces the lie the gate was designed to prevent.

For these reasons, the recommended design discipline is conservative: use one base coordinate (such as angle) to control one perceptual axis (such as hue), and use one or two invariants to control saturation and brightness. Prefer monotone, interpretable mappings over clever palettes. Then treat any remaining need for “beauty” as something that should emerge from structure and accumulation rather than from ornamental color design.

When this discipline is followed, color behaves like a natural transformation in the practical sense: it is a consistent translation of behavior into perceptual space that respects the local morphisms of the base. The result is a color image that is not simply “more attractive,” but more readable. It allows the viewer to see coherence, failure, stability, and transition—exactly the structural motif that the paper argues should replace fractal depiction as the central aesthetic primitive.

7. Interpreting the Visible Features

A coherence-driven image should not require the viewer to know the code in order to read it. The entire point of treating the artwork as a partial section of a behavior bundle is that the visible forms are meant to be interpretable as structural evidence: evidence of where the assignment $c \mapsto \sigma(c)$ behaves smoothly, where it collapses distinctions, and where it refuses representation because local continuity fails. This section offers a reading grammar for the most common macroscopic features produced by the construction—plumes, annuli, spines, and gaps—framed explicitly as categorical signals.

The viewer should think of the image not as a single "fractal object," but as a map of local representability under a chosen signature and coherence gate. Large features are not painted; they are emergent from many small glyph decisions. Therefore the right interpretive question is not "what shape is this?" but "what does this region tell me about coherence, stability, and distinguishability of behavior under parameter change?" In this sense the image is closer to a field diagram than to a decorative motif, with the crucial difference that it is built from refusal as well as from accumulation.

7.1 Plumes, annuli, and spines as categorical signals

The most salient large-scale patterns in radial atlas renderings often include plume-like fans, ring-like annuli, and spine-like ridges. These are not "hidden pictures." They are signatures of how coherence constraints interact with the radial restriction geometry.

Plumes typically appear as fan-shaped regions of dense glyphs that broaden or taper as they move radially. Categorically, a plume indicates a direction in parameter space along which the functor-like assignment behaves predictably over a range of neighborhood scales. Along that angular sector, small changes in $c$tend to produce small changes in $\sigma(c)$, and the coherence gate therefore allows marks to accumulate. The plume's internal texture—whether it is smooth, striated, or speckled—suggests how uniformly the signature varies within that sector. A smooth plume corresponds to a region where the signature changes are not only small but also consistent in direction (in signature space). A striated plume suggests the presence of substructures: alternating zones where coherence holds more strongly or weakly, often reflecting the influence of nearby dynamical features (e.g., proximity to bifurcation boundaries or internal "components" in the parameter landscape at the sampled scale).

Annuli (ring-like bands) are especially natural in a radial encoding because radius is literally a scale coordinate. When an annulus appears as a distinct band of higher or lower mark density, it suggests a threshold scale: a neighborhood size at which the coherence behavior changes qualitatively. For example, the inner region may be sparse because coarse neighborhoods around c_0 mix behaviors that cause the signature field to vary too abruptly; then, beyond a certain radius, neighborhoods become small enough that the signature field stabilizes and coherence begins to hold, producing a denser ring. Conversely, an annulus of thinning near the outer region may indicate that even at small neighborhood scales, the chosen signature has become sensitive to transient or boundary-adjacent effects, causing the coherence gate to suppress marks. In categorical terms, annuli are visual evidence of how restriction to smaller neighborhoods changes the effective “regularity” of the assignment.

Spines (thin ridges or sharp lines of dense or sparse activity) often arise at angular directions that align with structural axes of the family or with the chosen symmetry cues. In the simplest case, a spine may coincide with the real axis direction in parameter space, where conjugation symmetry is strongest and the signature may exhibit distinct behavior. More generally, spines can mark directions of heightened sensitivity: parameter rays along which the signature changes rapidly with angle, leading to coherence failure on either side but relative stability along the spine itself. In that case the spine is the trace of a preferred “morphism direction” in the base: a direction along which the mapping is comparatively coherent even when neighboring directions are not. If the spine is a region of suppression rather than density, it may represent a direction where the signature crosses a regime boundary repeatedly under small perturbations, causing persistent coherence failure.

In summary, plumes suggest coherent directions, annuli suggest scale thresholds, and spines suggest axial structure or anisotropic stability. These are categorical signals because they are about how structure is preserved (or not) under morphisms in the base, not about the incidental geometry of a fractal boundary.

7.2 Non-faithfulness regions: many parameters, similar signatures

A key property of the signature map $\sigma(c)$ is that it is lossy. It compresses a complex dynamical system into a handful of numbers. This is not a weakness; it is a design decision. But it has a categorical consequence: the assignment will be non-faithful in the everyday sense. Many distinct parameters will map to signatures that are close or even identical within numerical tolerance. In the image, these become regions where the glyph field looks uniform even though parameter values vary across a substantial region of the base.

Such regions are important to recognize because they prevent the viewer from misreading uniformity as triviality. Uniformity can mean one of two things:

1. The underlying dynamics genuinely varies little across that region at the scale being sampled, so the behavior is locally stable in a strong sense.
2. The signature is too coarse to distinguish meaningful differences, so distinct behaviors are being collapsed by the representation.

Both are “non-faithfulness,” but of different moral types.

The first type is a genuine stability plateau. In such a region, the coherence gate will tend to allow marks densely because signature variation is small. The uniformity will appear as a consistent mark thickness, curvature, and color over a broad patch. This is the image’s way of showing that, relative to the chosen signature, the system behaves like a locally trivial bundle: the fiber does not change much as you move in the base.

The second type is representational collapse. Here, the signature may be stable, but only because it ignores distinctions that matter in the full dynamics. The image will again look uniform, but one expects this uniformity to break when the signature is refined or when a different signature component is used (for example, a period estimator or a more sensitive orbit statistic). This matters for interpretation: the image is not claiming “nothing changes here.” It is claiming “nothing changes here according to the witness we chose.”

In categorical language, these are two flavors of “failure of faithfulness.” The functor-like assignment collapses different base points into similar fiber objects. In a strict setting, one might try to repair this by enriching the fiber category or refining the functor. In an artwork, the point is to make the collapse visible. A region of uniform texture can be read as an invitation: either it is a real stability basin, or it is the signature telling you where it cannot resolve differences. That ambiguity is not a flaw if the paper makes the signature explicit and the coherence rules transparent. It becomes part of the work’s epistemic posture.

One practical cue helps distinguish the two types. If uniformity persists across multiple scales (from center to edge in the radial atlas), it is more likely to reflect genuine stability. If uniformity holds at coarse scales but fragments at finer scales, it suggests that coarse neighborhoods were hiding differences that become resolvable under restriction. This is precisely the kind of distinction the radial atlas makes visible.

7.3 Failure loci: where coherence breaks and silence appears

If non-faithfulness is the collapse of distinctions, failure loci are where the attempt to preserve coherence breaks down. In the artwork, failure loci appear as thinning, tearing, gaps, or abrupt changes in density. They are the visible trace of the coherence gate suppressing marks because signature variation exceeded admissible thresholds under small parameter steps.

There are several reasons such loci appear, and each has a distinct interpretation.

Regime boundaries and bifurcation-adjacent behavior. Near the true Mandelbrot boundary, the boundedness classification itself becomes sensitive at finite depth. The signature can change abruptly because the orbit escapes in one case and remains bounded in another. Silence in such regions is a direct signal: the base category regards these parameters as close, but the fiber objects differ too much under the chosen signature. The local morphism does not lift to a small morphism in the fiber. The image refuses to claim continuity.

