

To Plot a Stone with Six Birds: A Geometry is A Theory

Ioannis Tsiokos¹

5 February 2026

¹Automorph Inc., 1207 Delaware Ave #4131, Wilmington, DE 19806, USA

Abstract

Geometry is usually treated as a background container: points, distances, and curvature are assumed and then modeled. From the Six Birds Theory (SBT) perspective, this conflates the object with the interface: we lack an auditable way to *derive* space-like structure as a stable closure artifact of finite observers acting on micro-dynamics, or to distinguish genuine geometric signal from packaging artifacts. We present an SBT-native construction in which *points* are packaged macro-states (equivalence classes under a lens) and *distance* is optimized accounting: costs derived from induced macro dynamics are composed into protocols and minimized over paths. Operationally, we implement a reproducible pipeline (substrate \rightarrow lens ladder \rightarrow closed macro kernel \rightarrow cost metric) together with coherence audits (idempotence defect, prototype stability, route mismatch, inter-scale distortion, and connectivity) and a curvature diagnostic based on loop holonomy. Across controlled substrates, we obtain three checkable results: (i) on an isotropic grid, the induced metric is connected and remains coherent across refinement (bounded closure and distortion audits), while constraints (anisotropic gating) induce a systematic deformation of the emergent metric; (ii) under a fixed deterministic holonomy protocol, loop residue is near-zero on the plane-like grid (median 0.0479) but shifts strongly upward on a sphere-like substrate (median 0.5980), separating flat from curved regimes; (iii) defining cost as negative log transition probability under staged isotropic diffusion yields an approximately quadratic, separable accounting law in which the median Pythagorean residual drops from 33.19 at $\tau=4$ to 0.0586 at $\tau=128$, while a Manhattan (L1) control produces diamond contours and does not exhibit this collapse. These results provide a falsification-first, reproducible workflow for comparing emergent geometries (flat versus curved versus constrained versus fractal) without assuming coordinates, and for attaching each claim to audits and explicit controls. We emphasize that stronger claims (e.g., continuum-limit convergence to manifolds or exact curvature tensors) are not established here; our conclusions rely on diagnostic proxies under controlled substrates with documented failure modes. This work extends the SBT foundation [Tsiokos \[2026a\]](#) and complements our prior treatment of emergent mathematics [Tsiokos \[2026b\]](#).

keywords: geometry; emergence calculus; closure; audit; packaging

1 Introduction

To *plot* an object is to treat it as if it already lives in a space: we assign locations, measure distances, and speak of surfaces and curvature as though these were primitive facts. Yet in practice we never access “space itself”—we access only interactions, constraints, and limited interfaces: what an observer can reliably distinguish, compress, and act upon. Six Birds Theory (SBT) [Tsiokos \[2026a\]](#)

takes this seriously: points, distances, and geometric laws are not axioms but *closure artifacts* of stable compression under the six primitives.

This paper asks a concrete question: *how do you plot a stone if you are not allowed to assume space?* A stone has microstructure (grains, pores, microfractures) and dynamics (how disturbances propagate, how configurations change). An agent with finite bandwidth cannot retain all micro-detail; it must package. If repeated packaging produces a stable notion of “the same” (macro states), if feasible moves compose into protocols, and if the costs of effecting or distinguishing moves become coherent across refinement, then a space-like layer becomes available. In one line:

Geometry is what you get when repeated packaging produces a stable notion of nearby, composition of moves, and cost of moving, and that structure closes under refinement.

The narrative of this paper parallels our companion work on emergent mathematics [Tsiokos \[2026b\]](#): there we study how counting-like structure becomes stable under SBT closure; here we study how *plotting-like* structure becomes stable. The difference is one of emphasis. “Counting” stabilizes a notion of quantity; “plotting” stabilizes a notion of *location and transport*. Both are instances of the same emergence calculus: staging (multi-scale refinement) and packaging (quotienting by indistinguishability) are forced by finite interfaces, while accounting turns feasibility into a ledger that can be optimized over protocols.

What we do (high level). We start from a finite substrate of microstates with micro-dynamics (a Markov kernel). At each stage we construct a lens (a packaging map) and prototypes that realize closure, yielding a macro kernel. We then define distance from accounting: transition likelihood induces cost, and distance is the minimal cost of a protocol connecting macro states. A geometric layer is said to exist when this induced metric structure is coherent across refinement, as quantified by a small set of diagnostics (idempotence and stability defects, route mismatch, inter-scale distortion, connectivity), and when curvature is detected as loop holonomy (protocol noncommutativity made geometric).

Reader map: exhibits. We emphasize falsification-first: each exhibit has a corresponding “where it breaks” analysis. The core experimental exhibits are:

- **E1 (plane-like):** on a grid substrate with isotropic local moves, the emergent metric is coherent across scales (low defects, connected macro graph, bounded distortions).
- **E2 (curvature as holonomy):** a loop-based holonomy diagnostic concentrates near zero on plane-like substrates but shifts upward on sphere-like substrates, revealing curvature as protocol residue.
- **E3 (fractal regime):** on a Sierpiński substrate, refinement does not smooth toward Euclidean neighborhoods but stabilizes in scale-space, separating fractal from smooth regimes.
- **E4 (constraints deform geometry):** imposing directional feasibility constraints induces anisotropic deformation of the emergent metric and its cross-scale coherence.
- **E5 (Pythagoras emergence):** using negative log transition probability as cost under isotropic staged diffusion, the stable accounting law becomes approximately quadratic and separable, yielding a Pythagorean structure; a Manhattan control fails as expected.

Reproducibility. All experiments are runnable from configuration files in the repository (including end-to-end harness runs and a canonical holonomy run). For writing stability, we also commit “run packs” under `docs/notes/runs/` containing the exact summary JSONs and plots referenced by the paper, together with paper-ready comparison figures and a quotables table. This anchors the narrative to concrete, auditable artifacts rather than ad hoc screenshots.

Code availability. The repository for this paper is available at:

- <https://github.com/ioannist/six-birds-space>

Roadmap. Section 2 recaps the six primitives and how they specialize to geometry. We then present the emergent geometry construction (points as equivalence classes; distance as optimized accounting), the diagnostic toolkit, the experimental pipeline, and the exhibits, followed by robustness and failure modes and minimal Lean anchors.

2 Six Birds Recap: how the primitives specialize to geometry

This paper applies Six Birds Theory (SBT) Tsiokos [2026a] to the emergence of space-like structure. We do not treat “geometry” as a background container. Instead, we treat it as a stable higher-layer theory: a description that becomes available when repeated compression makes notions of *indistinguishability* (points), *feasible transformation* (moves), and *accounting* (cost) coherent under refinement. The six primitives are unchanged; what changes is which stable objects they induce. (Compare with our earlier application to emergent mathematics Tsiokos [2026b].)

P1: Operator rewrite (closure). A layer is usable only if it closes: the dynamics can be written in the layer’s own variables without constantly reopening micro-detail. In this paper, P1 appears as the induced *macro dynamics* obtained after packaging and staging. Closure is not assumed; it is audited and quantified. When closure holds, it supports geometry-like operators (transport, diffusion, and in smooth regimes, differential structure).

P2: Constraints (feasibility). Constraints define what transitions and interventions are admissible. Geometrically, P2 shapes the reachable neighborhood structure and the anisotropy of motion: it decides which directions are “easy,” “hard,” or impossible. In our exhibits, isotropic local moves yield plane-like metrics, while directional gating deforms the emergent metric in a predictable way.

P3: Protocols and holonomy (composition matters). P3 is the primitive of ordered composition: doing *A* then *B* need not equal doing *B* then *A*. Geometry reads this as curvature: after packaging, *noncommutativity of local transport* becomes a stable loop residue (holonomy). When small loops return to their starting point misaligned, curvature has emerged as a macroscopic footprint of protocol order.

P4: Staging (multi-scale refinement). Staging introduces a ladder of resolutions and time scales. Here it is realized as (i) a stage parameter τ (how much micro-dynamics we allow before repackaging) and (ii) a family of lenses indexed by resolution. Geometry is said to exist only when the induced structure is coherent across this refinement ladder; fractal regimes correspond to scale-stable closure without smoothing.

| Primitive (Bird) | Role in emergent geometry (this paper) |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| P1 Operator rewrite | Closed macro dynamics and induced operators on packaged states; “geometry” exists only when the macro description closes and stabilizes under refinement. |
| P2 Constraints | Feasibility of moves; controls locality, anisotropy, reachable neighborhoods, and deformations of the emergent metric. |
| P3 Protocols / holonomy | Composition of moves; curvature as a stable residue of noncommuting local transports (loop holonomy). |
| P4 Staging | Resolution ladder and time-scale parameter τ ; enables refinement tests, smooth versus scale-stable (fractal) fixed points. |
| P5 Packaging | Points and regions as equivalence classes under indistinguishability; the effective state space X is created by quotienting microstates. |
| P6 Accounting | Distance as minimal cost of protocols; scaling exponents as “bits per scale”; Pythagoras as stable quadratic accounting in isotropic diffusion regimes. |

Table 1: How the six SBT primitives specialize to geometry and space in this paper.

P5: Packaging (quotienting indistinguishability). Packaging maps many microstates to the same macro state: it constructs the layer’s effective state space by collapsing distinctions the interface cannot reliably maintain. In geometric language, this is where “points” come from: a point is an equivalence class of microstates that are indistinguishable at the chosen resolution. Refinement corresponds to sharpening the quotient.

P6: Accounting (cost is real). Accounting assigns resource meaning to distinctions and transformations. In this paper, distance is not primarily a ruler; it is primarily a ledger. We define costs from transition likelihoods and obtain distances as minimal protocol cost (shortest paths) in the induced macro dynamics. In the isotropic staged-diffusion regime, this accounting becomes approximately quadratic and separable, yielding a Pythagorean structure as the stable law of cost.

The remaining sections make these correspondences concrete. We first define the packaging and closure constructions that yield a macro kernel and an induced cost metric, then introduce diagnostics that quantify coherence across refinement, and finally present reproducible exhibits separating plane-like, curved, anisotropic, and fractal regimes.

3 Core construction: from packaging to an emergent metric

We now present the minimal SBT-native construction that turns a substrate with micro-dynamics into a space-like layer with points and distances. The key move is to treat distance as *accounting* (P6) optimized over *protocols* (P3), rather than as a primitive ruler.

