The Birth of Morphotopy:
From Earth-Measure to the Spaces of Life and Mind

PIATRA . INSTITUTE

June 2026

Abstract

Across physics, biology, cognitive science, and machine learning the same move recurs: a domain that is not spatial is rendered as a structured space of possible forms, and its behaviour is treated as constrained movement within that space. Configuration spaces, morphospaces, conceptual spaces, and the latent spaces of trained networks are instances. This paper argues that the recurrence is the latest phase of one long process in the history of geometry, the progressive detachment of form from physical extension, and that reading it as a set of independent metaphors misses that continuity. We trace that history stratigraphically, from the surveyor's field and the Euclidean figure through coordinates, the collapse of the parallel postulate, Riemann's intrinsic manifold, Klein's transformation groups, topology, phase space, and information geometry, to biological morphospace and the representational geometry of brains and machines, each stratum remaining in use beneath the next. We then name the current phase *morphotopy* and, so that the name does more than relabel, give five conditions a treatment must satisfy to count as a morphotopic one rather than a spatial figure of speech, and apply them to cases that pass and cases that fail. One structural obstruction is definable in any of these spaces, geometric frustration, the failure of locally satisfiable constraints to compose into a global solution; we exhibit it exactly in a minimal spin model, the one stratum where it can be computed cleanly, and argue rather than demonstrate its recurrence elsewhere. We date the strata to show that the abstraction accelerates, three quarters of them falling in the last twentieth of the span, an acceleration that is in part a feature of how finely the recent strata are individuated. The claim is historical and organizational rather than a new mathematics, and it is bounded: rendering the structure of a mind as navigable geometry is not an account of why there is experience.

1. A Name That Outlived Its Object

Geometry means earth-measure, from *gē* and *metron*, and the etymology records a practice the field abandoned long ago, the survey of land for boundary, tax, and the recovery of plots after a flood. The name now sits over an object that has changed past recognition. A working geometer studies manifolds that no surveyor could walk, probability families, latent spaces inside neural networks, and the space of possible animal forms, and none of these is land. The persistence of the old name is evidence rather than an oversight. A field keeps the term that coined it the way a body keeps a vestigial bone, and the bone records the function that is gone.

The temptation is to ask what all these spaces have in common, and the temptation should be refused, because the answer at that level is empty: they are all called spaces, and the word explains nothing it is applied to. The productive question is genealogical. Geometry did not always range over morphospaces and embeddings; it came to, through a sequence of definite changes, each of which moved its object one step further from the measured field, until the object was a space with no field, then a space with no figure, and at last a space with no points. To recover that sequence is to see why the disparate modern uses are continuous with land survey rather than loose analogies to it, and to locate the change now underway.

Two commitments organize the reading. The first treats each stratum as a regime of what could be made visible rather than as a step toward a single correct geometry. The question for each is not whether it was true, since Euclidean and hyperbolic geometry are both true of their own spaces, but what it allowed a thinker to see that the previous stratum could not, and what it put out of view in exchange. The second takes the sequence to be driven by a tension rather than fixed by a definition. The tension is between the figure held fixed and the transformation that deforms it, and it does not resolve. It migrates, stratum by stratum, from a regime in which the figure is primary and transformation is a threat to its identity, to one in which transformation is primary and the standing figure is only a slow trajectory. The historical claim is narrow and checkable in its parts: each stratum detaches form one step further from physical extension, and the detachment can be dated, which section 8 does. The conceptual proposal, in section 7, is that the stratum now forming deserves a name and, more importantly, a test that keeps the name from being a slogan.

2. Survey and Figure

2.1 The administered earth

Geometry begins as an instrument of administration, and its first objects are the ones the name still holds: land, edge, distance, surface. The Greek word for the Egyptian surveyors, *harpedonaptai*, means rope-stretchers, and the reconstruction long attached to them is a rope knotted into twelve equal intervals, which when drawn taut into sides of three, four, and five units closes on a right angle. Whether or not the knotted rope was used exactly so, the surviving documents leave no doubt about the character of the knowledge. The Rhind papyrus, copied by the scribe Ahmes around 1650 BCE, takes the area of a circle to be the square of eight-ninths of its diameter, a rule accurate enough for granary and field and indifferent to any notion of proof. The Moscow papyrus computes the volume of a truncated pyramid, the shape of a real earthwork. Across Mesopotamia the Old Babylonian tablet known as Plimpton 322, from around 1800 BCE, tabulates what amount to Pythagorean triples in base sixty, a millennium before Pythagoras, as a working table rather than a theorem.

What unites these is that geometry is here a technology of place in the service of the state and the sky. It fixes boundaries so that possession is stable, computes areas so that grain can be taxed, and orients temples and calendars to the stars. Its propositions are recipes, justified by working, and its space is the one the surveyor occupies, a space whose defining property is that one can in principle cross it and measure it. That property, the walkable-and-measurable, is the intuition every later stratum will have to detach from, and its persistence underneath the abstractions is why a mathematician still draws a diagram of a space that has no diagram.