Signature instability and transient sensitivity. Even inside the bounded regime, the expansion proxy or other orbit-derived measures may fluctuate due to transient phases of the orbit. Two nearby parameters can produce similar long-term behavior but different finite-time proxies. When this triggers the gate, silence is reporting the instability of the chosen witness at that sampling resolution. This is not a mathematical boundary in the strict sense; it is a boundary of representability under the chosen measurement scheme. Again, the image remains honest by refusing to draw rather than smoothing.

Numerical horizon limits. Finite iteration and finite precision impose a horizon. A parameter may appear bounded at depth $N$and then escape at depth $2N$. Near such “late escape” situations, small changes in $c$can change the escape time dramatically. The coherence gate often suppresses marks here, producing a halo of silence. This halo is not meaningless; it is the visual footprint of computational uncertainty. In a work that cares about epistemic discipline, this is valuable. It tells the viewer where the representation cannot claim stability.

The most important interpretive principle is that silence is not emptiness. It is a categorical statement: the section cannot be extended across this locus without violating the coherence rule. In the language of gluing, the local sections do not agree; in the language of bundles, a continuous extension is obstructed; in the language of functoriality, small base morphisms are not mapped to small fiber morphisms. The image makes this failure visible rather than hiding it.

One can also interpret failure loci as the artwork’s “boundary of trust.” In ordinary images, the entire surface is equally confident. In this framework, confidence is spatially encoded: dense coherent regions are high-trust; silent regions are explicit low-trust. This is not merely aesthetic. It is a methodological stance about how computational art should behave when it is built from unstable systems.

7.4 What the viewer learns by walking: perception as a multiscale operation

The radial atlas is designed for the body. It assumes a large print where walking replaces zoom. This matters because it turns perception itself into a multiscale operation: the viewer does not just see “more detail up close.” The viewer performs a restriction process that mirrors the mathematical restriction encoded by radius.

From a distance, the viewer sees global forms: overall symmetry, dominant plumes, annular thresholds, major failure loci. At this distance, the work reads like a map of coherence: a macroscopic diagram of where stability holds and where it fractures. The viewer can already infer something about the basepoint c_0: whether its neighborhood is broadly stable, whether it sits near a boundary, whether certain directions in parameter space are privileged.

As the viewer approaches, the glyph field becomes readable. Strokes reveal their thickness modulation, curvature cues, and color logic. The viewer begins to read the signature not as numbers but as a consistent "visual language." This is where the natural-transformation motif matters: if color and geometry agree, the viewer feels that the microstructures are not arbitrary decoration but repeated instances of a law.

At very close range—within the distance at which individual strokes become objects—the viewer encounters the work as a surface of decisions. The coherence gate becomes visible as pattern: clusters of marks, sudden gaps, edges where density drops. The viewer sees the image as a record of a selection rule, not as a continuous field. In other words, the work teaches the viewer to see like a category theorist sees: not as a consumer of a final picture, but as an interpreter of how local pieces fit together and where they do not.

This "walking reading" produces a subtle educational effect. It trains the viewer to understand scale not as a mere zoom factor but as a change in what can be said coherently. At coarse scales, many differences collapse and some coherence fails. At finer scales, some coherence is restored and some failures sharpen. The viewer experiences, physically, that "truth of representation" is scale-dependent. That is a powerful lesson that applies beyond fractals: it is the same lesson encountered in measurement theory, model selection, and any attempt to compress complex systems into finite summaries.

The outcome is that the image functions as both art and methodological demonstration. It does not merely present the Mandelbrot–Julia relationship as a spectacle of complexity. It presents it as a disciplined story about coherence, witnesshood, and the boundaries of representation—made legible at human scale through the simple act of walking.

8. Novelty and Relation to Prior Practice

Any claim of novelty in fractal art must be made carefully. The field has decades of experimentation behind it: escape-time renderings, smooth coloring, distance estimation, orbit traps, perturbation methods, deep zooms, flame algorithms, hybrid formulas, shader-based explorations, and countless palette innovations. It is therefore easy to say something "new" that is actually a minor variation on a well-trodden technique. The novelty claimed in this paper is not a new fractal formula and not a new coloring trick. It is a shift in optimization target: from maximizing spectacle to enforcing coherence as an aesthetic primitive, and from treating instability as ornament to treating it as a boundary of honest representation.

The proposal is best understood as a recombination of familiar ingredients under an unfamiliar discipline. The ingredients—critical-orbit features, parameter-plane sampling, multiscale inspection, and color mapping—are not exotic. What is unusual is the insistence that the image be interpretable as a partial section of a derived bundle, and that the construction allow refusal (silence) where the underlying assignment does not behave coherently at the sampling scale. This section makes that claim precise and, equally important, states the limits and humility conditions that prevent it from becoming an overconfident manifesto.

8.1 What fractal art typically optimizes for

A useful way to summarize mainstream fractal art practice is to describe its implicit objective function. While artists vary widely in taste and intent, many of the most common workflows optimize for some combination of:

- **High visual yield per pixel.** The goal is dense, ornate structure that rewards staring and zooming. Rendering techniques that produce more "interesting detail" are preferred.
- **Immediate impact.** Strong contrast, vivid saturation, and high-frequency texture create a rapid affective response. The viewer should be impressed quickly.
- **Palette-driven mood.** Color gradients and carefully curated palettes are used to evoke emotion and depth, often independent of whether the palette corresponds to a meaningful invariant.
- **Boundary fetish.** The most elaborate structure occurs near boundaries (especially the Mandelbrot boundary), so compositions often emphasize boundary-adjacent regions and deep zooms into miniature "universes."
- **Continuous filling.** A completed image is often expected to have minimal emptiness. Blank regions are perceived as wasted potential, so techniques that fill the plane with texture are favored.
- **Technique signatures.** Many subcultures develop recognizable styles: orbit traps have their look, flame fractals have theirs, distance-estimated sets have theirs. The "signature" of technique can become the brand.

None of these aims are illegitimate. They simply define a particular aesthetic: the fractal as an inexhaustible ornament engine. The optimization, however, has a side-effect: it encourages artists to treat discontinuity and instability as decorative opportunities, and it discourages any rule that would intentionally suppress marks. In a culture of “fill the plane with wonder,” silence looks like failure.

This also explains why many fractal artworks converge visually. If the objective function rewards high detail density, high saturation, and boundary complexity, then different formulas and different artists can still end up producing images that are stylistically similar: the mathematics becomes a substrate for a stable set of aesthetic preferences. The work becomes fractal-themed rather than structure-expressive.

8.2 How coherence constraints differ from escape-time aesthetics

Escape-time aesthetics are built on a simple classifier: iterate a point until it escapes, and color based on escape time. This is powerful because it produces structure from a single scalar and because it captures the global boundary between bounded and unbounded behavior. But it also encourages a particular kind of representation: continuous fields of color that suggest continuity everywhere, even when the underlying dynamics is maximally sensitive.

The coherence-constrained approach in this paper differs in four structural ways.

First, it replaces “every sample produces a value” with “some samples are refused.” In escape-time plots, every pixel is assigned a color. Discontinuities are not shown as discontinuities; they are smoothed by the field representation. In the coherence approach, discontinuity becomes visible as suppression. If local coherence fails under the chosen signature and sampling resolution, the mark is withheld. This is not an aesthetic accident; it is the defining rule.

Second, it shifts the emphasis from membership to variation. Escape-time plots are excellent at showing the boundary of membership (inside vs outside). They are less explicit about the question that motivates the category-theoretic framing: how does behavior vary under small parameter change, and where does that variation behave like a structure-preserving mapping? The coherence gate makes that question primary. It turns “variation under morphisms” into the visible content.

Third, it changes what counts as “interesting.” In escape-time art, the most interesting regions are often those with maximal gradient complexity. In coherence-constrained art, maximal instability is not automatically interesting; it is often a reason for silence. Interest shifts toward the tension between coherent regions and failure loci, toward anisotropies (directions where coherence holds better), and toward scale thresholds (annuli where restriction changes behavior). The boundary is still important, but it is important as a coherence stress-test, not as a source of decoration.