3.1 Substrate and micro-dynamics

Let Z be a finite set of microstates (in code: indices $0, \dots, n-1$). Let $\Delta(Z)$ denote the probability simplex over Z . We model micro-dynamics by a Markov kernel $P \in \mathbb{R}^{Z \times Z}$ that is row-stochastic [Norris \[1997\]](#). We adopt the row-vector convention: for $\mu \in \Delta(Z)$,

$$\mu_{t+1} = \mu_t P, \quad \mu_t \in \Delta(Z). \tag{1}$$

Staging (P4) enters via a discrete time-scale parameter $\tau \in \mathbb{N}$: the τ -step evolution is $\mu \mapsto \mu P^\tau$.

3.2 Packaging as a lens: points are indistinguishability classes (P5)

A *lens* (packaging map) is a function

$$f : Z \rightarrow X, \quad (2)$$

where X is a finite set of macro labels (“macro states”). In geometric language, elements of X will play the role of *points* at this resolution: a point is not primitive; it is a quotient class created by collapsing microstates that are indistinguishable under the interface induced by f .

We represent the lens by the indicator (coarse) matrix $C \in \{0, 1\}^{Z \times X}$:

$$C[z, x] := \mathbf{1}\{f(z) = x\}. \quad (3)$$

This induces a pushforward map on distributions (the coarse-graining operator)

$$Q_f : \Delta(Z) \rightarrow \Delta(X), \quad Q_f(\mu) := \mu C. \quad (4)$$

3.3 Prototypes as lifts: closure representatives (P1, P5)

To close dynamics at the macro level we choose, for each $x \in X$, a *prototype* $u_x \in \Delta(Z)$ that serves as a representative micro-distribution for the macro state x . Collect these rows into a matrix $U \in \mathbb{R}^{X \times Z}$ with

$$U[x, z] := u_x(z), \quad \sum_{z \in Z} U[x, z] = 1, \quad U[x, z] \geq 0. \quad (5)$$

This induces a lift map

$$U_f : \Delta(X) \rightarrow \Delta(Z), \quad U_f(\nu) := \nu U. \quad (6)$$

A basic consistency condition is that macro states are recovered after lifting:

$$Q_f \circ U_f = \text{id}_{\Delta(X)}. \quad (7)$$

In words: if we choose a macro label distribution ν , lift it to micro via prototypes, and repackage, we recover ν . (Our implementations enforce this numerically by construction.)

3.4 The induced macro kernel

With staging τ , lens f , and prototypes U , we obtain an induced macro Markov kernel \hat{P} on X by evolving in micro, then packaging:

$$\hat{P} := U P^\tau C \in \mathbb{R}^{X \times X}. \quad (8)$$

Under the row-vector convention, macro distributions $\nu \in \Delta(X)$ evolve as

$$\nu_{t+1} = \nu_t \hat{P}. \quad (9)$$

This is the closure move (P1) in its simplest form: we rewrite micro-dynamics in macro variables by inserting (i) packaging and (ii) a chosen lift back to micro.

3.5 Distance is accounting (P6): costs from likelihood

We now turn macro dynamics into a costed move system. Let $\eta > 0$ be a small smoothing constant. Define a directed per-step cost by negative log likelihood Shannon [1948]:

$$c_\eta(x \rightarrow y) := -\log(\hat{P}(x, y) + \eta). \quad (10)$$

To obtain an undirected cost suitable for a symmetric metric, we may symmetrize the kernel (or the cost). A common choice is

$$W := \frac{1}{2}(\hat{P} + \hat{P}^\top), \quad c_\eta^{\text{sym}}(x, y) := -\log(W(x, y) + \eta). \quad (11)$$

Key principle: *Distance is accounting.* A pair of macro states is “near” if it is cheap (low ledger cost) to move influence between them under the induced dynamics; the pair is “far” if doing so is expensive.

3.6 Distance is optimized protocol cost (P3): shortest paths

Let G be the directed (or undirected) weighted graph on vertex set X with edge weights given by c_η (or c_η^{sym}). A *protocol* is a finite path $\gamma = (x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_k)$ in this graph. The protocol cost is additive:

$$\text{Cost}_\eta(\gamma) := \sum_{i=0}^{k-1} c_\eta(x_i \rightarrow x_{i+1}). \quad (12)$$

We define the induced distance as the minimal protocol cost:

$$d_\eta(x, y) := \inf_{\gamma: x \rightsquigarrow y} \text{Cost}_\eta(\gamma). \quad (13)$$

In computation we realize (13) via all-pairs shortest paths on the macro graph. In the appendices we record minimal Lean anchors that connect this construction to standard metric facts (triangle inequality via path concatenation and metric separation via quotienting by zero distance).

3.7 Mathematical status: extended (pseudo-)metrics, directed costs, and quotients

Because our distances are computed as minimal path cost on a weighted graph, the basic mathematical structure is determined by construction rather than by assumption.

Nonnegativity and triangle inequality. Edge costs defined from (10) are nonnegative. Therefore the shortest-path construction (13) yields a (possibly extended) distance function satisfying $d_\eta(x, x) = 0$ and the triangle inequality by path concatenation.

Directed versus undirected. If we use the directed cost $c_\eta(x \rightarrow y)$, then the resulting $d_\eta(x, y)$ need not equal $d_\eta(y, x)$; mathematically this is a directed shortest-path distance (a quasi-metric or directed pseudo-distance). In this paper we default to an undirected geometry by symmetrizing before converting to costs, yielding a symmetric distance suitable for comparison with familiar geometric regimes. Directed geometries induced by asymmetric feasibility (constraints) are natural in the SBT framework but are deferred to future work.

Extended distances and disconnection. If the macro move graph is disconnected, there exist pairs with no connecting protocol; in that case $d_\eta(x, y) = \infty$. We treat non-finite distances as a failure mode of the proposed geometric layer and explicitly record the count of non-finite entries (connectivity audit).

Pseudometric versus metric and separation. In general, a shortest-path construction yields a pseudometric: it is possible in principle that distinct points have zero distance if there exist zero-cost paths. Our likelihood-based costs typically make nontrivial zero-cost paths rare (they would require essentially deterministic transitions), but the correct mathematical statement is still a pseudo- or extended metric unless separation is verified. A standard way to obtain a metric is to quotient by the zero-distance relation ($x \sim y$ iff $d_\eta(x, y) = 0$), yielding a metric on the separation quotient. We include minimal Lean anchors formalizing (i) triangle inequality for shortest-path distance and (ii) the separation quotient metric construction (Appendix A).

3.8 When do we say a geometric layer exists? (Coherence schema)

A single computed distance matrix is not yet “geometry.” Geometry is a claim of *coherence under refinement*: the packaged points, induced dynamics, and cost structure must stabilize across staging and across a ladder of lenses. We use the following practical schema.

Given: A refinement ladder of lenses f_0, f_1, \dots, f_L (coarse to fine), corresponding coarse matrices C_j , prototype matrices U_j , a staging parameter τ , and induced macro kernels $\hat{P}_j = U_j P^\tau C_j$ with induced distances $d_\eta^{(j)}$.

We say a space-like geometric layer is present on this ladder when the following hold:

1. **Closure is nearly idempotent (P1/P5):** The package–evolve–repackage operator is close to idempotent at each scale (small idempotence defect δ_{τ, f_j}).
2. **Macro points are stable (P4/P5):** Prototypes do not drift strongly under the induced closure (small stability defect $s_{\tau, f_j}(x)$ on average and in the worst case).
3. **The induced metric is connected (P2/P6):** The macro graph has finite distances between essentially all pairs (no or few infinite distances).
4. **Refinement is coherent (P4/P3):** Distances persist across scale changes up to rescaling. Using the refinement map $r : X_{j+1} \rightarrow X_j$ (fine to coarse), inter-scale distortion

$$\max_{a, b \in X_{j+1}} \left| d_\eta^{(j+1)}(a, b) - \alpha d_\eta^{(j)}(r(a), r(b)) \right| \tag{14}$$

is bounded (for some fitted $\alpha > 0$), and route mismatch between alternative coarse-graining routes is small.

When these conditions hold across a range of refinements, we treat $(X_j, d_\eta^{(j)})$ as an emergent geometric layer at scale j . Smooth Euclidean-like geometry corresponds to a regime in which local neighborhoods stabilize toward Euclidean behavior as j increases, while fractal regimes correspond to scale-stable closure without such smoothing.

4 Diagnostics: when geometry is coherent (and when it breaks)

SBT emphasizes falsification-first: a higher-layer theory is not declared real because it is appealing, but because it is *stable under repeated closure* and *auditable*. In this section we introduce the diagnostics used throughout the experiments. Each diagnostic answers a concrete question about coherence of the packaged layer under staging and refinement. We treat these quantities as falsification-first audits and gates, not as necessary and sufficient conditions for any particular continuum-limit notion of convergence.

4.1 Idempotence defect δ : does closure stabilize?

Fix staging τ and a lens $f : Z \rightarrow X$ with coarse matrix C and prototype matrix U . Consider the package-evolve-repackage operator acting on micro distributions,

$$E_{\tau,f}(\mu) := U_f(Q_f(\mu P^\tau)) = (\mu P^\tau C) U, \quad (15)$$

where $\mu \in \Delta(Z)$. If this operator is idempotent, $E_{\tau,f} \circ E_{\tau,f} = E_{\tau,f}$, then repeated closure does not drift. We quantify deviation from idempotence by the total-variation extreme-point defect

$$\delta_{\tau,f} := \frac{1}{2} \max_{z \in Z} \sum_{z' \in Z} \left| (E_{\tau,f}^2 - E_{\tau,f})[z, z'] \right|. \quad (16)$$

Equivalently, $\delta_{\tau,f}$ is the maximum total-variation gap between one-step and two-step closure applied to a point mass: $\delta_{\tau,f} = \max_{z \in Z} \text{TV}(\delta_z E_{\tau,f}, \delta_z E_{\tau,f}^2)$. A small $\delta_{\tau,f}$ indicates that closure is approximately stable under repetition; a large $\delta_{\tau,f}$ indicates that the “macro description” is not closed and will keep changing when reapplied.

4.2 Prototype stability $s_{\tau,f}$: do points persist under closure?

Each macro label $x \in X$ is represented by a prototype distribution u_x (row x of U). Stability asks whether the prototype returns to itself after closure:

$$s_{\tau,f}(x) := \text{TV}(E_{\tau,f}(u_x) - u_x). \quad (17)$$

We report mean and worst-case stability across x . Intuitively, stability measures “persistence of points”: if prototypes drift substantially, the intended points are not robust carriers at that stage.

