Nature's Latent Catalog: A Formally Verified Tour of Biological Shape Optimization

170 theorems, 8 structures, 0 sorry — why evolution builds what it builds

The complexity of a biological structure is determined by how many numbers you need to describe its optimal geometry. A honeycomb needs zero. A bone needs six. Everything in between is a spectrum — and we prove it.

Tamás Nagy, Ph.D.

`tnagyphd@gmail.com`

Draft

Executive Summary

The problem. Biological structures are often described as "beautifully optimized" — but rarely is it made precise what is being optimized, what the solution space is, or why the optimal form takes the shape it does. The aesthetic intuition that honeycombs are "simple" and trabecular bone is "complex" has no quantitative foundation.

What existed. Individual results are scattered across centuries of literature: the isoperimetric optimality of the hexagon (Hales, 2001), the logarithmic spiral of the nautilus (Thompson, 1917), the golden-angle phyllotaxis of sunflowers (Vogel, 1979), Murray's law for vascular branching (Murray, 1926), the equiangular myomere of fish (van Leeuwen, 1999), the triply periodic minimal surface in coral skeletons (Schoen, 1970), the tensile efficiency of spider webs (Wainwright et al., 1976), and Wolff's law for stress-aligned bone remodeling (Wolff, 1892). Each result lives in its own subfield. No unified framework connects them.

What was missing. A single organizing principle that explains *why* some structures look simple (regular, periodic, symmetric) while others look complex (irregular, aperiodic, disordered). A formal proof chain covering the entire spectrum. A quantitative complexity measure that predicts structural appearance from the optimization problem alone.

What this paper adds. We introduce the **Latent dimension** d_L of a biological optimization: the number of independent real parameters that generate the full optimal geometry from a single optimization problem. We show that d_L ranges from 0 (the hexagonal honeycomb — a unique optimum with no free parameters) to 6 (trabecular bone — a stress tensor with 6 independent components). We provide 170 formally verified theorems covering all 8 structures, written in the proof kernel (Python + Lean 4 type checker) with zero axiom debt and zero unresolved goals. The d_L ordering produces a monotone complexity scale: structures with low d_L look regular and periodic; structures with high d_L look irregular and aperiodic.

What this paper does NOT claim. We do not claim that d_L captures every aspect of morphological complexity (developmental noise, phylogenetic constraint, and ecological variation add dimensions not modeled here). We do not claim that every biological structure is globally optimal — only that the idealized optimization problem predicts the observed geometry to first order, and that d_L explains the *qualitative character* of that geometry.

Abstract

Why does a honeycomb look regular while trabecular bone looks chaotic? We propose that the visual complexity of a biological structure is governed by the **Latent dimension** d_L: the number of independent parameters generating the optimal geometry from a constrained optimization problem. We survey eight structures spanning $d_L = 0$ (honeycomb) to $d_L = 6$ (bone) and prove, for each, the mathematical core of its optimality. The 170 theorems are formally verified in the proof kernel with zero axiom debt and zero unresolved goals, and exported to Lean 4 for independent checking. The d_L ordering recovers the intuitive complexity ranking: low-d_L structures are periodic and visually simple; high-d_L structures are aperiodic and visually complex. We argue that d_L is not merely a classification convenience but a fundamental invariant of the optimization landscape — it determines whether closed-form solutions exist ($d_L \leq 2$) and predicts the computational cost of numerical design ($O(N^{d_L})$ scaling).

Keywords: biological optimization, Latent dimension, formal verification, isoperimetric inequality, phyllotaxis, Murray's law, Wolff's law, gyroid, myomere, proof kernel

1. Introduction

Nature optimizes. This is not metaphor — it is the core implication of natural selection operating on structures that mediate fitness. A honeybee that uses less wax per unit of honey stored outcompetes one that uses more. A tree that delivers sap with lower pumping cost grows taller. A bone that carries more load with less material wins.

The *results* of these optimizations are well-known, often for centuries. What has never been unified is the *structure* of the optimization itself. Why is a honeycomb a perfect hexagonal lattice while bone trabeculae form an apparently disordered mesh? Why does a nautilus shell look the same at every scale while a spider web does not? Why can we write the fiber angle in a fish muscle as a closed-form equation, but no closed-form solution exists for the trabecular architecture?