Fourth, it introduces a compositional multiscale grammar that does not rely on zoom. Escape-time aesthetics often assume a single scale per image (even if that scale is a deep zoom). The radial atlas explicitly encodes scale in the geometry. This changes the relationship between detail and meaning: microstructure is not merely “more stuff”; it is evidence of behavior under restriction to smaller neighborhoods.

In short, escape-time aesthetics render the set. Coherence constraints render the honesty of an assignment.

8.3 What is genuinely new here and what is a recombination

To avoid overclaiming, it helps to separate novelty into three tiers: mathematical novelty, technical novelty, and conceptual/disciplinary novelty.

Mathematical novelty.
This work does not claim new results in complex dynamics or category theory. The quadratic family, the critical orbit criterion, and the connectedness relationship are classic. The categorical vocabulary (base/fiber, functor-like assignment, partial section) is standard. The novelty is not in new theorems.

Technical novelty.
Some components are technically common in the fractal community (escape rates, derivative-based proxies, parameter-plane sampling). The radial multiscale encoding is not unheard of as a design idea, but it is far less common than zoom-based presentation, and it is not a canonical fractal rendering pattern. The explicit use of a coherence gate as a first-class rendering rule is, in practice, uncommon. Artists often apply smoothing and interpolation to reduce artifacts; few deliberately suppress marks to preserve a coherence claim. So there is a real technical distinction: the gate is not a post-process; it is the generator of negative space.

Conceptual and disciplinary novelty.
This is where the strongest claim belongs. The genuinely new contribution is the redefinition of the aesthetic objective:

- The artwork is constructed to be interpretable as a **partial section** of a derived behavior bundle.
- The primary aesthetic primitive is **coherence under small changes**, not maximal detail.
- The boundary of stability is treated as a **boundary of representability** rather than as a decoration source.
- Silence is used as **structured refusal** rather than as empty background.
- Color is disciplined as a **structured translation** (natural-transformation motif) rather than than as palette play.

Each of these, taken alone, might be described as a recombination or as a conceptual reframing. Taken together, they form a coherent methodology that differs sharply from typical fractal-art practice. The novelty is not "a new picture of the Mandelbrot set." The novelty is "a new contract between computation and representation."

It is also worth naming what is explicitly a recombination.

- Using the critical orbit to drive rendering is standard.
- Using derivative magnitude along orbits as an expansion cue is standard in dynamical diagnostics.
- Mapping geometric coordinates into color is standard in generative art.
- Encoding scale into layout is standard in scientific visualization (though less common in fractal art's mainstream aesthetic).

The paper's contribution is to bind these familiar components under one governing law: coherence preservation and refusal where coherence fails.

8.4 Limits: numerical artifacts, parameter choices, and interpretive humility

A coherence-driven framework makes strong ethical demands of the renderer, but it must also make strong humility statements about what the image can and cannot claim. There are at least four limit classes.

Finite iteration and false stability.
Boundedness classification at depth N is not equivalent to true membership in the Mandelbrot set. Late escape can masquerade as bounded behavior. Therefore any claim of "inside" is operational, not absolute. In the framework of this paper, this is addressed not by pretending to certainty, but by letting unstable regions become sparse under the coherence gate. Still, the viewer should be told explicitly: the image is built from finite horizon witnesses.

Threshold and metric dependence.
The coherence gate depends on thresholds ε and δ, and on a choice of norm in signature space. Changing these will change where marks are suppressed. There is no single "correct" threshold; there are only thresholds that implement a chosen standard of honesty. A too-strict gate can over-suppress and produce silence that reflects the gate more than the dynamics. A too-lenient gate can flood the image and erase meaningful discontinuities. This is an unavoidable design dependency and should be acknowledged as such.

Basepoint dependence.
The radial atlas is organized around a chosen c_0. Different basepoints yield different atlases. Near some c_0, restriction to smaller neighborhoods will increase coherence; near others, even

tiny neighborhoods remain unstable. Thus the image is not “the Mandelbrot set.” It is “a multiscale coherence atlas around c_0 under this signature and gate.” This local nature is a feature, but it requires interpretive clarity: viewers should not generalize beyond what the construction supports.

Signature insufficiency and representational collapse.
The signature $\sigma(c)$ is a lossy compression. Non-faithfulness regions are expected. The image cannot claim to distinguish all relevant dynamical differences. It can only claim to distinguish what the signature witnesses. The proper interpretive humility is: this is one possible fiber object attached to parameters, not the canonical one.

These limits are not embarrassing; they are integral to the paper’s thesis. The entire point of coherence constraints and partial sections is to build artwork that admits the boundary between what is representable and what is not. Interpretive humility is therefore not an apology appended at the end; it is part of the aesthetic contract. The viewer is invited to enjoy the piece not as a declaration of total knowledge, but as a disciplined witness of a complex system under a stated set of measurement and coherence rules.

In that sense, the paper’s novelty is inseparable from its humility. A work that claims categorical structure without admitting representation limits risks becoming exactly what this framework seeks to avoid: an oracle image, confident everywhere, accountable nowhere. Here the image is accountable by construction. It shows where it can speak and where it cannot, and it treats that boundary as the most honest and therefore the most meaningful part of the art.

9. Implementation Notes for Reproducibility

The goal of this section is to make the construction reproducible without turning the paper into a code dump. Reproducibility here means: a technically competent reader should be able to rebuild the artwork with a minimal set of algorithmic steps, a short parameter list, and clear tradeoff guidance. Because the piece is deliberately a “partial section” rather than a dense continuous field, the implementation must also specify the coherence gate and its thresholds; otherwise, another implementation might silently revert to a conventional “draw everything” workflow and lose the conceptual content.

The notes below assume the quadratic family $f_c(z) = z^2 + c$, the radial atlas sampling map around a basepoint c_0, a signature $\sigma(c)$ built from the critical orbit, and a stroke-based glyph vocabulary with optional HSV color encoding. However, the overall skeleton is modular: one can swap in other families and other signatures while retaining the coherence and partial-section discipline.

9.1 Minimal algorithm and parameter list

A minimal reproducible implementation can be described as five stages: sampling, orbit computation, signature extraction, coherence gating, and rendering.

Stage A: Choose global parameters.
The following are the core parameters that define the image:

- **Canvas geometry**: width W, height H, and whether the drawing region is a disk or a full rectangle with a disk mask.
- **Sampling density**: number of angular samples N_θ, number of radial samples N_r. Total samples $\approx N_\theta N_r$.
- **Basepoint**: $c_0 \in \mathbb{C}$.
- **Neighborhood schedule**: $\rho(r)$ with $\rho_{\max}$, $\rho_{\min}$, and exponent α controlling how quickly neighborhoods shrink with radius.
- **Orbit computation**: max iterations N, escape radius R, and numerical precision (standard double precision is usually sufficient for moderate depths).
- **Signature extraction**: choice of components and normalization ranges.
- **Coherence gate**: parameter adjacency threshold ε(often implicit from sampling) and signature-change threshold δ.
- **Glyph mapping**: stroke length range $[L_{\min}, L_{\max}]$, thickness range $[T_{\min}, T_{\max}]$, curvature scaling, and orientation rules.
- **Color mapping (optional)**: HSV mapping functions and clamping ranges.

A compact "default" parameter list might be:

- $W = 6000$, $H = 6000$ pixels (for print-ready), disk mask radius $\approx 0.48\min(W, H)$.
- $N_r \in [800,2000]$, $N_\theta \in [1200,4000]$ depending on desired density.
- $c_0$chosen near a region of interest (e.g., near but not on a boundary).
- $\rho_{\max}$ and $\rho_{\min}$ such that $\rho_{\max}$ captures visible coarse variation and $\rho_{\min}$ probes fine neighborhoods without collapsing to numerical noise (typical ranges depend strongly on c_0).
- $N \in [200,2000]$, $R \in [2,4]$.