2.2 Euclid's subtraction

The second stratum is founded by an act of removal. Euclid's *Elements*, composed around 300 BCE, does not measure a field. It begins from twenty-three definitions, five postulates, and five common notions, and derives some 460 propositions in a deductive order in which each holds by proof from what precedes it and nothing is granted to inspection of a drawn figure. The point that has no part, the breadthless line, the circle every radius of which is equal: these are not idealizations of marks on papyrus but objects available only to reason, and the diagram serves merely as an aid to them while the objects remain elsewhere. Form has separated from matter. The triangle of the *Elements* cannot erode, cannot be redrawn larger, cannot be wrong, because everything about a surveyed plot that could change has been subtracted from it, and what remains is necessity.

The cost of that necessity is contact with the world, a debt the later strata spend 2,000 years repaying, and Euclid himself marks the place where the repayment will begin. Four of his postulates are immediate. The fifth, which states in effect that through a point not on a given line there passes exactly one parallel, is long, conditional, and was felt from antiquity to be of a different character, more a theorem in want of a proof than an axiom. Proclus records the unease in the fifth century, and it does not subside. The postulate sits in the foundations as a hostage to the future, and the attempt to discharge it, to show it follows from the rest, is what eventually breaks the assumption that there is only one geometry.

3. Coordinate and the Plural of Space

3.1 The algebraic capture of form

Across the two millennia after Euclid little of comparable depth is added, a silence section 8 measures, and the third stratum opens abruptly with Descartes' *La Géométrie* of 1637. Its decisive move is small and total. Descartes fixes a unit length, and with it the product of two segments can be construed as another segment rather than as an area, so that the operations of algebra apply directly to geometric magnitudes. A curve becomes the set of points whose coordinates satisfy an equation, a geometric problem becomes the manipulation of that equation, and the degree of the equation classifies the curve. Fermat reaches an equivalent method in the same years. Form has become computable, which is the precondition for every later space that will be given by numbers rather than drawn.

The capture has a consequence Descartes did not need but later strata would seize. Once a figure is an equation, its tie to the visible plane is severed, for an equation in three unknowns describes a surface and an equation in more describes a figure no eye will ever see. The coordinate both captures form and begins to release it from the dimensions of sight, and the higher-dimensional spaces that the nineteenth and twentieth centuries will treat as ordinary are already licensed, in principle, the moment a curve becomes a formula.

3.2 The fifth postulate breaks

The pluralization of space proceeds through a failed proof. In 1733 the Jesuit Girolamo Saccheri, in *Euclides ab omni naevo vindicatus*, sets out to vindicate Euclid by assuming the parallel postulate false and deriving a contradiction. He assumes, in the relevant case, that through the point there pass many parallels, and grinds out theorem after theorem, expecting absurdity. None comes. He proves, without recognizing them as such, a long series of true propositions of the geometry he means to refute, and abandons the effort only by declaring the conclusion repugnant to the nature of the straight line, an aesthetic judgment standing in for the contradiction he never reached. Saccheri has done the geometry and refused to believe it.

What he refused, others affirm. Lobachevsky publishes a geometry in which the postulate fails in the Kazan Messenger of 1829, Bolyai independently in an appendix to his father's textbook in 1832, and Gauss, it later emerges from his papers, had reached the same results privately and withheld them, fearing the outcry. For decades the status of these systems is unsettled, since a body of theorems with no evident contradiction is not yet known to be consistent. The settlement comes from Beltrami. In 1868 he shows that the geometry of Lobachevsky is realized on a surface of constant negative curvature, the pseudosphere, and constructs the disk and half-plane interpretations now bearing his and Poincaré's names, in which the lines and distances of the new geometry are defined in terms of the old. The construction establishes a relative consistency: if Euclidean geometry harbours no contradiction, neither does the hyperbolic, because any contradiction in the one would translate into a contradiction in the other. With that, the new geometries are as sound as Euclid's, intuition is dethroned as the criterion of geometric truth, and the singular space becomes a plurality (Torretti, 1978). The episode is the visible front of a larger change that historians of mathematics read as the onset of mathematical modernism, the turn by which mathematics takes its own constructions, rather than the physical world, as its subject (Gray, 2008).

3.3 Riemann's intrinsic space

The plurality is given its general form in a single lecture. Riemann's *Habilitationsvortrag* of 1854, printed only in 1868, proposes that space is a particular case of what he calls a multiply extended magnitude, and that the geometry of such a magnitude is set by a rule for infinitesimal distance, the line element $ds^2 = \sum_{i,j} g_{ij}\, dx^i dx^j$, in which the coefficients g_{ij} may vary from point to point. Curvature is then a function of position computed from the g_{ij} alone. The precedent is Gauss, whose *Theorema Egregium* of 1828 had shown that the curvature of a surface is intrinsic, determinable by measurements made within the surface and unchanged by any bending that preserves distances, so that an ant confined to the surface could find it without reference to a surrounding space. Riemann removes the surrounding space entirely. A manifold of any dimension has its own metric and curvature from within, owing nothing to an embedding it may not possess, and the metric is not dictated in advance but is, in his words, a matter to be determined within the manifold itself (Riemann, 1868; trans. Clifford, 1873). The fixed ground of Euclid has become a field that varies, and the figure that the surveyor took as the immovable backdrop is now itself the variable.