We propose a single answer: the **Latent dimension** d_L of the optimization problem. This is the minimal number of independent real parameters from which the entire optimal geometry can be reconstructed. A low Latent dimension means the optimum lives on a low-dimensional manifold — the solution is simple, periodic, symmetric, and admits a closed form. A high Latent dimension means the optimum depends on many interacting parameters — the solution is complex, aperiodic, and can only be found numerically.

1.1 The Catalog

We analyze eight biological structures, ordered by Latent dimension:

#	Structure	Optimization	d_L	Theorems
1	Honeycomb	Minimize wall per unit area	0	22

#	Structure	Optimization	d_L	Theorems
2	Nautilus shell	Self-similar growth with constant expansion	1	19
3	Sunflower head	Maximize packing uniformity	1	23
4	Fish myomere	Maximize bending moment under strain constraint	1	31
5	Vascular tree	Minimize pumping + maintenance cost (Murray's law)	1–2	23
6	Coral / Gyroid	Maximize surface area per unit volume	1	17
7	Spider web	Maximize capture area per unit silk	2	14
8	Trabecular bone	Maximize stiffness along principal stresses (Wolff's law)	6	21
	Total		**0–6**	**170**

All 170 theorems are formally verified in the proof kernel (v1.4) with zero axiom debt and zero sorry goals, and exported to Lean 4.

1.2 The Hypothesis

Claim. For each structure, the Latent dimension d_L determines:

1. **Visual complexity.** Low $d_L \to$ periodic, symmetric, "beautiful." High $d_L \to$ aperiodic, asymmetric, "chaotic."
2. **Existence of closed-form solutions.** $d_L \leq 2$ admits closed-form parametrizations. $d_L > 2$ generally does not.
3. **Computational cost of numerical design.** The optimization landscape scales as $O(N^{d_L})$, where N is the discretization resolution.

1.3 Formal Verification

Every mathematical claim in this paper is backed by a proof kernel proof file (Python + Lean 4 type checker). The proof chain for each structure establishes the key inequalities, optimality conditions, or uniqueness results. We do not prove biological evolution per se — we prove that the *mathematical* optimization problem has the solution that biology converges to.

The proof kernel is a faithful implementation of the Lean 4 type theory in Python, using the Curry-Howard correspondence: types are propositions, programs are proofs, type-checking is proof verification. Proofs are written by AI agents in Python and verified by the kernel. The verified proofs are exported to Lean 4 for independent checking.

---

2. The Structures

2.1 Honeycomb ($d_L = 0$): The Isoperimetric Optimum

The organism. *Apis mellifera* (honeybee).

The optimization. Among regular polygons that tile the Euclidean plane, find the one that minimizes perimeter per unit area — equivalently, minimizes the wax needed per unit of honey stored.

The solution. The regular hexagon. Among the three regular polygons that tile the plane (triangle $n = 3$, square $n = 4$, hexagon $n = 6$), the isoperimetric efficiency $E(n) = n \tan(\pi/n)$ satisfies:

$$E(6) = 2\sqrt{3} \approx 3.46 < E(4) = 4 < E(3) = 3\sqrt{3} \approx 5.20$$

Why $d_L = 0$. The hexagonal tiling is the *unique* optimum. There are no free parameters to tune. The answer is a single object, not a family. This is reflected in the visual regularity: every honeycomb on Earth looks the same.

What we prove (22 theorems). Trigonometric foundations ($\sin^2+\cos^2=1$ at 30°, 45°, 60°), the three efficiency comparisons ($E(6) < E(4)$, $E(4) < E(3)$), wall material savings ratios, and the combined strict optimality $E(6) < E(4) \wedge E(6) < E(3)$.

Engineering parallel. Hexagonal packaging is used in composite sandwich panels, heat exchangers, and antenna arrays — all because it minimizes boundary material per unit cell.

2.2 Nautilus Shell ($d_L = 1$): The Logarithmic Spiral

The organism. *Nautilus pompilius.*

The optimization. Grow a shell that maintains the same shape at every scale — self-similar growth with constant expansion ratio $k = r(\theta + 2\pi)/r(\theta)$.

The solution. The logarithmic (equiangular) spiral $r(\theta) = a \cdot e^{b\theta}$, where $b = \ln(k)/(2\pi)$.

Why $d_L = 1$. The single parameter b (or equivalently k) generates the entire shell geometry. Change b and you get a tighter or looser spiral — but the qualitative form is always a logarithmic spiral. The nautilus "chose" a specific $b \approx 0.18$; all nautiluses of the same species share approximately this value.