- Coherence δ tuned so that stable regions render densely but boundary-adjacent regions tear.

Stage B: Sample the parameter plane via radial atlas.
For each radial index $i = 0, \ldots, N_r - 1$and angular index $j = 0, \ldots, N_\theta - 1$, define:

$$r_i = \frac{i}{N_r - 1}, \theta_j = 2\pi \frac{j}{N_\theta}.$$

Compute neighborhood radius:

$$\rho_i = \rho(r_i).$$

Map to parameter:

$$c_{ij} = c_0 + \rho_i e^{i\theta_j}.$$

Map to canvas position:

$$x_{ij} = x_0 + R_{\text{pix}}\, r_i \cos\, \theta_j, y_{ij} = y_0 + R_{\text{pix}}\, r_i \sin\, \theta_j,$$

where (x_0, y_0) is the canvas center and $R_{\text{pix}}$is the pixel radius of the disk region.

Stage C: Compute critical orbit and extract signature.
Iterate:

$$z_0 = 0, z_{n+1} = z_n^2 + c_{ij}$$

until $n = N$or $| z_n |> R$. Record:

- escape time τ,
- escape magnitude $| z_\tau |$ if escaped,
- derivative magnitudes $| 2z_n |$ along the computed segment.

Compute signature components (example):

- $b = 1$ if no escape by N, else 0,
- $e = \frac{\log |z_\tau|}{\tau}$ if escaped, else 0 or a bounded-regime placeholder,

- $\lambda = \frac{1}{m} \sum_{n=0}^{m-1} \log(|\, 2\, z_n \,| + \epsilon_0)$ where $m = \min(\tau, N)$ and ϵ_0 avoids log of zero,
- $s = 1$ if $|\, \mathrm{Im}\,(c_{ij}) \,| < \eta$, else 0.

Normalize each component into a stable range (e.g., $[0, 1]$) using clamping and known plausible bounds.

Stage D: Apply coherence gate.
Define a notion of adjacency. The simplest is to compare successive angular samples at fixed radius, and/or successive radial samples at fixed angle. For example, along angle at radius i:

$$d_{ij} = \| \sigma(c_{ij}) - \sigma(c_{i,j-1}) \|.$$

If $d_{ij} > \delta$ (and optionally $b = 1$ for both points), then suppress the mark at (i, j). One can symmetrize by also checking radial neighbors.

This gate can be expressed as: draw a mark only if local signature variation is admissible under the base adjacency implied by the sampling grid.

Stage E: Render glyph.
For each non-suppressed sample, compute glyph parameters:

- length Las a function of b, e, λ,
- thickness Tas a function of $e$or $|\, \lambda \,|$,
- curvature Kas a function of λ,
- orientation ϕ as a function of θ and symmetry cue s,
- color via HSV mapping if enabled.

Draw a short stroke (line segment or small Bezier curve) centered at (x_{ij}, y_{ij}) with the computed parameters. Optionally add alpha blending to allow accumulation.

That is the complete minimal algorithm. The paper's conceptual structure is encoded almost entirely in Stage D (the gate) and in the fact that Stage A explicitly treats $\rho(r)$ as a scale axis.

9.2 Complexity and rendering tradeoffs

The computational cost is dominated by orbit iteration. If there are $S = N_r N_\theta$ samples and each sample iterates up to Nsteps, then worst-case time is $O(SN)$. In practice, many points escape early, so average cost can be lower, but parameters inside the Mandelbrot set do not escape and thus cost the full Nsteps.

This leads to the main tradeoffs.

Iteration depth vs. density.
Increasing N makes boundedness classification more faithful (fewer false bounded points near the boundary) and can stabilize some signature components. But it increases cost linearly and can increase sensitivity: deeper iterations can reveal instabilities that produce more suppression under the coherence gate. A higher N can therefore paradoxically produce “more silence” unless $\delta$is adjusted.

Sampling density vs. coherence readability.
Higher N_r, N_θ increases visual density and produces smoother macroscopic forms. But it also increases the number of adjacent comparisons, potentially amplifying suppression in unstable regions (because more opportunities exist for threshold exceedance). Density should therefore be tuned together with δ.

Gate strictness vs. interpretability.
A strict δyields strong, legible silence boundaries and emphasizes the partial-section concept. Too strict, and the image becomes underdrawn, dominated by gaps that reflect threshold choice more than dynamics. Too lenient, and the image becomes a conventional dense texture field, losing the conceptual distinction between coherent and incoherent regions.

Stroke complexity vs. print clarity.
Curved strokes and variable thickness can carry more structure, but they also increase the risk of muddying at print scale when marks overlap. For large canvases, thinner, simpler strokes often read better from a distance and still reward close inspection. A practical rule is to design for two distances: (i) across-room readability (macroscopic plumes/annuli), and (ii) within-arm’s-length readability (glyph texture).

Vector vs raster rendering.
Vector output (SVG/PDF) is attractive because it scales cleanly, but very large numbers of strokes can make vector files heavy and slow to render. Raster output (PNG/TIFF) is straightforward but demands high pixel dimensions for large print. A hybrid approach is common: render at high-resolution raster, or render in vector but with stroke batching and simplification.

9.3 File formats, print considerations, and installation scaling

Because the work is explicitly intended to be walk-up legible, print planning is part of the technical design, not an afterthought. There are three questions: resolution, format, and physical substrate.

Resolution and pixel dimensions.
For large-format prints, a typical target is 150–300 DPI at final size, depending on viewing distance. If the canvas is intended to be examined closely, higher DPI is beneficial, but practical limits arise quickly. For example, a 60" × 60" print at 200 DPI requires 12,000 × 12,000 pixels, which is feasible but heavy. A key design choice is to ensure that individual glyph strokes have enough pixel support that they do not alias into noise.

Color management.
If color carries structural meaning, it must survive printing. HSV-designed mappings can shift under printer profiles. If the color scheme relies on subtle differences, those differences may be lost. Practical mitigation includes: (i) ensuring the brightness channel carries meaningful structure so grayscale remains legible, and (ii) testing a small print proof to verify that hue/saturation distinctions survive.

File formats.

- **PNG** is convenient and lossless but can be large.
- **TIFF** is common for print workflows and supports high bit depth.
- **PDF/SVG** are useful for vector workflows but can become unwieldy if millions of strokes are included.
- For installation-scale work, a tiled workflow (exporting in quadrants or strips) can be used to avoid single gigantic files.

Installation scaling and the "walk" requirement.
The strongest experiential version of the piece is physically large: the viewer should be able to step back to see annuli and plumes, then step forward to see glyph texture. If the piece must be smaller, the same effect can be preserved by increasing glyph size and reducing density, ensuring that microstructure remains inspectable. Conversely, if the piece is very large (wall-scale), density can be increased and stroke thickness reduced so that close inspection reveals many layers.

A practical note: the radial atlas is inherently circular. If printed on a rectangular canvas, one must decide whether to mask to a disk (leaving corners blank) or to extend the mapping into corners with a secondary rule. The disk mask is conceptually clean: it reinforces the "atlas" interpretation and avoids forcing the geometry to fill every corner.

9.4 Variants: other dynamical families and other signature choices

The framework is not tied to the quadratic family. What is essential is the existence of a parameterized family of dynamical systems with meaningful stability variation, and the ability to compute a compact behavioral signature. Several variants are natural.

Other dynamical families.