4.3 Route mismatch (RM): does refinement commute?

Geometry is a multi-scale claim. Given three adjacent scales $j \rightarrow j+1 \rightarrow j+2$, there are two natural ways to go from the finest to the coarsest: a direct two-step packaging route or a two-stage composition. Route mismatch measures the disagreement between them:

$$\text{RM}(j) := d\left(\text{Pack}_{j \leftarrow j+2}, \text{Pack}_{j \leftarrow j+1} \circ \text{Pack}_{j+1 \leftarrow j+2}\right), \quad (18)$$

where d is a chosen operator distance. In our implementation we use a total-variation supremum proxy (extreme-point distance) and optionally a Frobenius norm. A small RM indicates that refinement and packaging routes approximately commute; a large RM indicates a failure of scale coherence.

4.4 Inter-scale distortion: does distance persist across refinement?

Even if closure is stable at each scale, “geometry” requires that the *induced distances* align across scales. Let $d_\eta^{(j)}$ be the induced macro distance at scale j , and let $r : X_{j+1} \rightarrow X_j$ be the refinement map (fine \rightarrow coarse). We measure inter-scale distortion by

$$\text{Dist}(j+1 \rightarrow j) := \max_{a,b \in X_{j+1}} \left| d_\eta^{(j+1)}(a,b) - \alpha d_\eta^{(j)}(r(a),r(b)) \right|. \tag{19}$$

The scalar $\alpha > 0$ may be fitted (e.g., by least squares through the origin) to account for the expected rescaling of distance between resolutions. Small distortion indicates that the notions of “near” and “far” are compatible across refinement; large distortion indicates that the geometry is not stable as the observer sharpens resolution.

4.5 Dimension diagnostics: “bits per scale” and ball growth

A smooth Euclidean-like regime is characterized not only by coherence, but also by *local growth laws*. SBT suggests a natural interpretation of dimension as an accounting exponent: how many distinguishable macro cells must be maintained as resolution increases.

Information (entropy) versus scale. Fix a micro distribution μ (uniform or stationary). At scale j , packaging yields $\nu_j = Q_{f_j}(\mu) \in \Delta(X_j)$ with entropy

$$H_j := - \sum_{x \in X_j} \nu_j(x) \log \nu_j(x). \tag{20}$$

Using a scale proxy ε_j (taken from typical nearest-neighbor distance in the induced macro metric), we estimate an information-dimension slope by regressing H_j against $\log(1/\varepsilon_j)$ Rényi [1959], Cover and Thomas [2006]. In smooth regimes this slope trends toward an integer-like value; in fractal regimes it tends toward non-integer behavior Falconer [2003].

Ball growth in the emergent metric. Given a macro distance $d_\eta^{(j)}$, define metric balls $B_x(r) = \{y : d_\eta^{(j)}(x,y) \leq r\}$ and the mean ball count $N(r) = \mathbb{E}_x |B_x(r)|$. We estimate a ball-growth dimension by fitting the slope of $\log N(r)$ versus $\log r$ in a non-saturated radius range. Smooth Euclidean-like regimes exhibit integer-like slopes, while fractal regimes exhibit non-integer-like scaling.

4.6 Connectivity: does the induced metric disconnect?

Because distance is computed from the induced macro dynamics, a failure mode is disconnection: if the macro graph lacks paths between regions, the induced shortest-path distance is infinite. We explicitly record the count of non-finite distances. Large numbers of infinities indicate that the stage, lens, or prototype choice has produced a broken geometric layer (often due to overly aggressive thresholding, inappropriate staging, or incoherent packaging).

4.7 Holonomy: curvature as protocol residue (P3 made geometric)

We detect curvature not by assuming a manifold but by measuring loop residue in local transport. We construct local neighborhoods using metric k -nearest neighbors, embed each neighborhood into \mathbb{R}^2 via local classical multidimensional scaling (MDS), and define transport between neighboring

neighborhoods via Procrustes alignment (a best-fit rotation). Given a small loop (triangle) (x, y, z) , the holonomy matrix is

$$H_{x \rightarrow y \rightarrow z \rightarrow x} := R_{x \rightarrow y} R_{y \rightarrow z} R_{z \rightarrow x}, \tag{21}$$

and we extract its rotation angle magnitude as the holonomy score. Plane-like substrates concentrate near zero holonomy; sphere-like substrates exhibit a shifted distribution. In SBT terms, this is P3: noncommutativity of protocols survives packaging as a stable geometric invariant.

4.8 Checklist: a practical “geometry birth” audit

For a lens ladder $\{f_j\}$ and staging τ , we treat a geometric layer as coherent over a scale range when the following conditions jointly hold:

- **Closure:** δ_{τ, f_j} is small across scales (package–evolve–repackage stabilizes).
- **Point stability:** $s_{\tau, f_j}(x)$ is small on average and in the worst case (prototypes persist).
- **Connectivity:** Induced shortest-path distances are finite (macro graph connected).
- **Refinement coherence:** Route mismatch is small and distortion (19) is bounded after rescaling.
- **Regime signature:** Dimension diagnostics separate smooth from fractal behavior; holonomy separates flat from curved behavior.

These diagnostics are deliberately conservative: they are designed to fail loudly. The robustness sweeps in later sections report where they break (e.g., staging too small or too large, overly fine or coarse ladders, prototype mismatch, threshold-induced disconnection, holonomy neighborhood instability, and aliasing effects in the Pythagoras experiment). In this sense, “where it breaks” is part of the theory: a layer is real only to the extent that it survives its own closure tests.

Lean anchors (minimal). We do not formalize the entire emergence pipeline in Lean. However, we include minimal anchors consistent with the above diagnostics: (i) shortest-path costs satisfy triangle inequality via path concatenation (pseudometric structure) and (ii) quotienting by zero-distance yields a metric space (standard separation quotient construction). These anchors support the narrative without turning the paper into a full formalization project.

5 Experimental pipeline and reproducibility

This paper is intentionally empirical in the SBT sense: we propose concrete closure constructions and then test whether they stabilize under repetition and refinement. To keep the narrative auditable, every figure and quoted metric in this paper is tied either to (i) a configuration file that regenerates the run or (ii) a committed “run pack” snapshot of the run outputs.

5.1 Substrates (microstate generators)

All experiments begin with a finite microstate space Z and a micro Markov kernel P .

Grid (plane-like substrate). We use a lazy random walk on an $n \times n$ grid with von Neumann neighborhood (up/down/left/right) and stay-put probability. This is a canonical substrate with isotropic local moves. It supports E1 (plane-like metric) and serves as the flat baseline for holonomy comparisons.

Sphere-like substrate (curved regime). To obtain a curved macro regime without assuming coordinates in the closure, we generate a point cloud on the unit sphere S^2 only as a means to build a micro transition graph: we sample points, build a k -nearest-neighbor weighted graph with Gaussian weights, and row-normalize to obtain a Markov kernel. Coordinates are not used by the closure pipeline; they are used only to generate the micro connectivity. This substrate supports E2 (holonomy and curvature separation).

Sierpiński gasket (fractal regime). We generate the Sierpiński gasket graph by a recursive corner-identification construction and define a lazy random walk on its adjacency graph. This yields a scale-stable but non-smooth neighborhood structure, supporting E3 (fractal regime).

Anisotropic gating (constraints as geometry deformation). To demonstrate the role of constraints (P2), we apply a directional gating transformation to a grid kernel: moves opposing a chosen direction are suppressed and the kernel is renormalized. This produces an anisotropically deformed macro geometry (E4).

5.2 Lens ladders (packaging families) and refinement maps

To test coherence across scales we require not a single lens but a refinement ladder f_0, \dots, f_L (coarse \rightarrow fine) together with canonical refinement maps. Operationally we construct lens ladders from *diffusion coordinates* of the micro kernel [Coifman and Lafon \[2006\]](#), [von Luxburg \[2007\]](#):

1. Compute a low-dimensional spectral embedding from a symmetrized kernel (diffusion/Laplacian coordinates).
2. Cluster the embedded points at the finest requested resolution (e.g., $m = 128$) using a deterministic k -means routine [Lloyd \[1982\]](#).
3. Build coarser levels by clustering cluster centroids, yielding a nested partition and an explicit refinement map $r : X_{\text{fine}} \rightarrow X_{\text{coarse}}$.

This produces (i) a label array $f_j : Z \rightarrow X_j$ at each scale and (ii) refinement maps that satisfy consistency $f_{\text{coarse}}(z) = r(f_{\text{fine}}(z))$ by construction.

Lens choice and (non-)circularity. Because our lens ladders are built from diffusion/spectral summaries of P , a natural concern is circularity: are we “injecting” geometry by using a geometry-learning method? In our setting, diffusion coordinates are computed from the micro Markov kernel itself and can be interpreted as extracting slow modes of the dynamics rather than importing an ambient coordinate system. We nevertheless emphasize the SBT stance: a lens family is an observer/interface choice, and different packaging strategies can yield different macro geometries. For that reason we do not claim an intrinsic geometry independent of packaging; instead we treat closure and refinement coherence as auditable properties of the induced layer and report robustness and failure modes alongside headline exhibits.

5.3 Prototypes

We use two prototype choices:

- **Uniform-on-block:** u_x is uniform over microstates assigned label x .

- **Stationary-conditional:** compute a stationary distribution π of P and set u_x proportional to π restricted to the block $f^{-1}(x)$.

Both satisfy the macro identity condition $Q_f \circ U_f = \text{id}_{\Delta(X)}$ by construction.

5.4 Macro dynamics, cost, and distance

Given a lens (C, U) at scale j and staging τ , we compute the induced macro kernel

$$\hat{P}_j = U_j P^\tau C_j, \tag{22}$$

then define costs from negative log likelihood and distances as all-pairs shortest-path cost on the resulting weighted macro graph. By default we symmetrize the kernel before converting to costs (to obtain an undirected metric), and we explicitly record disconnection as a failure mode via non-finite distances.