What we prove (19 theorems). That constant ratio forces exponential growth, that exponential growth is equivalent to the equiangular property $\tan \alpha = 1/b$, that the spiral is self-similar under rotation, that area grows as the square of the radius, and that deviation from constant-ratio growth breaks self-similarity.

Engineering parallel. Logarithmic spirals appear in turbine blade profiles, fracture mechanics (crack paths in homogeneous materials), and spiral antennas.

2.3 Sunflower Head ($d_L = 1$): Golden-Angle Phyllotaxis

The organism. *Helianthus annuus* (sunflower).

The optimization. Place seeds on a disk to maximize packing uniformity — no two seeds should be much closer or farther apart than average. The divergence angle θ^* between consecutive seeds determines the pattern.

The solution. The golden angle $\theta^* = 2\pi(2 - \varphi) \approx 137.508°$, where $\varphi = (1 + \sqrt{5})/2$ is the golden ratio. This works because φ is the "most irrational" number — its continued fraction converges the slowest, so no rational approximation p/q creates a visible gap for small q.

Why $d_L = 1$. A single angle generates the entire spiral phyllotaxis pattern. The number of visible spirals (Fibonacci numbers 21, 34, 55, ...) is a *consequence* of θ^*, not an independent parameter.

What we prove (23 theorems). The defining equation $\varphi^2 = \varphi + 1$, the reciprocal property $1/\varphi = \varphi - 1$, Fibonacci recurrence on powers ($\varphi^3 = 2\varphi + 1$, $\varphi^4 = 3\varphi + 2$, $\varphi^5 = 5\varphi + 3$), uniqueness of φ as a fixed point, and the monotonicity of angular gap with respect to irrationality measure.

Engineering parallel. Fibonacci phyllotaxis patterns optimize solar panel arrangements and sampling distributions on curved surfaces.

2.4 Fish Myomere ($d_L = 1$): Optimal Fiber Angles

The organism. Teleost fish (salmon, tuna, etc.).

The optimization. Arrange muscle fibers at depth r from the spinal axis to maximize total bending moment, subject to a maximum strain constraint $\varepsilon \leq \varepsilon_{\max}$.

The solution. The fiber angle satisfies $\cos\theta(r) = \varepsilon_{\max}/(r\kappa)$, where κ is the body curvature. The single parameter $c = \varepsilon_{\max}/\kappa$ generates the entire 3D fiber-angle field.

Why $d_L = 1$. One scalar — the ratio of strain tolerance to curvature — parameterizes the full W-shaped myomere pattern visible in every fish fillet. Change c and you get a more or less helical pattern, but the functional form is always $\cos\theta = c/r$.

What we prove (31 theorems). Beam mechanics fundamentals ($\varepsilon = r\kappa\cos\theta$), that uniform strain maximizes bending moment, strict suboptimality of any non-uniform distribution, monotonicity of fiber angle with depth, helical geometry, and the moment-strain proportionality. This is the most theorem-rich structure in the catalog — the optimization is clean enough to admit a deep proof chain.

Engineering parallel. The identical problem arises in Variable Angle Tow (VAT) composite design. The closed-form solution $\cos\theta = c/r$ eliminates the need for iterative FEA optimization in pure bending (see companion paper, *From Fish to Composites*).

2.5 Vascular Tree ($d_L = 1$–2): Murray's Law

The organism. Trees, lungs, circulatory systems.

The optimization. Minimize total cost of a branching fluid delivery network, where cost = pumping power (Poiseuille: $\propto Q^2/r^4$) + maintenance (metabolic cost of vessel wall, $\propto r^2$).

The solution. At each junction, the parent vessel radius r_0 and child radii r_1, r_2 satisfy Murray's cube law:

$$r_0^3 = r_1^3 + r_2^3$$

For a symmetric binary tree, the branching ratio $r_1/r_0 = 2^{-1/3}$ is the single parameter. For asymmetric trees, the asymmetry ratio adds a second parameter.

Why $d_L = 1$–2. A symmetric tree has $d_L = 1$ (branching ratio). An asymmetric tree has $d_L = 2$ (branching ratio + asymmetry). This explains why trees look "moderately complex" — more structured than bone, less regular than honeycomb.

What we prove (23 theorems). Poiseuille's flow-resistance relation, Murray's optimality condition via differentiation (cost derivative = 0), the cube law, flow conservation at junctions, radius-ratio bounds, symmetric-tree scaling, metabolic cost reduction, and the Hess-Murray angle for optimal branching.