- **Higher-degree unicritical polynomials**: $z^d + c$. These preserve the "single critical orbit controls much" flavor while changing symmetry and boundary structure.
- **Multicritical polynomials**: families with multiple critical points, requiring a signature that aggregates multiple critical orbits. This can produce richer "fiber complexity," but also complicates coherence because there are multiple sources of instability.
- **Rational maps**: e.g., $z \mapsto z^2 + \frac{\alpha}{z^2+\beta}$ or similar families. These introduce poles and more intricate parameter phenomena, potentially producing more frequent failure loci—which may be aesthetically powerful if treated honestly.
- **Real one-dimensional families** (logistic map): these offer a different kind of atlas (interval parameter base, invariant measures as fibers), potentially enabling a strip-based rather than radial layout.
- **Iterated function systems** (IFS): parameters control contraction ratios and rotations; fibers are attractor sets. Coherence constraints can be extremely interpretable here because parameter continuity often produces attractor continuity until bifurcation-like thresholds.

Other signature choices.
The signature $\sigma(c)$ used in this paper is deliberately lightweight. It can be enriched in many ways:

- **Period detection**: attempt to detect attracting cycle period of the critical orbit in bounded regimes. This yields categorical regime labels that can drive glyph types.
- **Distance estimation**: include a distance-to-set proxy, providing smoother fields than escape time and potentially more stable coherence behavior.
- **Orbit trap statistics**: measure how often orbits pass near certain geometric traps; these produce rich signatures but can become decorative if not gated.
- **External angle/ray information**: for parameters where rays are computable, encode angle data as a structural coordinate rather than as decoration.

- **Entropy-like proxies**: approximate complexity of orbit segments with simple compressibility measures; these can create a direct bridge to MDL-style interpretations.

In each variant, the categorical discipline remains unchanged: define base morphisms (small parameter moves or restrictions), define fiber objects (signatures), enforce coherence via a gate, and render as a partial section with visible refusal where coherence fails. That discipline is the portable core.

A final note on reproducibility: because the gate and the signature are tunable, it is valuable to treat each artwork instance as a "parameterized experiment" and to record the parameter list alongside the image (even if only in metadata or an appendix). This aligns with the paper's ethos: the work is not an oracle image that arrived by inspiration; it is a reproducible witness generated under explicit constraints.

10. Conclusion: Coherence as an Aesthetic Primitive

This paper began with a dissatisfaction that is easy to feel and difficult to articulate: the usual visual culture of Mandelbrot and Julia imagery, for all its beauty, often treats mathematics as an ornament engine. The picture is a harvest of detail, a spectacle of boundary complexity, a proof-by-awe that "infinity lives here." That aesthetic is not wrong, but it is incomplete. It tends to hide the deeper question that makes these systems intellectually and spiritually compelling: what does it mean for behavior to depend on parameters, and what does it mean for that dependence to be representable without lying?

The approach developed here proposes a different primitive: coherence. Instead of optimizing for maximal visual yield, we optimize for structure-preserving assignment. Instead of filling every pixel, we allow refusal. Instead of using color as mood, we design it as a disciplined translation. And instead of relying on zoom to solve scale, we build scale into the canvas as an atlas of restrictions. The resulting images are not "fractal pictures" in the usual sense. They are records of where a chosen witness can speak coherently and where it cannot.

10.1 Thesis restated: the artwork is a constrained section, not a fractal picture

The core thesis can be stated plainly.

We treat the parameter plane as a base. We attach to each parameter ca compact behavioral signature $\sigma(c)$ derived from the dynamics of $f_c(z) = z^2 + c$. This attachment is the functor-like assignment $F: c \mapsto \sigma(c)$. We then render $\sigma(c)$ as a glyph in a perceptual codomain (stroke geometry and optionally color). Crucially, we do not draw everywhere. We impose a coherence gate: when small changes in parameter produce changes in signature that exceed an admissible threshold, the mark is suppressed. The image is therefore a partial section: it exists where local representability holds under the chosen signature and scale, and it becomes silent where representability breaks.

This is the critical shift. A conventional fractal image claims ubiquity: every parameter yields a value; every point is colored; continuity is suggested everywhere. Our construction claims boundedness: it speaks only where it can do so coherently. It makes discontinuity visible as a boundary of honest representation rather than as a decorative shimmer.

In categorical terms, the canvas is a witness of where the base morphisms lift into admissible fiber morphisms under the chosen assignment. The presence of a mark is a local commutation claim: "this small move in the base did not produce a disproportionate move in the fiber, so a coherent glyph is permitted." The absence of a mark is a refusal to assert that claim. The work is therefore fundamentally different in kind from a depiction of the Mandelbrot set as a filled region. It is a constrained section through a derived space of parameter–behavior pairs, rendered with an explicit ethic of coherence.

10.2 What is gained: legibility, restraint, and meaningful discontinuity

Optimizing for coherence rather than spectacle yields three gains that are both aesthetic and methodological.

Legibility.
When marks are governed by a small vocabulary and by coherence constraints, the viewer can learn the image. The work becomes readable rather than merely impressive. Large forms—plumes, annuli, spines—can be interpreted as stable features of the mapping rather than as incidental artifacts of a palette. Color, when used under the natural-transformation motif, reinforces this legibility by acting as a second witness rather than as an emotional overlay. The viewer is not trapped in "infinite texture." The viewer is invited into a grammar: direction, scale, stability, failure.

Restraint.
Restraint is not the absence of content; it is the refusal to overstate. Fractal art often overwhelms because it tries to show too much at once, producing a field where everything is equally "interesting." The coherence approach produces hierarchy. Stable regions are allowed to accumulate density; unstable regions tear; truly incoherent loci become silent. This restores contrast—not merely luminance contrast, but epistemic contrast. The work looks confident where it can be confident and quiet where it cannot. That quietness is not a defect. It is part of the meaning.

Meaningful discontinuity.
Perhaps the most important gain is the rehabilitation of discontinuity. In many generative aesthetics, discontinuity is treated as an artifact to be smoothed away. Here it becomes the

central evidence: a sign of boundary effects, instability, representational collapse, or numerical horizon limits. The viewer learns that some breaks are mathematical (the system is genuinely sensitive), some breaks are representational (the signature is too coarse), and some breaks are computational (finite iteration and precision impose limits). Instead of hiding these layers, the work makes them visible. This is a disciplined form of honesty that is rare in both art and visualization.

Together these gains produce a different kind of beauty: not the beauty of maximal ornament, but the beauty of a system that tells the truth about where it can speak and where it cannot. The image becomes an object of witnesshood rather than an object of seduction.

10.3 Future directions: richer fibers, better morphisms, and physical installations

The framework developed here is intentionally minimal, which means there is ample room for extension without abandoning the central ethic.

Richer fibers.
The behavioral signature $\sigma(c)$ can be enriched in principled ways. Period detection, distance estimation proxies, multiple critical orbits in multicritical families, orbit-trap statistics, compressibility measures, or even learned embeddings of orbit segments could serve as fiber objects. The key constraint is that richer fibers must not become an excuse for decoration. Each added component should carry interpretive weight and should be accompanied by a corresponding revision of the coherence gate. In categorical language: a richer fiber category should be paired with clearer morphisms, not merely with more dimensions.

Better morphisms.
The base morphisms used here are simple: adjacency on a sampling grid, radial restriction, angular stepping. More nuanced morphisms could encode meaningful parameter deformations: paths designed to cross known structural features, rays aligned with external angles, or morphisms defined by dynamical equivalence relations rather than by Euclidean proximity. On the fiber side, morphisms could be upgraded from "distance below threshold" to structured comparisons: monotone changes in escape potentials, stable invariants under conjugacy, or quantified disagreements between orbit statistics. This would make the "functor-like" language less metaphorical and more operational: coherence would be tested along morphisms that reflect actual dynamical structure rather than generic geometric closeness.

Physical installations.
The radial atlas was introduced because the work wants to be walked. This points naturally toward installation-scale futures: canvases large enough to be read at multiple distances, wrapped bands on cylindrical surfaces, multi-panel atlases that behave like galleries of charts, or even architectural placements where the viewer's movement is choreographed as a restriction

journey. One can imagine a room-scale “parameter chapel” where each wall is a chart around a different basepoint c_0, and walking around the room becomes a traversal of parameter space. In such settings, the conceptual payoff becomes stronger: the viewer does not merely look at multiscale structure; the viewer inhabits it.