3.4 The physical payoff

A construction of this abstraction might have remained internal to mathematics, and for sixty years it did. Then Einstein, in the general relativity of 1915 and 1916, identified the gravitational field with the curvature of a four-dimensional spacetime manifold, its metric determined by the distribution of matter and energy through the field equations, and the motion of bodies with the geodesics of that metric. Gravitation, which Newton had treated as a force acting across space, became a property of the shape of space itself, and the prediction of light bending near the Sun was confirmed in 1919. The geometrization is of the force and not of its cause, since matter and energy enter the equations from outside, as the source that bends the metric rather than as something themselves geometric, and the honest statement keeps the two apart. With that qualification, the episode is the strongest evidence the stratum offers that its abstractions were not idle: a piece of pure geometry, built with no physical motive, turned out to be the language in which gravity is written.

4. The Regime of Transformation

4.1 Klein's inversion

By the late nineteenth century there are many geometries, and the question is what unifies them. Klein's Erlangen program of 1872 answers by inverting the order of explanation. A geometry, he proposes, is given by a space together with a group of transformations acting on it, and its subject is precisely those properties "not altered by the transformations of the group." The figure is no longer the primary thing whose properties one studies; the group is primary, and a property counts as geometric insofar as the group preserves it. This dissolves the tension of the earlier strata into a single structure, because an invariant is now defined by reference to the transformations it survives, and has no content apart from them. The fixed and the moving are two faces of one object.

The inversion orders the geometries by inclusion of their groups, and the ordering runs against naive expectation. Projective geometry has the largest group and therefore the fewest invariants, keeping incidence and the cross-ratio while surrendering distance and angle and even the distinction between parallel and intersecting lines. Affine geometry, with a smaller group, recovers parallelism and ratios along a line. Euclidean geometry, with the smallest group, the rigid motions, has the most invariants of all, length and angle and congruence. A property belongs to the more general geometry the larger the group that leaves it alone, so that the metric notions which feel most basic are in fact the most special, the residue left when the group is shrunk as far as it will go. The program is not only a classification; through Klein's collaboration with Sophus Lie it ties geometry to the theory of continuous transformation groups, and the symmetry that will later organize physics and machine learning enters geometry here, as its defining principle rather than a feature of particular figures.

4.2 Identity without measure

Topology pushes the regime to the point where measurement disappears and identity remains. Its oldest result is Euler's, that for a convex polyhedron the vertices, edges, and faces satisfy $V - E + F = 2$, a number untouched by any deformation that preserves the pattern of incidences while distorting every length and angle. For a closed orientable surface the quantity generalizes to $V - E + F = 2 - 2g$, where the genus g counts the handles, and the genus alone classifies such surfaces up to continuous deformation: a sphere is not a torus because no bending without tearing turns zero handles into one, and a coffee cup is a torus because the handle is all that topologically matters. Poincaré's *Analysis Situs* of 1895 founds the systematic study, attaching to a space algebraic invariants, the fundamental group and the homology groups, that record its connectedness and its holes and are preserved under continuous maps. Form has become deformable identity, a sameness that holds through stretching and breaks only at a cut, and the metric that the previous strata had labored to make intrinsic is now set aside as inessential to what a shape is.

4.3 The Erlangen program returns

Klein's definition proved durable beyond geometry and beyond anything its author could have foreseen. The architectures that dominate machine learning are, in the reading of Bronstein and colleagues (2021), so many instances of the Erlangen principle, each fixed by the symmetry group of the data it is built to process. A convolutional network is built around translation, sharing its weights across positions so that a pattern recognized in one part of an image is recognized in any other; a graph network and the set-processing core of a transformer are built around permutation, so that the output does not depend on an arbitrary ordering of the inputs; networks on curved domains are built around the rotations of a local frame, the gauge symmetry of the manifold. The choice of which transformations a network is required to leave its outputs invariant or equivariant under is the choice of its inductive bias, the prior knowledge wired into its form. The design principle of the most recent stratum's machines is, almost verbatim, a definition that Klein framed in 1872 for figures in a plane, which is one measure of how far the regime of transformation reaches.

5. Spaces of States

5.1 Phase space

The next stratum makes the possible states of a system into a space and its history into a path. In the formulation consolidated by Gibbs in 1902, a mechanical system of n degrees of freedom is a single point in a phase space of $2n$ dimensions, one axis for each position and each momentum, and the system's entire evolution is a curve through that space governed by the Hamiltonian flow. The geometry of the flow is the physics. Liouville's theorem states that the flow preserves volume in phase space, so that a region of possible states is stirred but never compressed, which is the geometric content of the conservation of information in a reversible dynamics. Poincaré's recurrence theorem, from his work of around 1890, states that almost every trajectory in a bounded energy surface returns arbitrarily close to its start, a statement about the topology of the flow with no counterpart in the language of forces. The vocabulary that the study of such systems generates, the attractor toward which trajectories are drawn, the basin of states that share an attractor, the bifurcation at which the portrait of the flow changes qualitatively, is frankly geometric, and its referents are possible evolutions rather than figures in a space. With them becoming itself has become geometric.

5.2 Information geometry

The same gesture reaches a domain with no evident spatial character, the families of probability distributions. Information geometry, in the form Amari and Nagaoka give it (2000), treats a parametrized family of distributions as a manifold whose points are the distributions and whose metric is the Fisher information, which measures how sharply a change of parameter changes the distribution and so how distinguishable nearby distributions are from samples. The natural measure of separation between two distributions, the Kullback–Leibler divergence, is not a metric, since it is not symmetric and violates the triangle inequality, yet it induces on the manifold a pair of flat affine connections in duality with each other, a structure richer than the single Levi-Civita connection of ordinary Riemannian geometry. The practical payoff is the natural gradient: a learning rule that follows the steepest descent with respect to the Fisher metric, the geometry of the statistical manifold, rather than the steepest descent in the arbitrary coordinates of the parameters, and so moves in a way invariant to how the model happens to be parametrized. Inference and learning have become motion on a manifold of beliefs.