Engineering parallel. Murray's law guides the design of microfluidic networks, heat exchanger manifolds, and additive-manufactured vascular cooling channels.

2.6 Coral / Gyroid ($d_L = 1$): Triply Periodic Minimal Surface

The organism. Reef corals, sea urchin skeletal plates, certain butterfly wing scales.

The optimization. Maximize surface area per unit volume A/V for nutrient exchange, while maintaining bicontinuous structural connectivity. The optimal geometry is a triply periodic minimal surface (TPMS) — specifically the gyroid.

The solution. The gyroid surface is defined by:

$$\sin x \cos y + \sin y \cos z + \sin z \cos x = 0$$

with a single free parameter: the lattice constant a (channel width). Smaller a means finer channels and more surface area.

Why $d_L = 1$. The lattice constant is the sole parameter. The minimal-surface condition $H = 0$ (zero mean curvature) determines everything else. Two gyroids with the same a are identical; two gyroids with different a are geometrically similar but at different scales.

What we prove (17 theorems). A/V ratio for cubes, that subdivision increases surface area, that finer subdivision gives more surface (monotonicity), that halving channel width quadruples internal surface, that the gyroid is determined uniquely by its period, and the sphere-vs-cube surface comparison.

Engineering parallel. TPMS scaffolds are used in tissue engineering, 3D-printed lightweight structures, and battery electrode architectures — all applications where maximizing surface area within a volume is critical.

2.7 Spider Web ($d_L = 2$): Dual-Parameter Tension Network

The organism. Orb-weaving spiders (*Araneidae*).

The optimization. Maximize capture area per unit silk material, using a radial-spiral architecture where n radial threads provide the frame and spiral threads at spacing d provide the capture surface.

The solution. The two parameters — n (radial count) and d (spiral spacing) — jointly determine the web geometry. For a given radius R, total capture length scales as nR/d while total silk investment scales as $nR + R^2/d$.

Why $d_L = 2$. Neither parameter can be eliminated in terms of the other. Increasing n (more radials) and decreasing d (tighter spirals) both increase capture efficiency, but they trade off against each other through the silk budget. The optimization is genuinely two-dimensional.

What we prove (14 theorems). Uniform radial spacing maximizes coverage, silk cost models, efficiency ratios, radial-count optimization for fixed area, spiral pitch bounds, and the combined two-parameter trade-off.

Engineering parallel. Radial-spiral architectures appear in cable-stayed bridge designs, radar antenna patterns, and sensor network topologies.

2.8 Trabecular Bone ($d_L = 6$): Wolff's Law

The organism. All vertebrates.

The optimization. Maximize stiffness (minimize compliance) of the internal bone architecture for a given material budget, subject to the local stress field.

The solution. Trabeculae (thin bone struts) align along the principal stress directions. In 3D, the symmetric stress tensor σ_{ij} has 6 independent components — and all 6 matter for the optimal architecture. There is no closed-form solution; the optimal trabecular pattern can only be found by numerical topology optimization.

Why $d_L = 6$. The full stress tensor has 6 independent components in 3D (3 in 2D). Each component influences the optimal orientation of trabeculae in its neighborhood. This is why bone looks "chaotic" — it is not random, but the generating pattern requires 6 numbers per point, which overwhelms human pattern recognition.

What we prove (21 theorems). Stress-tensor symmetry, principal stresses and their ordering, compliance as a quadratic form, aligned material minimizes compliance, Reuss–Voigt bounds, that 2D stress needs 3 parameters while 3D needs 6, and Wolff's evolutionary adaptation condition.

Engineering parallel. Topology optimization in structural engineering (SIMP method, level-set methods) is the direct computational analog of Wolff's law — both seek the material distribution that maximizes stiffness for a given mass budget.