Across all these directions, the non-negotiable core remains the same. Coherence is the primitive. The artwork must earn its marks. It must treat silence as a truthful outcome. It must discipline color as translation rather than mood. And it must preserve the distinction between what is computed, what is represented, and what is claimed.

If fractal art has often been the celebration of infinite complexity, this approach proposes a complementary celebration: the beauty of constraint, the dignity of refusal, and the emergence of form from a strict commitment to coherence-preserving assignment. That is the aesthetic thesis—and the methodological invitation—of this work.

References

Complex dynamics, Mandelbrot/Julia foundations

1. Douady, A., and Hubbard, J. H. *Étude dynamique des polynômes complexes.* Publications Mathématiques d'Orsay 84-02 (1984) and 85-04 (1985). WorldCat+1
2. Douady, A., and Hubbard, J. H. *Exploring the Mandelbrot Set (The Orsay Notes).* (Notes presenting results from 1981–82; widely circulated in PDF form.) Pi Math Cornell
3. Douady, A., and Hubbard, J. H. "On the dynamics of polynomial-like mappings." *Annales scientifiques de l'École Normale Supérieure* (4) 18(2) (1985): 287–343. Numdam
4. Milnor, J. *Dynamics in One Complex Variable.* 3rd ed., Princeton University Press, 2006. Barnes & Noble+1

Fractals, visualization, and "fractal art" technique lineage

5. Peitgen, H.-O., and Richter, P. H. *The Beauty of Fractals: Images of Complex Dynamical Systems.* Springer, 1986. Springer+1

6. Peitgen, H.-O., Jürgens, H., and Saupe, D. *Chaos and Fractals: New Frontiers of Science.* Springer, 1992. Springer+1

7. Pickover, C. A. *Computers, Pattern, Chaos and Beauty: Graphics from an Unseen World.* St. Martin's Press, 1990. Amazon+1

8. Barnsley, M. F. *Fractals Everywhere.* 2nd ed., Morgan Kaufmann / Academic Press Professional, 1993. Amazon+1

Category theory and "coherence / functor / natural transformation" background

9. Mac Lane, S. *Categories for the Working Mathematician.* 2nd ed., Springer (Graduate Texts in Mathematics 5), 1998. Amazon+1

10. Awodey, S. *Category Theory.* 2nd ed., Oxford University Press (Oxford Logic Guides), 2010. Google Books+1

11. Riehl, E. *Category Theory in Context.* Dover Publications, 2016. Barnes & Noble+1

12. Leinster, T. *Basic Category Theory.* Cambridge University Press, 2014. Cambridge University Press & Assessment+1

13. Mac Lane, S., and Moerdijk, I. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory.* Springer, 1992. Amazon+1

14. Borceux, F. *Handbook of Categorical Algebra, Volume 1: Basic Category Theory.* Cambridge University Press, 1994. Amazon+1

15. Street, R. "Fibrations in bicategories." *Cahiers de Topologie et Géométrie Différentielle Catégoriques* 21(2) (1980): 111–160. EuDML+1

Computer graphics / color mapping background (HSV as an implementation codomain)
16. Hughes, J. F., van Dam, A., McGuire, M., Sklar, D. F., Foley, J. D., Feiner, S. K., and Akeley, K. *Computer Graphics: Principles and Practice.* 3rd ed., Addison-Wesley Professional, 2013. Pearson+1

The author has published over 4,700 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.