Computational note. At macro size m , distances are computed by all-pairs shortest paths on the macro move graph. In the canonical runs used for this paper we cap $m \leq 128$, so this step is inexpensive and stable. Holonomy adds repeated local neighborhood embeddings and alignments; we therefore bound the number of sampled loops (typically ≤ 800). Scaling the same diagnostics to $m \gg 10^3$ would require sparse/approximate shortest-path routines and more careful numerical conditioning, which we treat as future engineering work rather than a prerequisite for the present emergence exhibits.

5.5 Holonomy pipeline (curvature diagnostic)

To detect curvature as protocol residue, we compute a loop-based holonomy score from the induced macro metric:

1. Construct local neighborhoods using metric k -nearest neighbors.
2. Embed each neighborhood into \mathbb{R}^2 via local classical MDS [Borg and Groenen \[2005\]](#).
3. Define transport between overlapping neighborhoods by Procrustes alignment (best-fit rotation) [Schönemann \[1966\]](#).
4. For sampled small loops (triangles), compute the product of local transports around the loop and extract its rotation angle magnitude as the holonomy score.

We provide a canonical, deterministic holonomy configuration (including neighborhood parameters and loop sampling seeds) that regenerates the committed holonomy figure and summary JSON.

Two embeddings, two roles. We use diffusion/spectral embeddings only to construct the lens ladder from P (packaging); the emergent metric itself is defined from the induced macro kernel via costs and shortest paths. Separately, the holonomy diagnostic uses local MDS embeddings of *metric neighborhoods* as intermediate coordinate charts used only to estimate loop transport via alignment. Neither embedding is used to define the distance; both are auxiliary constructions whose outputs are audited for stability and sensitivity.

5.6 Reproducibility: configs, run folders, and committed run packs

All runs are generated from configuration files via the experiment harness:

```
python experiments/run.py --config experiments/configs/grid_plane.yaml
python experiments/run.py --config experiments/configs/sphere_knn.yaml
python experiments/run.py --config experiments/configs/sierpinski.yaml
python experiments/run.py --config experiments/configs/anisotropic.yaml
python experiments/run.py --config experiments/configs/holonomy_demo.yaml
python experiments/run.py --config experiments/configs/pythagoras_rw_grid.yaml
```

Each command produces a run folder `results/<run_id>/` containing `config.json`, `metrics.json`, `pointer.json`, a short log, and per-run artifacts and plots. These run folders are intentionally *not* committed.

For writing stability, the paper instead references committed run packs under `docs/notes/runs/<run_id>/` that contain the exact summary JSONs and plots used in the exhibits, as well as paper-ready comparison figures and a quotables table under `docs/notes/figures/` and `docs/notes/tables/`. This makes every quoted number auditable without requiring access to ephemeral local run directories.

6 Results I: coherent metrics, fractal regimes, and constraint deformation (E1–E4)

We now report the core empirical claim of this paper: when packaging and staging yield a coherent closure, the induced cost structure supports a stable notion of “near” and “far” that behaves like geometry. We present four exhibits—plane-like (grid), sphere-like (curved regime), Sierpiński (fractal regime), and anisotropic deformation under constraints—and interpret them through the diagnostics introduced in Section 4. (Curvature as holonomy is deferred to Section 7.)

Throughout, the figures and quoted summary numbers refer to committed run packs (Section 5), so the narrative is anchored to auditable artifacts rather than ad hoc screenshots. The overlay figures in this section provide a “regime map” across substrates: idempotence defect (closure stability), inter-scale distortion (refinement coherence), and the mean induced distance scale.

6.1 Overlay summaries across substrates

Table 2 reports representative per-run summary metrics for the four geo-pipeline exhibits (E1–E4), including δ and stability defects at the finest scale, connectivity (infinite-distance count), and maximum inter-scale distortion after rescaling. We also compute route mismatch (RM) for adjacent triples in the refinement ladder and record it in the committed run-pack summary JSONs; in the canonical runs it tracks the same coherence and failure modes as distortion (small when refinement is compatible, large when ladders become too fine or closure degrades).

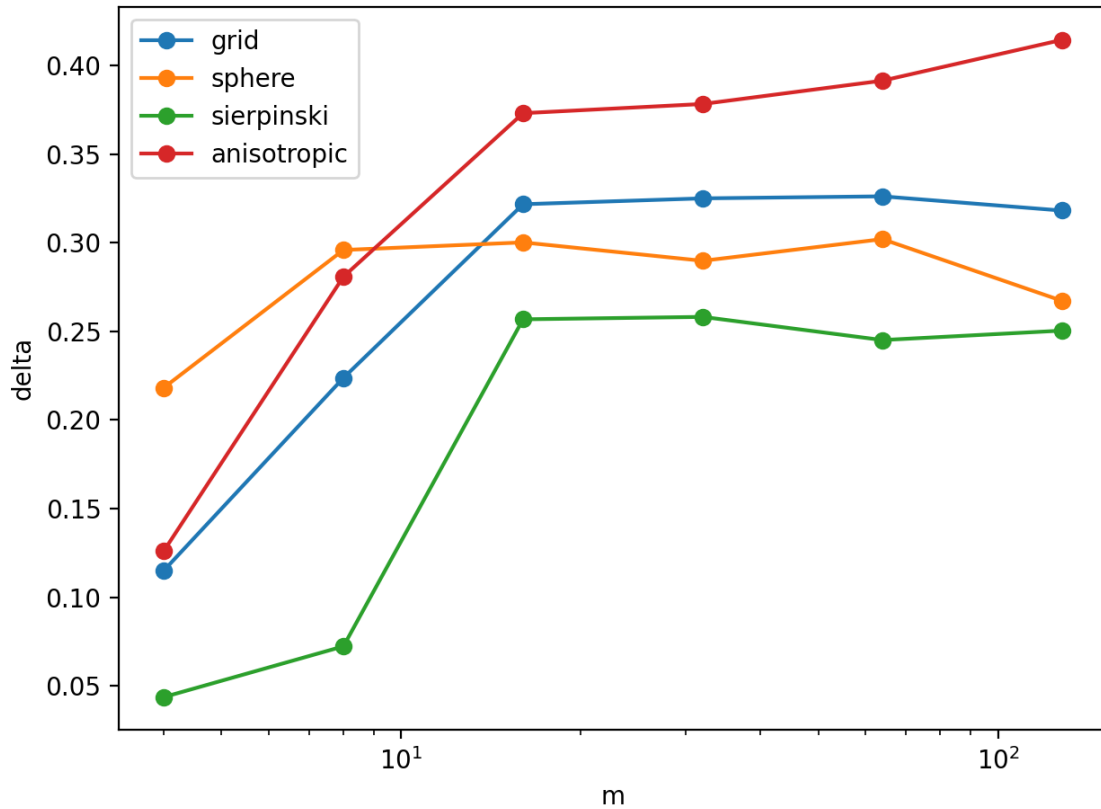


Figure 1: Idempotence defect δ across refinement, plotted against the macro state count m for four canonical substrates (grid, sphere-like, Sierpiński, anisotropic). Bounded δ indicates that repeated package–evolve–repackage does not drift arbitrarily; large or rapidly growing δ signals failure of closure.

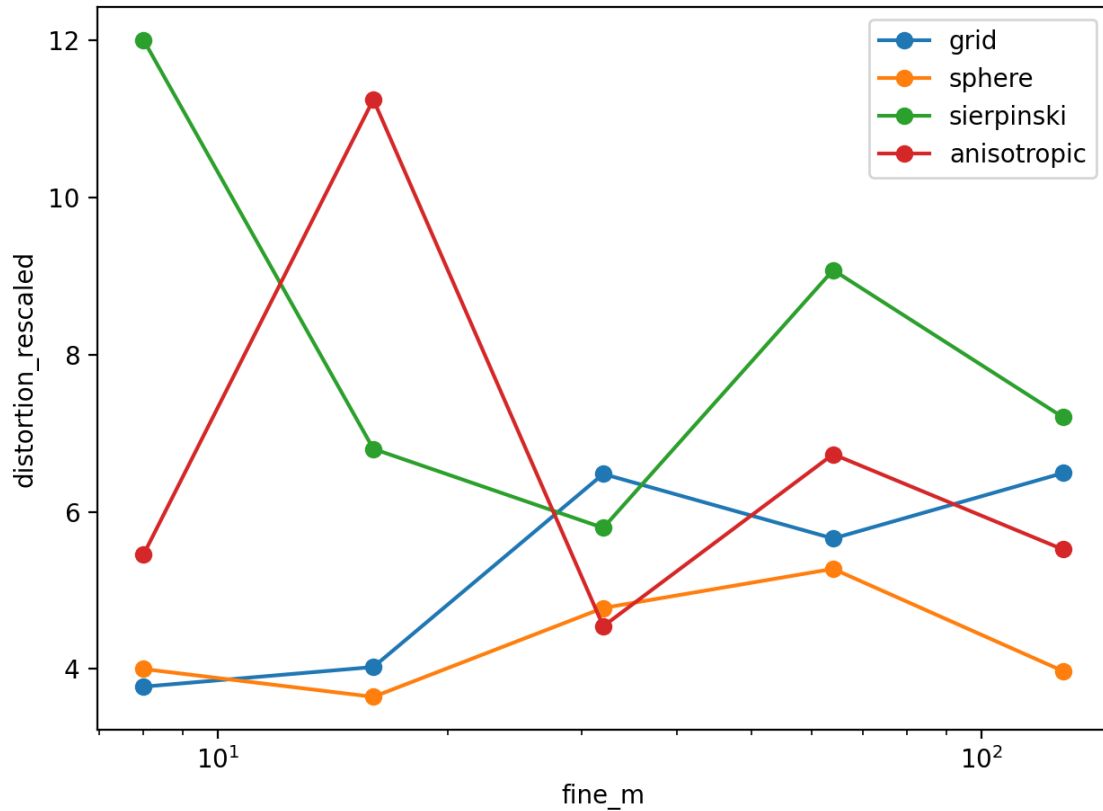


Figure 2: Inter-scale distortion across adjacent refinement levels (after fitted rescaling), plotted against the finer-level macro state count. Low distortion indicates that the induced notion of distance is compatible across refinement; high distortion indicates incoherence of “near” and “far” under scale change.