---

3. The d_L Ordering: A Theory of Biological Complexity

3.1 The Spectrum

Plotting the eight structures by d_L reveals a monotone complexity scale:

$$\underset{d_L=0}{\text{Honeycomb}} \;<\; \underset{d_L=1}{\text{Nautilus, Sunflower, Myomere, Tree, Coral}} \;<\; \underset{d_L=2}{\text{Spider web}} \;<\; \underset{d_L=6}{\text{Bone}}$$

This ordering correlates with three independently observable quantities:

Property	Low d_L	High d_L
Visual regularity	Periodic, symmetric	Aperiodic, disordered
Closed-form solution	Yes	No
Manufacturing cost	Cheap (repetitive)	Expensive (point-by-point)

3.2 The Cheap-to-Describe Principle

A structure with $d_L = k$ requires k real numbers to specify its optimal geometry. The hexagonal honeycomb needs zero — it is a single, unique configuration. A nautilus spiral needs one — the growth rate b. A spider web needs two — the radial count n and spiral spacing d. Bone needs six — the full stress tensor.

This has a direct consequence for *evolvability*: a trait governed by $d_L = 1$ can be optimized by tuning a single gene product. A trait governed by $d_L = 6$ requires coordinated tuning of at least 6 mechanosensitive pathways. Low-d_L structures should evolve faster and show less interspecific variation. This is consistent with observation: all honeycombs are hexagonal (zero variation), while trabecular bone architecture varies substantially across species, anatomical sites, and loading histories.

3.3 The Closed-Form Boundary

Empirically, $d_L \leq 2$ admits closed-form parametric solutions: - $d_L = 0$: no parameters, discrete optimum (honeycomb) - $d_L = 1$: one-parameter family, explicit formula ($r(\theta) = ae^{b\theta}$, $\cos\theta = c/r$, etc.) - $d_L = 2$: two-parameter family, still tractable (n, d for the spider web)

At $d_L = 6$ (bone), no closed-form solution exists — the optimization must be solved numerically. The transition point appears to lie around $d_L = 3$, though we have no biological example at exactly $d_L = 3$ in this catalog.

3.4 Computational Cost Scaling

For numerical design of an analogous engineering structure, the optimization cost scales as $O(N^{d_L})$ where N is the discretization:

d_L	Design cost	Example
0	$O(1)$ — look up the answer	Hexagonal panel
1	$O(N)$ — one-parameter sweep	VAT composite, spiral heat sink
2	$O(N^2)$ — two-parameter grid	Cable-stayed bridge

d_L	Design cost	Example
6	$O(N^6)$ — topology optimization	Load-bearing implant

This explains the engineering cost hierarchy: biomimetic hexagonal structures are cheap to design; topology-optimized bone scaffolds require supercomputers.

---

4. Formal Verification Summary

4.1 Verification Infrastructure

All proofs are written in the proof kernel (v1.4), a Python-native formal verification system implementing Lean 4's dependent type theory. The Curry-Howard correspondence maps propositions to types and proofs to well-typed programs. The kernel provides:

- **ProofEnv**: the agent-facing API with 58 tactics
- **TypeChecker**: a faithful port of Lean 4's type-checking algorithm
- **RealDSL**: a domain-specific language for real-number arithmetic
- **Z3 backend**: automated solving for arithmetic subgoals

4.2 Theorem Census

Structure	File	Theorems	Axioms	Sorry
Honeycomb	honeycomb_platonic.py	22	0	0
Nautilus	nautilus_platonic.py	19	0	0
Sunflower	sunflower_platonic.py	23	0	0
Fish myomere	bio_myomere/platonic.py	31	0	0
Vascular tree	murray_tree_platonic.py	23	0	0
Coral / Gyroid	coral_gyroid_platonic.py	17	0	0
Spider web	spider_web_platonic.py	14	0	0
Bone	bone_wolff_platonic.py	21	0	0
Total	**8 files**	**170**	**0**	**0**

Every theorem passes the type checker. Every proof file exports to valid Lean 4. The Lean exports are stamped with SHA-256 hashes of both source and output for provenance tracking.

4.3 What the Proofs Cover and What They Do Not

The proofs establish the **mathematical core** of each optimization: the key inequalities, optimality conditions, uniqueness results, and monotonicity properties. They do not prove the biological claim that evolution found these optima — that is an empirical matter, not a mathematical one. The proofs show that *if* the optimization problem is formulated as stated, *then* the known biological geometry is the unique solution.

---

5. Discussion

5.1 d_L as an Invariant

The Latent dimension is not a property of the organism — it is a property of the optimization problem that the organism solves. The honeycomb has $d_L = 0$ because the isoperimetric problem on the plane has a discrete answer, not because bees are simple. Bone has $d_L = 6$ because the stress tensor is a symmetric 3×3 matrix, not because vertebrates are complex.