5.3 Space without points

At the far end of the stratum the geometric object sheds even its points. Grothendieck's reconstruction of algebraic geometry, systematized in the topos theory presented by Mac Lane and Moerdijk (1992), replaces a space considered as a set of locations with the system of its parts and the ways they cover one another. A sheaf on such a system records consistent local data that glue along overlaps, a topos is a category of such sheaves, and it behaves in every formal respect like a space, possessing its own internal logic, while in general having no underlying set of points at all. Space has become a mode of coherence among local pieces rather than a container of positions. By the close of this stratum the geometric object need not be visible, need not carry a metric, need not even have points; it need only carry structure that some class of transformations preserves, which is the most that survives of what the surveyor began with.

6. Spaces of Life and Mind

6.1 Morphospace

Biology supplies the stratum in which a form does not sit in a space but constitutes one. Raup's analysis of shell coiling (1966) parametrizes the logarithmic-spiral shells of molluscs and other groups by four numbers: the whorl expansion rate W, the distance D of the generating curve from the coiling axis, the rate T of translation along that axis per revolution, and the shape S of the generating curve. A choice of the four specifies a shell, and the resulting space, of which Raup plotted slices as a cube, contains every shell that this developmental rule can make. The striking content of the construction is the parts of the cube that are empty. Real ammonoids, gastropods, and bivalves cluster in some regions and avoid others entirely, and the vacancies are evidence, since a form that is geometrically available but never realized points to a developmental or functional constraint that forbids it, a wall in the space of the possible that selection never crosses. Theoretical morphology, the program McGhee (1999) names and surveys, makes this its method: construct the space of forms a generative rule allows, then ask why life occupies the part of it that it does.

The stratum comes with its own correction, and the correction matters for the argument to come. Mitteroecker and Huttegger (2009) show that morphospaces are not in general Euclidean, that the distance between two forms and the meaning of a straight line between them depend on how the space was built, and that the geometric intuition of evenly spaced coordinates can mislead precisely where it is most inviting. A morphospace is a genuine space only to the extent that its construction licenses the distances and directions one reads off it, and where that license is absent the spatial picture is decoration. The point returns in section 7 as the sharpest of the conditions a morphotopic treatment must meet.

6.2 Conceptual and semantic spaces

Cognition supplies the next stratum, and meaning is its territory. Gärdenfors's conceptual spaces (2000) model a concept not as a list of necessary and sufficient features but as a region in a space whose axes are quality dimensions, and the model's substantive claim is that natural concepts occupy convex regions, so that any form lying between two instances of a concept is itself an instance. The clean case is colour, whose dimensions of hue, saturation, and brightness form a space in which the colours grouped under a single word do fall in a convex region, and in which a prototype, the focal example, sits at the region's centre with membership grading off toward its boundary. Similarity has become proximity, and a category has acquired a shape.

Distributional semantics makes the same structure empirical and scalable, on the principle that a word's meaning is reflected in the company it keeps. Latent semantic analysis (Landauer & Dumais, 1997) extracts a space from the patterns of co-occurrence of words across documents; the neural language model of Bengio and colleagues (2003) learns a distributed vector for each word as a by-product of predicting text; and the word vectors of Mikolov and colleagues (2013) make such embeddings cheap and sharp enough that analogies appear as displacements, the step from man to woman roughly parallel to the step from king to queen. The parallelogram is a real and repeatable structure, and it is also a limited one, fragile under change of corpus and prone to spurious cases, which is the right note on which to hold the result. The embedding shows that meaning, learned at scale, leaves a geometric trace rich enough to compute with. It does not yet show that the geometry is doing the work rather than recording it, and the distinction is the hinge of the next subsection.

6.3 From trace to mechanism

The decisive recent development is that the geometry has become something one can read off a working system and intervene upon, in brains and in machines, so that it is no longer a description laid over the computation but the computation made visible. In neuroscience, representational similarity analysis (Kriegeskorte, Mur, & Bandettini, 2008) characterizes a population of neurons by the matrix of dissimilarities among its responses to a set of stimuli, turning the representational geometry of a brain region into a measured object that can be compared across regions, animals, and models. In artificial networks the evidence is sharper still, because the system is fully observable and can be edited. Gurnee and Tegmark (2023) recover from a language model linear directions that encode the latitude and longitude of places and the date of events, with individual units that track these coordinates, the raw material of a world model. Li and colleagues (2023) train a network on nothing but sequences of legal moves in the board game Othello, never showing it a board, and find within its activations a representation of the current position that a probe can read and that, when edited, changes the model's subsequent play to match the altered board.