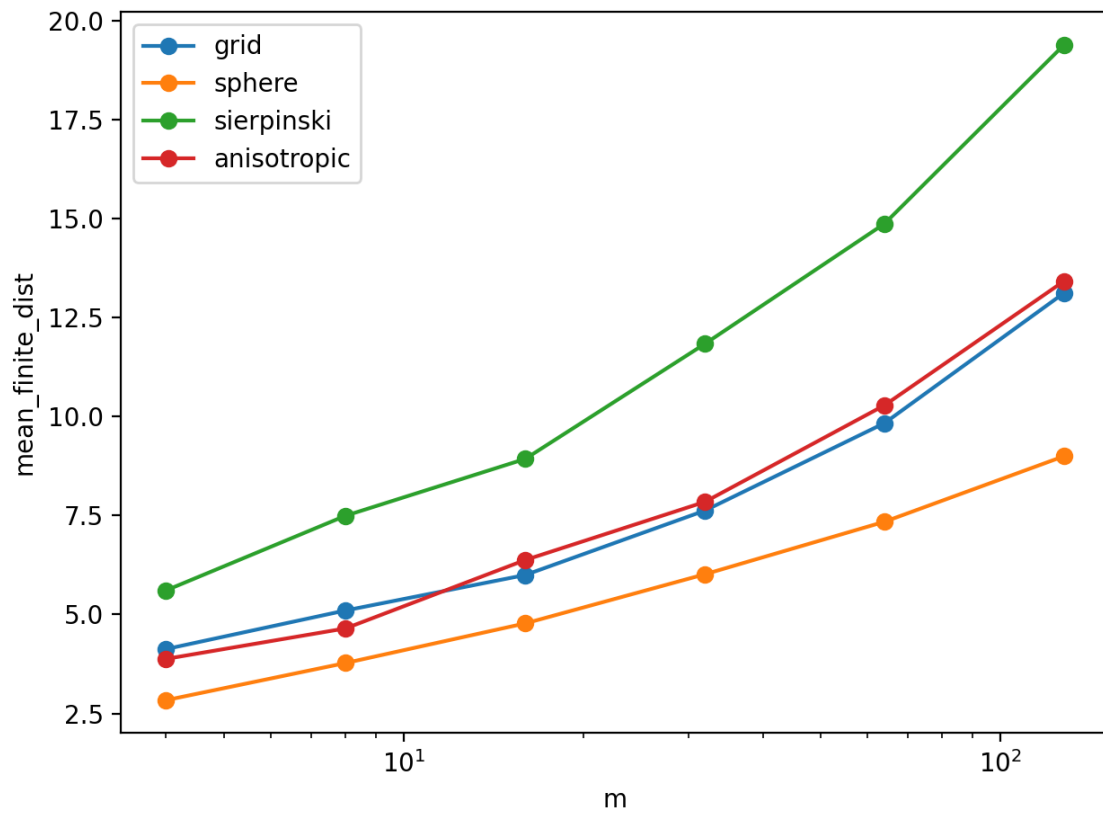


Figure 3: Mean finite induced distance as a function of macro resolution m . This summarizes the overall scale of the induced metric and helps interpret the rescaling factor in distortion measurements. Disconnection would appear as non-finite distances (recorded separately).

Table 2: Exhibit quotables (core exhibits E1–E4).

| exhibit | run_id | n_μ | m | τ | δ | stab_mean | stab_max | mean_dist | distort | n_∞ |
|-------------|-----------------------------------|---------|-----|--------|----------|-----------|----------|-----------|---------|------------|
| grid | grid_plane_20260202T193540Z_4b14 | 625 | 128 | 5 | 0.3182 | 0.6291 | 0.8687 | 13.12 | 6.493 | 0 |
| sphere | sphere_knn_20260202T193547Z_7681 | 500 | 128 | 5 | 0.2672 | 0.8552 | 0.9851 | 9.00 | 5.273 | 0 |
| sierpinski | sierpinski_20260202T193551Z_95ab | 366 | 128 | 5 | 0.2504 | 0.6348 | 0.8492 | 19.39 | 12.01 | 0 |
| anisotropic | anisotropic_20260202T193556Z_310f | 625 | 128 | 5 | 0.4145 | 0.6154 | 0.8558 | 13.42 | 11.25 | 0 |

Table 3: Exhibit quotables (supplements: holonomy).

| exhibit | run_id | plane_median | sphere_median | median_diff | plane_evaluated | sphere_evaluated |
|---------------|-------------------------------------|--------------|---------------|-------------|-----------------|------------------|
| holonomy_demo | holonomy_demo_20260202T215655Z_c66b | 0.0479 | 0.5980 | 0.5501 | 791 | 800 |

Table 4: Exhibit quotables (supplements: pythagoras).

| exhibit | run_id | τ_{\min} | τ_{\max} | fit_rms _{min} | fit_rms _{max} | pyth_med _{min} | pyth_med _{max} | axis_q/l |
|--------------------|------------------------------------------|---------------|---------------|------------------------|------------------------|-------------------------|-------------------------|----------|
| pythagoras_rw_grid | pythagoras_rw_grid_20260202T193600Z_ed66 | 4 | 128 | 12.10 | 0.1509 | 33.19 | 0.0586 | 0.0168 |

6.2 E1: Plane-like emergent metric on a grid

The grid substrate is the canonical “flat” case: micro moves are local and isotropic, and the lens ladder packages microstates into increasingly refined macro regions. Across the refinement ladder, the induced closure remains bounded under repetition (Figure 1), and the induced distances remain compatible across refinement up to rescaling (Figure 2). The macro metric graph remains connected in the canonical run (finite distances throughout), confirming that the induced notion of distance is not an artifact of disconnection or thresholding.

In SBT terms, this is the regime in which packaging (P5) and staging (P4) yield a usable closed operator (P1) and accounting (P6) produces a stable ledger over protocols (P3). The outcome is not “coordinates” but a coherent notion of proximity: macro states that exchange probability mass easily are near; those that do not are far.

6.3 E2: Sphere-like macro metric (curvature expected; holonomy deferred)

The sphere-like substrate is constructed to be locally similar to the grid in terms of local connectivity and random-walk dynamics, but globally incompatible with a flat embedding. At the level of closure and distance coherence, the sphere-like run is comparable to the grid case: the induced macro kernel yields a connected macro graph, and the core defects remain bounded over the ladder (Figures 1–2). On these diagnostics alone, one might call the geometry “well-formed.”

However, coherence of distances does not by itself detect curvature. Curvature is a statement about *transport around loops*: whether the order of local moves matters after packaging. This is exactly P3 (protocol composition) made geometric. In Section 7 we therefore measure loop holonomy and show a clear separation between plane-like and sphere-like substrates: small-loop holonomy concentrates near zero for the grid but shifts upward for the sphere-like case.

6.4 E3: Sierpiński gasket (fractal regime)

The Sierpiński gasket substrate exhibits stable multi-scale structure without smooth local Euclidean neighborhoods. In this regime, the emergence question is not “does it become a manifold” but “does it become coherent under refinement at all, and if so, what invariants stabilize?” The overlay diagnostics show that closure and induced distance can remain meaningful (bounded defects, finite distances) while refinement does not smooth toward a single Euclidean tangent picture. This is the SBT distinction between a smooth fixed point and a scale-stable (fractal) fixed point: both are valid higher-layer theories, but they stabilize different invariants.

In practice, scaling diagnostics such as entropy-versus-scale and ball-growth in the induced metric are useful but can be sensitive in finite graphs and depend on the chosen cost-to-distance convention. We therefore treat “dimension” estimates as secondary corroboration rather than a headline claim: our main point in E3 is that refinement can yield a coherent layer without smoothing toward Euclidean local neighborhoods. Recovering precise Hausdorff-like dimensions is out of scope here; we instead emphasize auditable coherence and document scaling behavior and its sensitivities in the accompanying artifacts and robustness notes.

6.5 E4: Constraints deform geometry (anisotropic gating)

To demonstrate that geometry is not merely “read off” from a substrate, we introduce constraints (P2) that alter feasibility of moves. Starting from the grid kernel, we apply directional gating that suppresses motion against a preferred direction and renormalizes the kernel. This produces a macro

geometry that is still coherent enough to support an induced metric, but is systematically deformed: neighborhoods and distances reflect the biased feasibility structure rather than an isotropic baseline.

Conceptually, this is a direct SBT statement: constraints are not secondary—they define what protocols exist and what accounting costs can be minimized. When P2 changes, the emergent geometry changes. The anisotropic exhibit thus serves as a controlled deformation of the plane-like regime, illustrating that the “geometry layer” is an induced theory of feasible transformations and their costs, not a fixed container.

Summary of E1–E4. Across these exhibits we see the same emergence calculus expressed in distinct refinement-stable regimes (analogous to fixed points in renormalization): (i) a plane-like regime where closure and distances are coherent and curvature is absent, (ii) a curved regime where distances remain coherent but loop transport accumulates residue, (iii) a fractal regime where refinement stabilizes scale-laws rather than smooth tangent structure, and (iv) constraint-driven deformation where feasibility reshapes the induced metric. Section 7 makes the curvature statement explicit by measuring holonomy and then presents the Pythagoras emergence experiment as a signature accounting-to-geometry bridge.

7 Results II: curvature as holonomy and the Pythagorean form (E2, E5)

Section 6 established that the induced cost metric can be coherent across refinement on multiple substrates. We now present two signature results that connect SBT directly to classical geometric structure: (i) curvature detected as loop holonomy (protocol residue) and (ii) the emergence of a Pythagorean form from staged isotropic diffusion under a cost-as-accounting definition.

7.1 E2: Curvature as protocol residue (holonomy)

In SBT terms, curvature is not a primitive tensor on a pre-existing manifold. It is the macroscopic footprint of P3: *protocol order matters*. After packaging, local “move” operators need not commute; when we compose local transports around a small loop, we may fail to return aligned. That loop residue is called holonomy [Lee \[1997\]](#).

Operationally (Section 4), we compute holonomy from the induced macro metric by (i) selecting local metric neighborhoods, (ii) embedding them into \mathbb{R}^2 via local MDS, (iii) defining transport between overlapping neighborhoods via Procrustes alignment (a best-fit rotation), and (iv) composing transports around sampled triangles. The resulting holonomy score is the rotation angle magnitude (in radians) of the loop transport product.