Python source code.

```
#!/usr/bin/env python3
"""
Coherence Atlas (Radial Parameterization) for the quadratic family f_c(z)=z^2+c.

Generates a single-canvas, multi-scale "radial atlas" centered at basepoint c0 in parameter
space.
Each sample parameter c is assigned a compact signature sigma(c) derived from the critical
orbit.
A coherence gate suppresses marks where local variation in sigma is too large. Marks are drawn
as
short strokes (optionally curved), with optional HSV colorization.

Outputs:
- PNG image (default)
- JSON sidecar with parameters used (for reproducibility)

No external dependencies beyond: numpy, pillow.

Example:
  python coherence_atlas.py --out atlas.png --size 6000 --nr 1400 --ntheta 2400 \
    --c0_re -0.1 --c0_im 0.65 --rho_max 0.35 --rho_min 0.0008 --alpha 2.2 \
    --N 800 --R 2.0 --delta 0.25 --color 1
"""

from __future__ import annotations
```

```
import argparse
import json
import math
import time
from dataclasses import dataclass, asdict
from typing import Dict, Tuple, Optional

import numpy as np
from PIL import Image, ImageDraw


# ----------------------------
# Utilities
# ----------------------------

def clamp(x: float, lo: float, hi: float) -> float:
    return lo if x < lo else hi if x > hi else x

def safe_log(x: float, eps: float = 1e-12) -> float:
    return math.log(max(x, eps))

def hsv_to_rgb_u8(h: float, s: float, v: float) -> Tuple[int, int, int]:
    """
    HSV in [0,1] -> RGB 8-bit.
    """
    h = (h % 1.0) * 6.0
```

```
    i = int(h)
    f = h - i
    p = v * (1.0 - s)
    q = v * (1.0 - s * f)
    t = v * (1.0 - s * (1.0 - f))
    if i == 0:
        r, g, b = v, t, p
    elif i == 1:
        r, g, b = q, v, p
    elif i == 2:
        r, g, b = p, v, t
    elif i == 3:
        r, g, b = p, q, v
    elif i == 4:
        r, g, b = t, p, v
    else:
        r, g, b = v, p, q
    return (int(clamp(r, 0.0, 1.0) * 255.0),
            int(clamp(g, 0.0, 1.0) * 255.0),
            int(clamp(b, 0.0, 1.0) * 255.0))


def norm2(v: np.ndarray) -> float:
    return float(np.sqrt(np.dot(v, v)))


# ----------------------------
```

```
# Signature
# ----------------------------

@dataclass
class Params:
    out: str
    sidecar: str
    size: int
    margin: int
    nr: int
    ntheta: int

    c0_re: float
    c0_im: float

    rho_max: float
    rho_min: float
    alpha: float

    N: int
    R: float

    eta_sym: float

    # Gate thresholds
    delta: float            # signature-delta threshold for coherence
```

```
    gate_inside_only: bool    # if True, apply coherence gate only when both points classified
bounded

    # Glyph controls (in pixels)
    L_min: float
    L_max: float
    T_min: float
    T_max: float
    curvature_scale: float
    orient_mode: str         # "theta" or "theta_plus"

    # Color controls
    color: bool
    hue_mode: str            # "theta" or "signature"
    sat_from: str            # "escape" or "lambda"
    val_from: str            # "bounded" or "escape" or "lambda"
    bounded_val: float
    bg: str                  # "black" or "white"
    alpha: int               # 0..255

    # Normalization clamps (for stability)
    escape_norm_max: float
    lambda_norm_max: float


def rho_schedule(r: float, rho_min: float, rho_max: float, alpha: float) -> float:
```

```python
    # rho(r) decreases as r increases: coarse in center, fine near edge
    return rho_min + (rho_max - rho_min) * ((1.0 - r) ** alpha)


def critical_orbit_signature(c: complex, N: int, R: float, eta_sym: float,
                             escape_norm_max: float, lambda_norm_max: float) -> np.ndarray:
    """
    sigma(c) = [b, e_norm, lam_norm, s]
    b: boundedness proxy in {0,1} at depth N
    e_norm: normalized escape-rate proxy (0 if bounded)
    lam_norm: normalized expansion proxy (mean log|2 z_k|)
    s: symmetry cue (near real axis)

    Returns float vector length 4.
    """
    z = 0.0 + 0.0j
    escaped = False
    tau = N
    z_tau = 0.0 + 0.0j

    # accumulate log|2 z|
    lam_sum = 0.0
    m = 0

    for n in range(N):
        # derivative magnitude proxy at current z
```

```
    # add eps in magnitude for stability near z=0
    lam_sum += safe_log(abs(2.0 * z) + 1e-12)
    m += 1

    # iterate
    z = z * z + c

    if (not escaped) and (abs(z) > R):
        escaped = True
        tau = n + 1
        z_tau = z
        break

b = 0.0 if escaped else 1.0

# escape-rate proxy
if escaped:
    e = safe_log(abs(z_tau)) / max(tau, 1)
    e_norm = clamp(e / escape_norm_max, 0.0, 1.0)
else:
    e_norm = 0.0

# expansion proxy: mean log|2 z_k|
lam = lam_sum / max(m, 1)
# map to [0,1] by clamping symmetric range around 0
lam_norm = clamp((lam + lambda_norm_max) / (2.0 * lambda_norm_max), 0.0, 1.0)
```

```python
    # symmetry cue: near real axis
    s = 1.0 if abs(c.imag) < eta_sym else 0.0

    return np.array([b, e_norm, lam_norm, s], dtype=np.float32)


# ----------------------------
# Glyph rendering
# ----------------------------

def stroke_params_from_sigma(theta: float, sigma: np.ndarray, p: Params) -> Tuple[float, float,
float, float]:
    """
    Returns (L, T, kappa, orient) where:
      L: length in px
      T: thickness in px
      kappa: curvature coefficient in px (controls bend amplitude)
      orient: stroke orientation angle in radians
    """
    b, e_norm, lam_norm, s = float(sigma[0]), float(sigma[1]), float(sigma[2]), float(sigma[3])

    # "energy" controls
    energy = e_norm
    # expansion centered around 0: lam_norm in [0,1] -> [-1,1]
    lam_center = 2.0 * lam_norm - 1.0
```

```
    # length: bounded marks slightly longer, escaping marks respond to energy
    L = p.L_min + (p.L_max - p.L_min) * (0.35 * b + 0.65 * energy)
    # thickness: energy + expansion magnitude
    T = p.T_min + (p.T_max - p.T_min) * clamp(0.6 * energy + 0.4 * abs(lam_center), 0.0, 1.0)

    # curvature amplitude: from expansion (signed)
    kappa = p.curvature_scale * lam_center

    # orientation
    if p.orient_mode == "theta":
        orient = theta
    else:
        # slight twist based on symmetry cue and expansion
        orient = theta + (0.7 * (s - 0.5)) + 0.35 * lam_center

    return L, T, kappa, orient


def color_from_sigma(theta: float, sigma: np.ndarray, p: Params) -> Tuple[int, int, int, int]:
    """
    Produces RGBA color for a mark.
    Uses HSV mapping with interpretable channels.
    """
    b, e_norm, lam_norm, s = float(sigma[0]), float(sigma[1]), float(sigma[2]), float(sigma[3])
```

```
if p.bg == "white":
    base_bg = 1.0
else:
    base_bg = 0.0

if not p.color:
    # grayscale: use value channel only
    if p.val_from == "bounded":
        v = p.bounded_val if b > 0.5 else clamp(0.3 + 0.7 * e_norm, 0.0, 1.0)
    elif p.val_from == "lambda":
        v = clamp(0.15 + 0.85 * lam_norm, 0.0, 1.0)
    else:
        v = clamp(0.15 + 0.85 * e_norm, 0.0, 1.0)

    # invert if background is white
    if p.bg == "white":
        v = 1.0 - 0.85 * v
    g = int(clamp(v, 0.0, 1.0) * 255.0)
    return (g, g, g, p.alpha)

# hue
if p.hue_mode == "theta":
    h = (theta / (2.0 * math.pi)) % 1.0
else:
    # hue from signature: blend lambda and escape
    h = (0.65 * lam_norm + 0.35 * e_norm) % 1.0
```

```
    # saturation
    if p.sat_from == "lambda":
        s_val = clamp(0.2 + 0.8 * abs(2.0 * lam_norm - 1.0), 0.0, 1.0)
    else:
        s_val = clamp(0.15 + 0.85 * e_norm, 0.0, 1.0)

    # value
    if p.val_from == "bounded":
        v = p.bounded_val if b > 0.5 else clamp(0.25 + 0.75 * e_norm, 0.0, 1.0)
    elif p.val_from == "lambda":
        v = clamp(0.20 + 0.80 * lam_norm, 0.0, 1.0)
    else:
        v = clamp(0.20 + 0.80 * e_norm, 0.0, 1.0)

    # adjust for white background (keep contrast)
    if p.bg == "white":
        v = 1.0 - 0.85 * v
        s_val = clamp(0.10 + 0.70 * s_val, 0.0, 1.0)

    r, g, b2 = hsv_to_rgb_u8(h, s_val, v)
    return (r, g, b2, p.alpha)


def draw_stroke(draw: ImageDraw.ImageDraw, x: float, y: float, L: float, T: float,
                kappa: float, orient: float, rgba: Tuple[int, int, int, int]) -> None:
```

```
"""
Draw a short curve-like stroke: approximate quadratic Bezier with 3 points.
For kappa ~ 0 -> straight segment.
"""
dx = math.cos(orient)
dy = math.sin(orient)

x0 = x - 0.5 * L * dx
y0 = y - 0.5 * L * dy
x2 = x + 0.5 * L * dx
y2 = y + 0.5 * L * dy

# Perpendicular for curvature control point
px = -dy
py = dx
x1 = x + kappa * px
y1 = y + kappa * py

# Sample curve
steps = 10
pts = []
for t in np.linspace(0.0, 1.0, steps):
    # Quadratic Bezier: (1-t)^2 P0 + 2(1-t)t P1 + t^2 P2
    a = (1.0 - t) * (1.0 - t)
    b = 2.0 * (1.0 - t) * t
    c = t * t
```

```python
        xt = a * x0 + b * x1 + c * x2
        yt = a * y0 + b * y1 + c * y2
        pts.append((xt, yt))