The pattern sharpens to a mechanism in the case of arithmetic. Nanda and colleagues (2023), reverse-engineering a small transformer trained on addition modulo a prime, find that it embeds each number as a rotation, a point on a circle given by the sine and cosine of an angle, and computes the sum by composing rotations through the trigonometric identities for the sine and cosine of $a + b$, then reading the angle of the result. Addition has been turned into rotation. Engels and colleagues (2024) find the same circular form arising on its own in large models for the days of the week and the months of the year, irreducibly two-dimensional features that the network turns to do modular arithmetic over the calendar. Kantamneni and Tegmark (2025) generalize the circle to a helix, showing that a model represents the integers along an axis that is linear in the number and wound by several periods at once, and that it adds two integers by rotating their helices into the helix of the sum and reading the answer off, an operation they name the Clock algorithm and confirm by intervening directly on the helical directions. Arithmetic, which looks nothing like geometry, is carried out as trigonometry, and the claim is not interpretive but causal, since perturbing the geometry changes the answer. Huh and colleagues (2024) argue that the convergence is general, that networks trained on different data and even different modalities come to measure similarity among inputs in increasingly the same way, as if toward one shared representational geometry of the world. In the sharpest of these cases the geometry is not a trace that meaning leaves behind but the machinery that does the work, open to measurement and to intervention.

6.4 Navigation as basal cognition

The strata of life and mind are gathered by Levin's framework (2022) into a single claim, that problem-solving is not confined to the behavioural space of an animal moving through its environment but runs equally in the anatomical, physiological, and transcriptional spaces of the body, so that a developing or regenerating tissue navigating toward a target form and correcting when displaced from it is performing a basal cognition, defined as the competent steering of a system through a problem space toward adaptive regions despite perturbation. The claim is concrete in the planarian flatworm. The target morphology toward which its cells rebuild after amputation is encoded not only in the genome but in a stable pattern of bioelectric voltage across the tissue, and a brief manipulation of that pattern can permanently rewrite the target, so that a cut fragment regenerates two heads rather than one and, left undisturbed, continues to regenerate two heads through later rounds of cutting, its genome untouched (Durant et al., 2017). Levin's own description states the geometry without prompting: the cells minimize the anatomical distance in morphospace between the current and the target form. The setpoint is a point in the space of possible bodies, the bioelectric field is the controller, and regeneration is navigation. With that, the disparate occupants of this whole stratum, the organism, the concept, the word, and the developing body, share a single form. Each is a trajectory in a space of possibilities, moving under constraint toward regions that count, for it, as success.

7. Morphotopy, and a Test for It

7.1 The name and its object

The strata no longer share an object with earth-measure. What they study is how forms become possible, how they persist and deform, and how systems move among them under constraint, and the move that produces a new one, of which the strata of section 6 are the latest examples, has no settled name. The name proposed here is *morphotopy*: the study of how forms, organisms, meanings, and agents occupy, generate, and traverse spaces of possible form under constraint. Its object is not a space but the production and navigation of spaces. Where a geometry asks what the structure of a given space is, morphotopy asks how a space of possibility comes to exist for some system, what forms can inhabit it, what transformations it admits, what constraints bound it, and what within it counts, for that system, as a destination.

7.2 Five conditions

A term of this breadth is justified only if it does more than relabel, and the literature already holds a long list of nominal geometries of this and that in which the spatial language does no work and could be deleted without loss. The proposal is therefore stated as a test, five conditions a treatment must meet to be morphotopic rather than decorative. There must be a space of possible forms or states, the unrealized ones included, so that what never happened still has a location. There must be a notion of nearness or cost, a topology or a metric, whose claims about distance and direction are licensed by the construction and not assumed, which is the condition the morphospace literature isolates (Mitteroecker & Huttegger, 2009). There must be a class of admissible transformations, the moves the system can make. There must be constraints that bound the space and forbid regions of it. And there must be a notion of navigation, the goals, attractors, or valued regions toward which trajectories run, so that motion has a direction that matters.

The conditions do work, because real cases pass and fail them in instructive ways. Raup's morphospace passes all five: the cube of unrealized shells is the space, the developmental parameters license a meaningful nearness, growth supplies the transformations, the empty regions are the constraints, and selection toward viable form is the navigation. A bare word embedding passes the first and, with care, the second, but fails the fifth until a task is added, since a static cloud of vectors has no destinations; this is why the embedding is an instrument and the trained network that uses it to compute is the morphotopic object. A loose claim that some domain is a landscape, with peaks and valleys named but no licensed metric and no admissible moves specified, fails at the second and third conditions and is exposed by the test as metaphor. The fifth condition, navigation, is what separates a morphotopic space from a merely descriptive one, and the second, the licensed metric, is the filter that catches the largest number of impostors.

7.3 What it is not

Stated this way, morphotopy is broader than several existing programs and contains them as cases. Geometric deep learning is the Erlangen principle applied to the architecture of networks, a morphotopy of one engineered domain. Representational geometry is the measurement of the spaces inside brains and trained models, the empirical arm. Conceptual spaces is the cognitive instance, and Levin's navigation of problem spaces is the biological one. The proposal does not displace or improve any of these, each of which is sharper in its own domain than a general framework could be. Its contribution is twofold and modest: to recognize that these are not independent borrowings of spatial language but instances of one historical movement with a datable trajectory, and to supply the criterion that distinguishes a substantive instance from a spatial figure of speech. It is an organizing claim with a test attached, and it leaves the mathematics of each domain where it found it.