Figure 4 shows a clear separation between plane-like and sphere-like substrates under a single canonical configuration that is reproducible from `experiments/configs/holonomy_demo.yaml`. In the committed summary JSON (`docs/notes/holonomy_demo_summary.json`), the median holonomy is 0.0479 on the plane-like case versus 0.5980 on the sphere-like case (median difference 0.5501, median ratio 12.49), with 791 and 800 evaluated loops, respectively. Importantly, the holonomy pipeline is deterministic under fixed seeds: rerunning the same config reproduces the same summary statistics exactly, and sensitivity to neighborhood parameters is treated as a first-class failure mode in the robustness analysis. We interpret this as diagnostic evidence of curvature-like loop residue in the packaged layer; we do not infer a smooth curvature tensor or sectional curvature from this estimator.

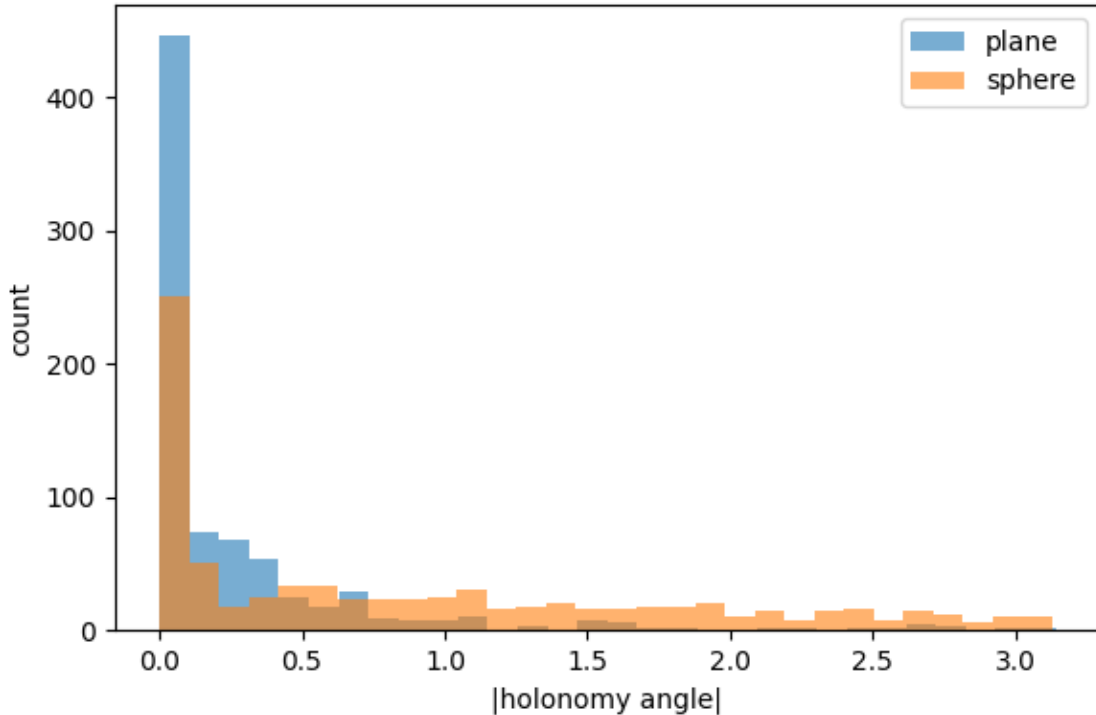


Figure 4: Holonomy (loop rotation angle magnitude) for a plane-like grid substrate versus a sphere-like substrate under the canonical holonomy configuration. The plane distribution concentrates near zero, while the sphere distribution is shifted upward, indicating curvature as stable protocol residue after packaging.

In geometric language, this is the promised bridge from SBT to curvature: P5 packaging produces macro “points,” P6 accounting produces a macro metric, and P3 holonomy appears as loop-dependent residue of local transports. The sphere-like substrate exhibits nonzero holonomy despite having otherwise coherent distance structure, demonstrating that curvature is not captured merely by distance coherence but by transport noncommutativity.

7.2 E5: Pythagoras as emergent accounting identity

The Pythagorean theorem is often presented as an immutable truth about Euclidean triangles. From the SBT perspective, it becomes a statement about a *stable accounting law*: when staged dynamics and isotropy force costs to become quadratic and separable across orthogonal directions, the Pythagorean form emerges as the closed relation between costs of composed moves.

Setup: cost from staged isotropic diffusion. Consider a lazy isotropic random walk on a large 2D torus (to avoid boundary effects). Let $p_\tau(\Delta x, \Delta y)$ denote the displacement probability after τ micro-steps. We define the accounting cost surface

$$C_\tau(\Delta x, \Delta y) := -\log(p_\tau(\Delta x, \Delta y)), \quad (23)$$

with standard numerical smoothing in implementation [Cover and Thomas \[2006\]](#). As τ increases, the central limit regime makes the displacement distribution approximately Gaussian [Durrett](#)

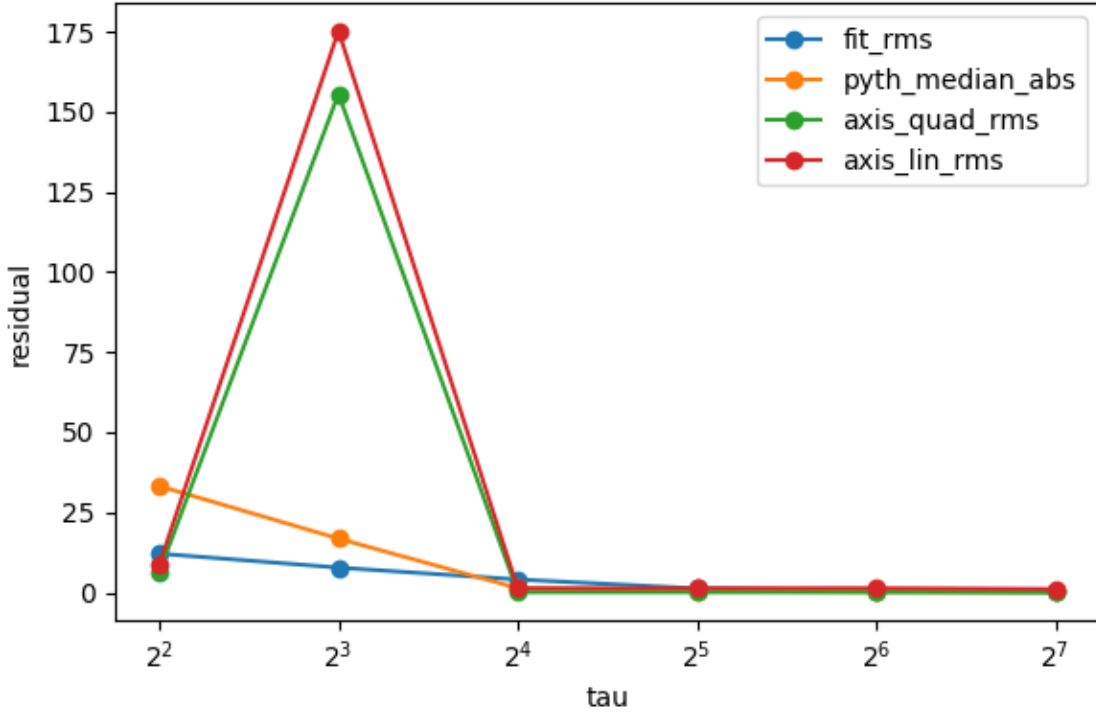


Figure 5: Pythagorean residual statistics versus stage τ for the isotropic torus random walk cost surface. Residuals drop sharply after a regime transition (around $\tau \approx 16$ in the canonical run), consistent with the onset of an approximately quadratic and separable accounting law.

[2019]; in the isotropic case, the negative log probability becomes approximately quadratic in displacement. Empirically we test the emergent quadratic form

$$C_\tau(\Delta x, \Delta y) \approx a_\tau(\Delta x^2 + \Delta y^2) + b_\tau, \quad (24)$$

and we quantify how closely it holds as τ varies.

Remark (mechanism exhibit). This experiment is intentionally stylized: rather than passing through a lens ladder, we measure displacement likelihood directly on a substrate where the diffusion mechanism is clean. The purpose is to isolate a concrete route by which an accounting-based cost becomes approximately quadratic and separable under staging and isotropy, not to assert that every emergent metric produced by our closure pipeline must be Euclidean.

Pythagorean residual as a protocol-composition test. To express “right-triangle additivity” in cost terms, we compare the diagonal cost to the composed axis costs. Define the Pythagorean residual

$$R_\tau(\Delta x, \Delta y) := C_\tau(\Delta x, \Delta y) - \left(C_\tau(\Delta x, 0) + C_\tau(0, \Delta y) - C_\tau(0, 0) \right). \quad (25)$$

If cost is separable and quadratic as in (24), then $R_\tau(\Delta x, \Delta y) \approx 0$ over a broad range of displacements, and level sets of C_τ become approximately circular.

In the canonical run (committed summary `docs/notes/pythagoras_rw_grid_summary.json`), the quadratic fit RMS decreases from 12.10 at $\tau = 4$ to 0.1509 at $\tau = 128$, and the median absolute




Figure 6: Cost contours $C_\tau(\Delta x, \Delta y)$ at small τ (left) and large τ (right). As staging increases, the cost surface “circularizes” toward an isotropic quadratic form, consistent with emergent Euclidean distance structure under accounting-as-distance.

Pythagorean residual decreases from 33.19 to 0.0586. Beyond the transition, axis scaling strongly favors quadratic growth: at $\tau = 128$, the quadratic-axis RMS is 0.0177 versus 1.0536 for a linear-axis model. These trends are precisely what the SBT framework predicts when P4 staging pushes dynamics into a stable diffusion regime and P2 isotropy enforces symmetry between directions.

Control: Manhattan cost fails as expected. To ensure the above is not a tautology of shortest paths, we include a control in which the cost is explicitly Manhattan:

$$C_{L1}(\Delta x, \Delta y) := |\Delta x| + |\Delta y|. \quad (26)$$

This control produces diamond-shaped contours rather than circles and exhibits linear axis scaling (near-zero axis-linear RMS by construction), while the quadratic-axis model performs poorly. Note that this is a geometric negative control rather than a matched dynamical control: it enforces linear additivity by definition rather than by altering the underlying random-walk dynamics. Matched dynamical controls (e.g., anisotropic or driven/nonreversible walks) are a natural next step.