    # Pillow's line uses integer width; clamp to >= 1
    width = max(1, int(round(T)))
    draw.line(pts, fill=rgba, width=width, joint="curve")


# ----------------------------
# Main render
# ----------------------------

def render(p: Params) -> Dict:
    t0 = time.time()

    size = int(p.size)
    margin = int(p.margin)
    R_pix = (size // 2) - margin
    cx = size / 2.0
    cy = size / 2.0

    # Background
    if p.bg == "white":
        img = Image.new("RGBA", (size, size), (255, 255, 255, 255))
    else:
```

```
    img = Image.new("RGBA", (size, size), (0, 0, 0, 255))
  draw = ImageDraw.Draw(img, "RGBA")

  c0 = complex(p.c0_re, p.c0_im)

  # Precompute radial and angular arrays
  r_vals = np.linspace(0.0, 1.0, p.nr, dtype=np.float32)
  theta_vals = np.linspace(0.0, 2.0 * math.pi, p.ntheta, endpoint=False, dtype=np.float32)

  # Optional: store previous signatures along theta for each radius
  # We'll gate along theta neighbors (wrap-around), and optionally along radial neighbors.
  prev_sigma_theta: Optional[np.ndarray] = None
  prev_sigma_rad = np.zeros((p.ntheta, 4), dtype=np.float32)  # last radius signatures (for radial
gating)

  suppressed_count = 0
  drawn_count = 0

  for i, r in enumerate(r_vals):
    rho = rho_schedule(float(r), p.rho_min, p.rho_max, p.alpha)

    # compute this ring signatures
    sig_ring = np.zeros((p.ntheta, 4), dtype=np.float32)

    # First pass: compute signatures
    for j, theta in enumerate(theta_vals):
```

```
    c = c0 + rho * complex(math.cos(float(theta)), math.sin(float(theta)))
    sig = critical_orbit_signature(
        c=c,
        N=p.N,
        R=p.R,
        eta_sym=p.eta_sym,
        escape_norm_max=p.escape_norm_max,
        lambda_norm_max=p.lambda_norm_max,
    )
    sig_ring[j, :] = sig

# Second pass: gating + render
for j, theta in enumerate(theta_vals):
    sig = sig_ring[j, :]

    # neighbors (theta wrap)
    sig_prev_theta = sig_ring[j - 1, :]
    # neighbor (radial)
    sig_prev_r = prev_sigma_rad[j, :] if i > 0 else None

    # gate measure: max of theta and radial deltas (where available)
    d_theta = norm2(sig - sig_prev_theta)
    d_rad = norm2(sig - sig_prev_r) if sig_prev_r is not None else 0.0
    d = max(d_theta, d_rad)

    # decide whether to apply gate only inside (bounded) or everywhere
```

```
        apply_gate = True
        if p.gate_inside_only:
            b = float(sig[0])
            b_prev = float(sig_prev_theta[0])
            b_prev_r = float(sig_prev_r[0]) if sig_prev_r is not None else b
            # apply gate only if all are bounded proxies
            apply_gate = (b > 0.5) and (b_prev > 0.5) and (b_prev_r > 0.5)

        if apply_gate and (d > p.delta):
            suppressed_count += 1
            continue

        # Canvas position
        x = cx + R_pix * float(r) * math.cos(float(theta))
        y = cy + R_pix * float(r) * math.sin(float(theta))

        # Glyph params
        L, T, kappa, orient = stroke_params_from_sigma(float(theta), sig, p)
        rgba = color_from_sigma(float(theta), sig, p)

        draw_stroke(draw, x, y, L, T, kappa, orient, rgba)
        drawn_count += 1

    prev_sigma_rad = sig_ring

# Disk mask (optional): keep only within radius R_pix (soft edge not applied here)
```

```
    # If you want corners blank in a square canvas, uncomment:
    #
    # mask = Image.new("L", (size, size), 0)
    # mdraw = ImageDraw.Draw(mask)
    # mdraw.ellipse((cx - R_pix, cy - R_pix, cx + R_pix, cy + R_pix), fill=255)
    # img.putalpha(mask)

    t1 = time.time()

    meta = {
        "runtime_seconds": t1 - t0,
        "counts": {
            "drawn": int(drawn_count),
            "suppressed": int(suppressed_count),
            "total_samples": int(p.nr * p.ntheta),
        },
        "params": asdict(p),
    }
    return img, meta


def parse_args() -> Params:
    ap = argparse.ArgumentParser(description="Generate coherence atlas fractal-art image.")
    ap.add_argument("--out", type=str, default="coherence_atlas.png", help="Output PNG
filename.")
    ap.add_argument("--sidecar", type=str, default="", help="Output JSON sidecar filename
(default: out + .json).")
```

```
    ap.add_argument("--size", type=int, default=6000, help="Image size (square) in pixels.")
    ap.add_argument("--margin", type=int, default=120, help="Margin (pixels) around disk.")

    ap.add_argument("--nr", type=int, default=1400, help="Radial samples.")
    ap.add_argument("--ntheta", type=int, default=2400, help="Angular samples.")

    ap.add_argument("--c0_re", type=float, default=-0.10, help="Basepoint c0 real.")
    ap.add_argument("--c0_im", type=float, default=0.65, help="Basepoint c0 imag.")

    ap.add_argument("--rho_max", type=float, default=0.35, help="Max neighborhood radius at center.")
    ap.add_argument("--rho_min", type=float, default=0.0008, help="Min neighborhood radius near edge.")
    ap.add_argument("--alpha", type=float, default=2.2, help="Neighborhood shrink exponent.")

    ap.add_argument("--N", type=int, default=800, help="Max iterations for critical orbit.")
    ap.add_argument("--R", type=float, default=2.0, help="Escape radius.")

    ap.add_argument("--eta_sym", type=float, default=0.03, help="Symmetry cue threshold on |Im(c)|.")

    ap.add_argument("--delta", type=float, default=0.25, help="Coherence gate threshold on ||delta sigma||.")
    ap.add_argument("--gate_inside_only", type=int, default=1, help="1: gate only when bounded proxy; 0: gate everywhere.")
```

```
    ap.add_argument("--L_min", type=float, default=1.6, help="Min stroke length (px).")
    ap.add_argument("--L_max", type=float, default=7.5, help="Max stroke length (px).")
    ap.add_argument("--T_min", type=float, default=1.0, help="Min stroke thickness (px).")
    ap.add_argument("--T_max", type=float, default=3.2, help="Max stroke thickness (px).")
    ap.add_argument("--curvature_scale", type=float, default=4.0, help="Curvature scale (px).")
    ap.add_argument("--orient_mode", type=str, default="theta_plus", choices=["theta",
"theta_plus"],
                    help="Stroke orientation mode.")

    ap.add_argument("--color", type=int, default=1, help="1: color (HSV), 0: grayscale.")
    ap.add_argument("--hue_mode", type=str, default="theta", choices=["theta", "signature"],
help="Hue mapping.")
    ap.add_argument("--sat_from", type=str, default="escape", choices=["escape", "lambda"],
help="Saturation mapping.")
    ap.add_argument("--val_from", type=str, default="bounded", choices=["bounded", "escape",
"lambda"], help="Value mapping.")
    ap.add_argument("--bounded_val", type=float, default=0.62, help="Value used for bounded
regime when val_from=bounded.")
    ap.add_argument("--bg", type=str, default="black", choices=["black", "white"],
help="Background color.")
    ap.add_argument("--alpha_mark", type=int, default=110, help="Mark alpha (0..255).")

    ap.add_argument("--escape_norm_max", type=float, default=0.9, help="Clamp scale for
escape proxy normalization.")
    ap.add_argument("--lambda_norm_max", type=float, default=2.0, help="Clamp half-range for
lambda normalization.")

    args = ap.parse_args()
```

```
out = args.out
sidecar = args.sidecar.strip()
if not sidecar:
    sidecar = out + ".json"

return Params(
    out=out,
    sidecar=sidecar,
    size=args.size,
    margin=args.margin,
    nr=args.nr,
    ntheta=args.ntheta,
    c0_re=args.c0_re,
    c0_im=args.c0_im,
    rho_max=args.rho_max,
    rho_min=args.rho_min,
    alpha=args.alpha,
    N=args.N,
    R=args.R,
    eta_sym=args.eta_sym,
    delta=args.delta,
    gate_inside_only=bool(args.gate_inside_only),
    L_min=args.L_min,
    L_max=args.L_max,
    T_min=args.T_min,
```

```
        T_max=args.T_max,
        curvature_scale=args.curvature_scale,
        orient_mode=args.orient_mode,
        color=bool(args.color),
        hue_mode=args.hue_mode,
        sat_from=args.sat_from,
        val_from=args.val_from,
        bounded_val=args.bounded_val,
        bg=args.bg,
        alpha=int(clamp(args.alpha_mark, 0, 255)),
        escape_norm_max=args.escape_norm_max,
        lambda_norm_max=args.lambda_norm_max,
    )


def main() -> None:
    p = parse_args()
    img, meta = render(p)

    # Save image
    img.save(p.out, format="PNG", optimize=True)

    # Save sidecar metadata
    with open(p.sidecar, "w", encoding="utf-8") as f:
        json.dump(meta, f, indent=2, sort_keys=True)
```

```python
    print(f"Wrote: {p.out}")
    print(f"Wrote: {p.sidecar}")
    print(f"Drawn marks: {meta['counts']['drawn']}, suppressed: {meta['counts']['suppressed']}, "
          f"total samples: {meta['counts']['total_samples']}")
    print(f"Runtime seconds: {meta['runtime_seconds']:.2f}")


if __name__ == "__main__":
    main()
```