7.4 One recurring obstruction: frustration

If the strata share a movement, one might ask whether they share any structural content, and at least one obstruction is definable wherever a space of forms is placed under competing constraints, though the computation of section 8 exhibits it in only one of the strata. In a frustrated system the geometry of the interactions makes it impossible to satisfy all local constraints at once. The canonical case is the triangular antiferromagnet, in which each pair of neighbouring spins is happiest pointing opposite ways, but three around a triangle cannot all disagree, so one bond is always unsatisfied. In the spin ice of certain pyrochlore crystals the analogous local rule leaves a vast number of equally good ground states rather than a single best one, and the resulting residual disorder gives a measurable zero-point entropy of the kind Pauling (1935) had identified in ordinary water ice (Bramwell & Gingras, 2001). Stripped of the physics the condition is

$$\min_{x \in X} \sum_i E_i(x) \ > \ \sum_i \min_{x \in X} E_i(x),$$

that the best attainable global state is strictly worse than the sum of the separately attainable local optima. The right side is always a lower bound on the left, and equality holds exactly when some single configuration minimizes every local term together, the unfrustrated case; strict inequality is frustration. The statement is about constrained spaces as such and not about magnets, which is why the same obstruction appears, under other names, in optimization, in constraint satisfaction, and in any system whose local preferences cannot be jointly met, and why a space of possible forms generically fails to admit a configuration that is locally ideal everywhere. Section 8 exhibits the inequality exactly, in the smallest system that can carry it.

7.5 The pattern beyond physics and code

The movement is not confined to physics, biology, and machine learning, and three further cases, drawn from domains that share nothing else, show the test being met outside any of the strata already traced. A chord of n notes is a point in the quotient of the n-torus by the symmetric group, T^n/S_n, an orbifold whose singular boundaries correspond to chords with doubled or evenly spaced notes; Tymoczko (2006) shows that the chords near these singularities are the ones that admit efficient voice leading, short moves to a neighbouring chord, so that the harmonic practice composers arrived at by ear tracks the geometry of the space, and consonant triads cluster near its centre. The possible evolutionary trees on a fixed set of species form the space of Billera, Holmes, and Vogtmann (2001), assembled from one orthant per tree topology, glued along the lower-dimensional faces where a branch length falls to zero and the topology can change; the space has globally non-positive curvature, which guarantees a unique shortest path between any two trees and hence a well-defined average of a collection of them, turning the comparison and summary of competing phylogenies into a geometric operation. The shape of a dataset is read by persistent homology (Carlsson, 2009), which grows a family of spaces from the data by connecting points within a radius that increases, and records the connected components, loops, and voids that appear and disappear as it grows, the features that persist across a wide range of radii being the robust shape and the short-lived ones being noise. Each of the three meets the five conditions, with a space of the possible, a licensed metric, admissible moves, constraints, and a notion of which regions matter, and that is why they belong to one stratum rather than to a coincidence of vocabulary.

8. The Strata, Measured

Two small computations accompany the argument, and their role is exact and limited. Both are reproduced by the code in `simulation/`, which writes every number below to `simulation/output/results.json` and the two figures to `simulation/output/figures/`. They illustrate two specific claims, the acceleration of the stratigraphy and the existence of frustration, and neither is a measurement of history or of mind; each reaches as far as its construction and no further.

8.1 Dating the strata

The first computation assigns to each of sixteen strata a representative work and date, from the earth-measure of the second millennium BCE to the present, and computes their order, span, and gaps. The dating is conventional and the earliest stratum deliberately coarse, since earth-measure has no single date, and what is exact is the arithmetic on the chosen dates rather than the dating of a diffuse history. The count of strata is a choice too, and a consequential one: individuating the recent phases finely while treating antiquity in a few broad blocks builds part of the acceleration into the stratification before any date is read, so the contraction reported below is a property of the chosen granularity as much as a fact about the history. So qualified, the strata span 3826 years, and the largest single interval between consecutive strata is the 1937 years from Euclid's figure to Descartes' coordinate, the long plateau that section 3 described. The distribution is sharply skewed toward the present. Taking the collapse of the parallel postulate in 1829 as the threshold of the modern regime, 12 of the 16 strata fall within the final 197 years, three quarters of them inside roughly the last twentieth of the span. Two further statistics fix the acceleration. Half the strata exist only by 1895, late in a history that opens before 1800 BCE. And the mean interval between consecutive strata, measured within each era, is about 1146 years before the threshold and about 18 years after it, a contraction of roughly 64 times. The detachment of form from extension is not a steady drift but an acceleration, and the recency of most of the strata, several still actively contested, is part of why the stratum now forming has no agreed name.