Bird-level interpretation. This experiment makes the “distance is accounting” principle concrete. P6 provides the ledger: cost is defined as negative log likelihood under the induced (staged) dynamics. P4 provides the stage parameter τ that pushes the system into a stable diffusion regime. P2 isotropy ensures symmetry and (approximately) independent contributions of orthogonal directions. Finally P3 is the protocol logic: the “hypotenuse” move is compared to the composed axis moves via (25). In this regime, the stable accounting law takes the Pythagorean form. Importantly, this does not claim Euclidean geometry is fundamental; it demonstrates that Euclidean structure is a *stable higher-layer description* in an isotropic staged-diffusion regime under SBT closure.




Figure 7: Control: Manhattan (L1) cost yields diamond contours, not circular ones. This explicitly fails the quadratic accounting law and serves as a negative control for the emergence of Pythagorean structure.

8 Robustness, failure modes, and limitations

SBT is falsification-first: a layer is not declared real because it is elegant but because it *survives its own closure tests*. In the geometry setting, this means we must treat “space birth” as conditional. The same pipeline that yields coherent metrics and curvature signatures in Sections 6–7 also provides knobs that can cause the layer to destabilize, disconnect, or become ambiguous. This section summarizes the main sensitivity findings and corresponding failure modes.

All robustness sweeps are documented in an internal note (`docs/notes/robustness_20260202.md`) and are reproducible via sweep configuration files under `experiments/configs/sweeps/`. Here we summarize the outcomes at the level appropriate for the paper: what breaks, why it breaks in SBT terms, and what this implies for claims of emergent geometry.

8.1 Representative failure modes (“where it breaks”)

Grid (E1): staging that is too large can inflate stability defects. On the plane-like grid substrate, increasing the stage parameter τ does not monotonically improve closure. While moderate staging pushes dynamics into a regime where the induced cost metric is coherent across refinement, overly large staging can inflate prototype stability defects: macro representatives drift under repeated closure. In SBT terms, P4 staging is not “more is better”: too much micro evolution

between repackaging can wash out the very distinctions the lens is attempting to preserve, so P5 packaging no longer yields persistent macro carriers.

Grid (E1): very fine ladders can amplify inter-scale distortion. Increasing the number of refinement levels or pushing to very fine macro resolutions can increase inter-scale distortion even when per-scale closure defects remain bounded. This is a coherence failure rather than a local failure: distances may exist at each scale, but the refinement ladder is not compatible enough for a single stable geometry to persist across the range. In SBT terms, P4 staging and P5 packaging have produced usable layers, but P3 cross-scale commutation (route coherence) becomes strained as the ladder becomes too fine for the fixed lens construction.

Connectivity failure: thresholding and smoothing can disconnect the metric. A direct failure mode of the induced metric is disconnection: if the macro move graph becomes disconnected (often due to overly aggressive edge thresholding or inappropriate likelihood smoothing choices), shortest-path distances become infinite. This is not a minor numerical artifact; it is a conceptual failure of the claimed geometry layer. In SBT terms, accounting (P6) has been applied to a move system that no longer supports global protocols (P3), so “distance” ceases to be a meaningful global invariant.

Sphere holonomy (E2): neighborhood choice can destabilize curvature estimation. Holonomy is intentionally a higher-order diagnostic: it depends on local neighborhoods, local embeddings, and overlap-based transport. Robustness sweeps show that holonomy magnitudes and even the number of *evaluated* loops can become unstable when neighborhood size is too small, overlap thresholds are too strict, or neighborhood expansion is disabled. In such regimes, the pipeline can undersample valid loops or produce noisy local embeddings that inflate angles. This sensitivity is not a flaw to be hidden; it is an expected feature of a curvature diagnostic built from finite, packaged neighborhoods. It also motivates our use of a canonical deterministic configuration (`holonomy_demo.yaml`) for the headline separation figure.

Pythagoras (E5): finite-size aliasing can stall the quadratic regime. The Pythagoras experiment relies on a staged diffusion regime in which displacement statistics approach a smooth, approximately Gaussian form. If the torus is too small relative to τ (or if the sampled displacement window is too wide for the available support), wrap-around aliasing can disrupt the emergence of a clean quadratic cost surface and stall the improvement of the Pythagorean residual. In SBT terms, this is a failure of staging-as-limit: P4 is being pushed beyond the regime where the substrate supports the intended diffusion approximation, so the accounting law (P6) does not stabilize into the Euclidean form.

8.2 Knobs that matter (practical guidance)

The following parameters materially affect whether a coherent geometric layer emerges and which regime it occupies:

- **Staging τ (P4):** Too small yields under-mixed, noisy macro costs; too large can inflate prototype drift or wash out distinctions.
- **Refinement ladder / macro resolutions (P4/P5):** Overly coarse ladders may hide structure; overly fine ladders can amplify distortion or route mismatch.

- **Prototype choice (P5/P1):** Uniform-on-block versus stationary-conditional prototypes can change stability and idempotence behavior.
- **Cost smoothing and thresholds (P6):** Parameters such as η (log smoothing) and edge thresholding can affect connectivity and metric stability.
- **Holonomy neighborhood parameters (P3):** k_{neigh} , overlap thresholds, and neighborhood expansion control the stability and evaluability of loop transport.
- **Finite-size parameters in Pythagoras (P4):** Torus size N , displacement window, and τ range determine whether the diffusion regime is clean or aliased.

8.3 Limitations and non-claims

We emphasize what the results do *not* claim.

- **Not a proof of Euclidean fundamentality.** The Pythagoras result does not assert that Euclidean geometry is fundamental; it provides evidence that Euclidean structure is a stable higher-layer accounting law in an isotropic staged diffusion regime under our cost definition.
- **Finite constructions.** All substrates here are finite and all limits are operational: refinement ladders, staging ladders, and stability are tested by diagnostics rather than by asymptotic theorems.
- **Lens dependence.** Different lens families (different packaging strategies) can produce different macro points and therefore different induced geometries. This is expected in SBT: a geometry is layer-relative.
- **Curvature estimation is diagnostic, not axiomatic.** Holonomy is measured via local embedding and alignment; it is a robust separator in the canonical configuration but remains sensitive to neighborhood choices, as the sweeps document.

Taken together, these limitations strengthen the SBT interpretation: a geometric layer is a *conditional closure artifact*. When it stabilizes across repetition and refinement, “space” becomes available as a reliable compression. When it does not, the diagnostics indicate which primitive (staging, packaging, constraints, accounting, or protocol coherence) has failed to support the intended layer.

9 Discussion and conclusion: what SBT predicts about space

This paper answers the title question in SBT terms: *to plot a stone* is not to recover coordinates from a pre-existing container, but to construct a stable notion of location, distance, and transport from finite-interface closure. When packaging, staging, and accounting cohere under refinement, a space-like layer becomes available. When they do not, the diagnostics indicate precisely which aspect of closure fails.

9.1 What we showed

We implemented and audited a concrete emergence pipeline from micro-dynamics to a macro metric, then tested it across multiple substrates. The central empirical claims are:

- **Points are packaged states.** A “point” at a given resolution is an equivalence class induced by a lens $f : Z \rightarrow X$ (P5). It is not assumed; it is constructed.
- **Distance is optimized accounting.** From the induced macro kernel $\hat{P} = UP^rC$ we define costs from negative log likelihood and distances as minimal protocol cost (shortest paths). This realizes P6 (ledger) optimized over P3 (composition of moves).
- **Geometry is coherence under refinement.** A single distance matrix is not yet geometry; geometry is a claim that the packaged layer stabilizes across staging and a refinement ladder. We operationalize this via closure idempotence and prototype stability defects, route mismatch, inter-scale distortion (after rescaling), and connectivity.
- **Curvature appears as holonomy.** Distance coherence alone does not detect curvature. Curvature appears as a stable loop residue of local transport (P3), and we empirically separate plane-like and sphere-like substrates by a holonomy distribution shift.
- **The Pythagorean form emerges in an isotropic diffusion regime.** Under staged isotropic diffusion, negative log transition probability becomes approximately quadratic and separable, and the Pythagorean residual collapses, while a Manhattan control fails as expected.

Taken together, these results support the SBT slogan stated in the introduction: *geometry is what you get when repeated packaging produces a stable notion of nearby, composition of moves, and cost of moving, and that structure closes under refinement*. Importantly, the paper is not a collection of isolated plots: it is a demonstration that SBT’s closure primitives can be made computationally concrete and audited in a way that distinguishes regimes (flat versus curved versus fractal versus constrained deformation).

9.2 What we did not claim

The results are deliberately scoped. In particular:

- **We did not claim that geometry is fundamental.** The Euclidean and Pythagorean structures observed here are not asserted to be axioms of reality; they are shown to be stable higher-layer accounting laws under specific staging, isotropy, and packaging choices.
- **We did not claim uniqueness.** Different lens families (different packaging strategies) can induce different macro points and therefore different geometries. This is expected in SBT: geometry is layer-relative.
- **We did not prove asymptotic limits.** All constructions are finite and operational. “Limit behavior” is inferred from coherence across ladders, not from a formal $\varepsilon \rightarrow 0$ theorem.
- **We did not reduce curvature to a single scalar.** Holonomy is a diagnostic built from local embeddings and alignments; it is a robust separator under canonical parameters, but it remains sensitive to neighborhood choices, as documented in the robustness sweeps.

These non-claims are not disclaimers; they are part of the SBT stance. A space-like layer is a conditional birth, and its claims should be phrased in terms of stability under the very closure operations that define it.

9.3 Predictions and next experiments

The six-birds framing makes concrete predictions about when different geometric regimes should appear.

Predictions.

- **Constraints shape geometry (P2).** Changing feasibility (gating, anisotropy, locality) should predictably deform the induced metric and transport invariants. In particular, strong directional constraints should produce anisotropic geodesics and may induce effective “cones” of reachability in the macro space.
- **Curvature is protocol noncommutativity (P3).** Whenever local transports are forced to be patchwise (because the layer is packaged), noncommutativity should generically appear, and holonomy-like residues should persist under refinement when the layer is coherent.
- **Smooth versus fractal is a refinement-regime question (P4/P5/P6).** Away from critical regimes, refinement should tend to smooth local neighborhoods toward Euclidean-like behavior. Near critical or self-similar regimes, refinement should stabilize scale laws (non-integer dimensions, anomalous diffusion) rather than smooth tangent structure. We use “fixed point” language only as an analogy: we do not implement a formal renormalization operator or prove convergence.
- **Pythagoras is an accounting identity, not an axiom (P6/P4/P2).** Quadratic-additive structure should appear precisely in regimes where staged dynamics yields approximately Gaussian displacement statistics with isotropy; it should fail under anisotropy, strong constraints, or finite-size aliasing.