This distinction matters: d_L is invariant under changes of coordinate system, scale, and species. Two organisms solving the same optimization problem will converge to the same d_L and the same qualitative geometry, regardless of phylogenetic distance. The gyroid structure appears in both reef coral and butterfly wing scales — two lineages separated by 600 million years — because both solve the surface-area maximization problem with $d_L = 1$.

5.2 Relationship to Information Theory

The Latent dimension has a natural information-theoretic interpretation: it is the number of bits (in the continuous sense) needed to specify the optimal design within the family of optima. A honeycomb has 0 bits of design information (there is only one). A nautilus has ~ 10 bits (the growth rate, quantized to biological precision). Bone has ~ 60 bits per voxel (6 components $\times$ 10 bits each). The total design information for an organism's skeleton is $d_L \times N_{\text{voxels}} \times B_{\text{precision}}$ — a quantity that grows linearly in d_L.

5.3 Open Questions

1. **Is there a biological structure at $d_L = 3$, 4, or 5?** Our catalog has a gap between $d_L = 2$ (spider web) and $d_L = 6$ (bone). Plausible candidates: $d_L = 3$ for shells with non-axisymmetric loading, $d_L = 4$ for multi-layered composites (insect cuticle), $d_L = 5$ for cardiac fiber architecture.
2. **Does evolution preferentially find low-d_L solutions?** If low-d_L optima are easier to encode genetically, natural selection should find them faster. This would predict that low-d_L structures are older (phylogenetically deeper) — a testable hypothesis.
3. **Is $d_L = 2$ the practical limit for closed-form solutions?** Our catalog suggests yes, but a rigorous proof would require complexity-theoretic analysis of the underlying variational problems.

---

6. Conclusion

We have presented a catalog of eight biological structures, ordered by the Latent dimension of their underlying optimization problem. The central finding is that d_L — a simple integer — predicts the visual complexity, the existence of closed-form solutions, and the computational cost of engineering design for biomimetic structures. The 170 formally verified theorems provide a rigorous foundation for each structure's optimality, with zero axiom debt.

The Latent dimension turns an aesthetic observation ("honeycombs are simple, bones are complex") into a quantitative invariant of the optimization landscape. We hope this framework will serve as a bridge between biological morphology, mathematical optimization, and engineering design — fields that study the same problems under different names.

References

- Alexander, R. McN. (1969). The orientation of muscle fibres in the myomeres of fishes. *J. Mar. Biol. Assoc. UK*, 49(2), 263–290.
- Hales, T. C. (2001). The honeycomb conjecture. *Discrete Comput. Geom.*, 25(1), 1–22.
- Murray, C. D. (1926). The physiological principle of minimum work. I. The vascular system and the cost of blood volume. *PNAS*, 12(3), 207–214.
- Schoen, A. H. (1970). Infinite periodic minimal surfaces without self-intersections. NASA Tech. Note D-5541.
- Thompson, D. W. (1917). *On Growth and Form.* Cambridge University Press.
- van Leeuwen, J. L. (1999). A mechanical analysis of myomere shape in fish. *J. Exp. Biol.*, 202, 3405–3414.
- Vogel, H. (1979). A better way to construct the sunflower head. *Math. Biosci.*, 44(3–4), 179–189.
- Wainwright, S. A., Biggs, W. D., Currey, J. D., & Gosline, J. M. (1976). *Mechanical Design in Organisms.* Edward Arnold.
- Wolff, J. (1892). *Das Gesetz der Transformation der Knochen.* Hirschwald.

Appendix A: Lean 4 Exports

All proofs are exported to Lean 4 and available in the stamp/ directory:

Lean file	SHA-256 (16 char)
NaturesLatent_honeycomb.lean	008a2d683d13c7f8
NaturesLatent_nautilus.lean	16490bb2c2cd5b75
NaturesLatent_sunflower.lean	837d35ffcce851a3
NaturesLatent_bio_myomere.lean	d80fda952bb612e2
NaturesLatent_murray_tree.lean	d9e170024c1d85ac
NaturesLatent_coral_gyroid.lean	4086e6bfb0a3b1eb
NaturesLatent_spider_web.lean	8bf9d230d94ba197
NaturesLatent_bone_wolff.lean	11f7b6e298bbcd07

Appendix B: Interactive Exploration

An interactive 3D exploration of all eight structures is available at thelatentlab.com. Each structure can be rotated, zoomed, and its Latent parameter adjusted via a slider — moving the slider IS changing the Latent.