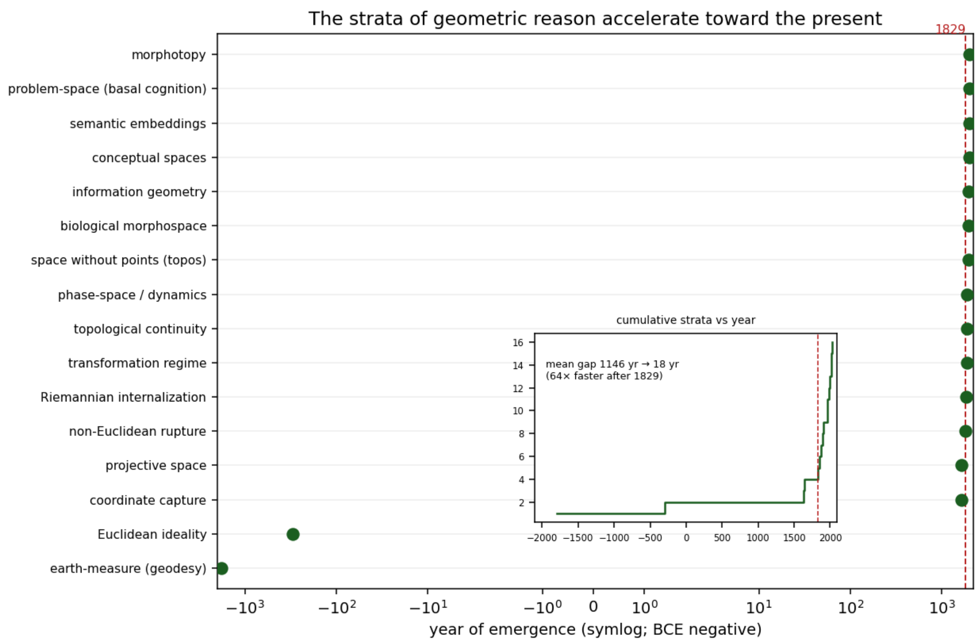

Figure 1: The sixteen strata of geometric reason by date of emergence (symlog axis; BCE negative). Millennia separate the early strata, and twelve of the sixteen fall after the non-Euclidean rupture of 1829. The inset plots the cumulative count against a linear time axis: flat for three millennia, then nearly vertical, the mean interval between strata contracting from about 1,146 years to about 18.

8.2 Frustration, exactly

The second computation exhibits the frustration inequality by exhaustively enumerating the states of two small antiferromagnetic systems, in which a bond between two spins is satisfied when they point opposite ways. On a triangle, an odd cycle, the three bonds cannot be satisfied together. The enumeration confirms it: the lowest energy attainable is -1.0 in units of the coupling, while the three bonds taken separately could reach -3.0, so the global optimum lies 2.0 above the sum of the local optima and the inequality holds strictly. The frustration leaves a signature in the set of best states, which is sixfold degenerate, with exactly one of the three bonds left unsatisfied in each, and a residual entropy of 0.5973 per spin. The square, an even cycle and therefore free of frustration, is the control: its four bonds are satisfied together, its lowest energy of -4.0 equals the sum of the local minima, no bond is left unsatisfied, and its best state is only twofold degenerate, with a residual entropy of 0.1733 per spin. A single triangle is enough to establish that local optima need not compose into a global one, which is the elementary and exact form of an obstruction that any sufficiently constrained space of possible forms can hold. That the triangle's behaviour is not an artifact of its size can be checked by scaling it up. On a periodic triangular lattice, enumerated exactly at $N = 9$, 12, and 16 spins, the fraction of bonds left unsatisfied in the ground state is exactly one third at every size, the same value the lone triangle shows, and the ground state is extensively degenerate, with ground-state degeneracy 42, 68, and 42 at the three sizes, large and not monotonic in N. The residual entropy per spin therefore stays finite, ranging across these small tori from 0.5973 down to 0.2336 and straddling the exact thermodynamic value of 0.3231 per spin that Wannier (1950) obtained for this lattice. The obstruction the single triangle exhibits is the seed of a macroscopic fact.

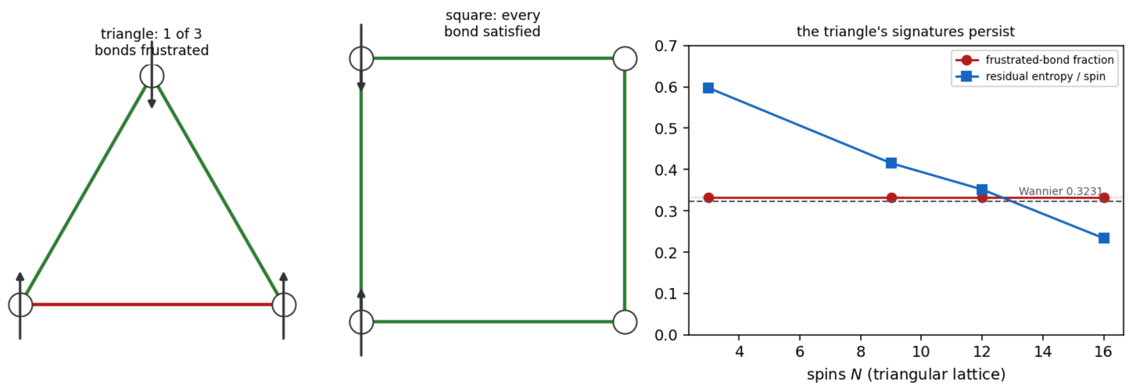

Figure 2: Geometric frustration by exact enumeration. Left: a triangle in a ground state, two spins up and one down, leaves one of its three antiferromagnetic bonds unsatisfied (red), while the bipartite square satisfies all four (green). Right: on the periodic triangular lattice the frustrated-bond fraction stays exactly one third at every size, and the residual entropy per spin remains finite, scattering around Wannier's thermodynamic value of 0.3231.

What the model does not do is measure the frustration of any particular applied space; it shows that the obstruction is real and structural, and leaves its appearance in semantic, developmental, or social spaces to be established case by case.