Next experiments (concrete). Several natural extensions would strengthen the emergence calculus and broaden applicability:

- **Operator emergence beyond distance (P1).** Study when induced macro operators become local and stable (discrete Laplacians, gradient-like operators, connections) and whether curvature can be recovered in multiple equivalent ways (e.g., via commutators of induced operators).
- **Learned lenses and prototypes.** Instead of fixing a lens ladder by diffusion clustering, learn packaging and prototype choices by explicitly minimizing closure defects and distortion across scales, turning “geometry discovery” into an optimization problem.
- **Alternative ledgers (P6).** Replace negative log likelihood with other accounting notions (control energy, communication cost, repair cost) and test which ledgers yield which emergent geometries on the same substrate.
- **Beyond synthetic substrates.** Apply the pipeline to real-world substrates that are naturally graph-like (transport networks, biological connectomes, interaction graphs) to test whether the same coherence diagnostics predict when a usable geometry layer appears.
- **Formal convergence targets.** Connect the empirical coherence criteria to formal notions of convergence of metric spaces (e.g., Gromov–Hausdorff style) and to renormalization-style fixed points for the lens ladder [Gromov \[1999\]](#).

9.4 Six birds, one end-to-end story

It is worth restating the complete pipeline in bird terms, because it is the core conceptual contribution.

- **P5 Packaging** creates the candidate “points” as equivalence classes under indistinguishability: $f : Z \rightarrow X$.
- **P4 Staging** provides the ladder (both in τ and in refinement levels) on which stability can be tested and on which smooth versus fractal regimes separate.
- **P1 Operator rewrite** is realized by closure: inserting prototypes to rewrite micro dynamics as a macro kernel $\hat{P} = UP^TC$ and auditing its idempotence and stability.
- **P6 Accounting** turns macro feasibility into a ledger (cost), from which distance is defined.
- **P3 Protocols** enter twice: distance is the minimal cost over composed move protocols (shortest paths), and curvature appears as the residue of composing local transports around loops (holonomy).
- **P2 Constraints** shape the feasible protocol space, thereby deforming the geometry and potentially changing which fixed point the refinement ladder approaches.

Closing. Plotting is therefore not a primitive act; it is a successful closure. When closure succeeds, we obtain geometry as a stable higher-layer description of what moves are possible and what they cost. When closure fails, the diagnostics tell us why. In this sense, *space is not where the stone is*; space is what becomes true about the stone when the six birds make a stable map possible.

A Lean anchors (minimal)

We include a small set of Lean 4 “anchors” that support the narrative without attempting to formalize the entire emergence pipeline. The intent is modest: to confirm that the two load-bearing mathematical moves used throughout the paper align with standard metric constructions.

Build. The Lean project lives under `lean/`. To build:

```
cd lean && lake build
```

Anchors proved.

- `GeoSBT/PathMetric.lean` proves `graph_edist_triangle`: triangle inequality for a shortest-path distance induced by nonnegative edge costs (path concatenation).
- `GeoSBT/QuotientMetric.lean` proves `separation_quotient_metric` and `separation_quotient_dist_eq_zero`: the standard construction turning a pseudo-metric into a metric by quotienting by the “zero distance” relation (separation quotient).
- `GeoSBT/Pythagoras.lean` proves `pythagoras_real`: a classical Euclidean sanity anchor (orthogonality implies squared-norm additivity in \mathbb{R}^2).

These theorems do not prove that our emergent metric is Euclidean. Rather, they formalize the minimal mathematical scaffolding that underlies our constructions: “distance as minimal path cost” is a pseudo-metric, and identifying points at zero distance yields a metric space.

B Reproducibility appendix

This repository is designed so that the paper can be audited from stable, committed artifacts *and* regenerated end to end from configuration files.

B.1 Configs (regenerating runs)

End-to-end runs are produced via the harness:

```
python experiments/run.py --config experiments/configs/grid_plane.yaml
python experiments/run.py --config experiments/configs/sphere_knn.yaml
python experiments/run.py --config experiments/configs/sierpinski.yaml
python experiments/run.py --config experiments/configs/anisotropic.yaml
python experiments/run.py --config experiments/configs/holonomy_demo.yaml
python experiments/run.py --config experiments/configs/pythagoras_rw_grid.yaml
```

Each command creates `results/<run_id>/` with `config.json`, `metrics.json`, `pointer.json`, per-run artifacts, and plots. The `results/` directory is intentionally not committed.

B.2 Config format and determinism

The configuration files under `experiments/configs/` use the `.yaml` extension but contain JSON content (valid YAML 1.2) so they can be parsed without external YAML dependencies. Each config includes explicit seeds; in particular, lens clustering uses seeded, deterministic routines and the holonomy demo uses deterministic loop sampling under fixed seeds (with sorted iteration to avoid ordering nondeterminism). As with any linear-algebra-heavy pipeline, minor floating-point differences across platforms/BLAS implementations are possible; the committed run packs should be treated as the reference artifacts for the numbers and figures quoted in this paper.

B.3 Committed run packs (stable references for writing)

For writing stability, we commit run packs under:

```
docs/notes/runs/<run_id>/
```

These packs contain the exact plots and summary JSONs referenced by the paper, along with paper-ready comparison figures (`docs/notes/figures/`) and a quotables table (`docs/notes/tables/`). The paper itself uses copies of these figures under `paper/figures/`.

B.4 Canonical configuration snapshot (major knobs)

For reader convenience we summarize the major knobs of the canonical configurations used in the paper. The full parameter lists are in the config files themselves; these tables are a snapshot of the most load-bearing settings.

Table 5: Canonical geo-pipeline configurations used for E1–E4. Common settings: $\tau = 5$, lens levels [4, 8, 16, 32, 64, 128] with $n_{\text{eigs}} = 6$, prototypes=`uniform`, symmetrize=`weight_avg`, $\eta = 10^{-12}$, and $\varepsilon_{\text{edge}} = 10^{-15}$.

| Exhibit | Config | Substrate (micro kernel) |
|----------------|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E1 plane | <code>grid_plane.yaml</code> | Grid random walk: <code>n_side=25</code> , <code>lazy=0.5</code> . |
| E2 sphere | <code>sphere_knn.yaml</code> | Sphere point cloud kNN: <code>n_points=500</code> , <code>k=10</code> , $\sigma = 0.5$, <code>self_loop=10⁻⁶</code> . (Coordinates are used only to generate the micro graph; closure never accesses them.) |
| E3 fractal | <code>sierpinski.yaml</code> | Sierpiński gasket random walk: <code>level=5</code> , <code>lazy=0.5</code> . |
| E4 constrained | <code>anisotropic.yaml</code> | Grid + anisotropic gate: direction east, strength 1.0 (renormalized). |

Table 6: Canonical holonomy configuration used for Figure 4 (E2).

| Item | Setting |
|------------------------|-----------------------------------------------------------------------------------------------|
| Config | <code>holonomy_demo.yaml</code> |
| Macro construction | $\tau = 5$, lens levels [128], $n_{\text{eigs}} = 6$ (seed 0). |
| Holonomy neighborhoods | $k_{\text{neigh}} = 24$, <code>expand_hops=1</code> , <code>min_overlap=4</code> . |
| Loop sampling | triangles via $k_{\text{loop}} = 8$, <code>max_loops=800</code> , seeds (plane=0, sphere=1). |

Table 7: Canonical Pythagoras run configuration (E5).

| Item | Setting |
|-----------|----------------------------------------------------------------------------------------------------------|
| Config | <code>pythagoras_rw_grid.yaml</code> |
| Substrate | 2D torus random walk with $N = 512$, <code>lazy=0.5</code> . |
| Stages | $\tau \in \{4, 8, 16, 32, 64, 128\}$. |
| Sampling | displacement window via <code>D_factor=3</code> and <code>D_max=30</code> ; right-triangle samples=2000. |
| Numerics | p -floor 10^{-300} and fit threshold $p_{\text{min}} = 10^{-20}$. |

B.5 Export and comparison scripts

Two helper scripts support writing-time reproducibility:

- `python scripts/export_run_artifacts_to_docs.py -recommended -overwrite` exports recommended run outputs from `results/` into committed run packs under `docs/notes/runs/`.
- `python scripts/make_paper_ready_comparisons.py` regenerates the overlay comparison figures and quotables tables from the committed run packs.

B.6 Paper build

To build the PDF:

```
cd paper && make pdf
```

References

- Ingwer Borg and Patrick J. F. Groenen. *Modern Multidimensional Scaling: Theory and Applications*. Springer Series in Statistics. Springer, 2 edition, 2005. doi: 10.1007/0-387-28981-X.
- Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. doi: 10.1016/j.acha.2006.04.006.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2 edition, 2006. ISBN 978-0-471-24195-9.
- Rick Durrett. *Probability: Theory and Examples*. Cambridge University Press, 5 edition, 2019.
- Kenneth Falconer. *Fractal Geometry: Mathematical Foundations and Applications*. John Wiley & Sons, 2 edition, 2003.
- Mikhail Gromov. *Metric Structures for Riemannian and Non-Riemannian Spaces*. Birkhäuser, 1999. Translated by Sean Michael Bates.
- John M. Lee. *Riemannian Manifolds: An Introduction to Curvature*, volume 176 of *Graduate Texts in Mathematics*. Springer, 1997.
- Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982. doi: 10.1109/TIT.1982.1056489.
- J. R. Norris. *Markov Chains*. Cambridge University Press, 1997. ISBN 978-0-521-48181-6.
- Alfréd Rényi. On the dimension and entropy of probability distributions. *Acta Mathematica Academiae Scientiarum Hungaricae*, 10(1-2):193–215, 1959. doi: 10.1007/BF02063299.
- Peter H. Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, 1966. doi: 10.1007/BF02289451.
- Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948. doi: 10.1002/j.1538-7305.1948.tb01338.x.
- Ioannis Tsiokos. Six birds: Foundations of emergence calculus, January 2026a. URL <https://doi.org/10.5281/zenodo.18365949>. Zenodo.
- Ioannis Tsiokos. To count a stone with six birds: A mathematics is a theory, January 2026b. URL <https://doi.org/10.5281/zenodo.18402004>. Zenodo.
- Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. doi: 10.1007/s11222-007-9033-z.