9. The Test, the Boundary, and the New Terrains

9.1 Predictions, and how the proposal could fail

The proposal is organizational and historical, and stating its limits is part of stating it. It is not a new mathematics, and it does not predict which domain will next be rendered as a space, though it expects, from the acceleration of section 8, that the rendering will recur and probably soon. Its empirical content lies in the test, which is what makes it more than a redescription. The test can fail, and its failure is informative. If a domain is dressed in spatial language but no licensed metric and no admissible transformations can be supplied for it, the test excludes the case, and a programme of nominal geometries that supplied no such structure would be revealed as decoration rather than refuted by argument. Conversely, the claim that the strata form one movement would be weakened if the cases that pass the test shared nothing beyond passing it, and strengthened, as section 7.4 argues, by a structural obstruction such as frustration that is definable across them. The criterion is the part of the proposal that can be held to account; the name is the part that can be argued over.

9.2 The boundary at experience

There is one further claim the sequence seems to invite and that must be refused. A morphotopic account can capture a great deal of the structure of a mind, the model a system maintains of itself, the integration of its information, the continuity of its trajectory through time, the value it attaches to regions of its space of actions, and the work of section 6.3 shows how much of that structure is already legible, and editable, in trained systems. None of it reaches the question Chalmers (1995) isolated, why any of that structure should be accompanied by experience at all rather than proceeding in the dark. The geometry of selfhood is not an explanation of consciousness, and the two are easy to conflate because the structural story is becoming so detailed. Keeping them apart strengthens the programme rather than conceding anything, and it is what the five conditions exist to enforce, since a morphotopy that claimed to have geometrized experience would have abandoned the demand that its spaces and metrics be licensed by what they actually explain.

9.3 The first terrain and the new ones

Geometry began with the earth because the earth was the first thing human knowledge was compelled to make measurable, the first space that had to be divided, owned, and predicted. Its long history is the discovery that measurability is not confined to the earth, nor to physical space, nor to anything one can walk across, but extends to ideal figures, to curved and pointless spaces, to the states of dynamical systems, to the forms a body can take, to the meanings a word can bear, and to the representations a brain or a machine builds of its world. Morphotopy is the name for the stratum in which that discovery turns on itself and the production and navigation of such spaces becomes the subject of the science, and the test set out here is what is meant to keep the name from outrunning what it can show. The earth was only the first terrain. Life, mind, meaning, and the machines now learning to represent them are the new ones.

References

Amari, S., & Nagaoka, H. (2000). *Methods of Information Geometry*. American Mathematical Society & Oxford University Press.

Beltrami, E. (1868). Saggio di interpretazione della geometria non-euclidea. *Giornale di Matematiche*, 6, 284–312.

Bengio, Y., Ducharme, R., Vincent, P., & Jauvin, C. (2003). A neural probabilistic language model. *Journal of Machine Learning Research*, 3, 1137–1155.

Billera, L. J., Holmes, S. P., & Vogtmann, K. (2001). Geometry of the space of phylogenetic trees. *Advances in Applied Mathematics*, 27(4), 733–767.

Bramwell, S. T., & Gingras, M. J. P. (2001). Spin ice state in frustrated magnetic pyrochlore materials. *Science*, 294(5546), 1495–1501.

Bronstein, M. M., Bruna, J., Cohen, T., & Veličković, P. (2021). Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv:2104.13478*.

Carlsson, G. (2009). Topology and data. *Bulletin of the American Mathematical Society*, 46(2), 255–308.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.

Clifford, W. K. (1873). On the hypotheses which lie at the bases of geometry [translation of Riemann, 1854]. *Nature*, 8, 14–17, 36–37.

Desargues, G. (1639). *Brouillon project d'une atteinte aux événements des rencontres d'un cône avec un plan* [Rough draft on conics].

Descartes, R. (1637). *La Géométrie* (appendix to *Discours de la méthode*).

Durant, F., Morokuma, J., Fields, C., Williams, K., Adams, D. S., & Levin, M. (2017). Long-term, stochastic editing of regenerative anatomy via targeting endogenous bioelectric gradients. *Biophysical Journal*, 112(10), 2231–2243.

Einstein, A. (1916). Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354(7), 769–822.

Engels, J., Liao, I., Michaud, E. J., Gurnee, W., & Tegmark, M. (2024). Not all language model features are linear. *arXiv:2405.14860*.

Euclid. (1956). *The Thirteen Books of Euclid's Elements* (T. L. Heath, Trans.; 2nd ed.). Dover. (Composed c. 300 BCE.)

Gärdenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. MIT Press.

Gauss, C. F. (1828). Disquisitiones generales circa superficies curvas. *Commentationes Societatis Regiae Scientiarum Gottingensis Recentiores*, 6, 99–146.

Gibbs, J. W. (1902). *Elementary Principles in Statistical Mechanics*. Charles Scribner's Sons.

Gray, J. (2008). *Plato's Ghost: The Modernist Transformation of Mathematics*. Princeton University Press.

Gurnee, W., & Tegmark, M. (2023). Language models represent space and time. *arXiv:2310.02207*.

Huh, M., Cheung, B., Wang, T., & Isola, P. (2024). The platonic representation hypothesis. In *Proceedings of the 41st International Conference on Machine Learning*. *arXiv:2405.07987*.

Kantamneni, S., & Tegmark, M. (2025). Language models use trigonometry to do addition. *arXiv:2502.00873*.

Klein, F. (1872). *Vergleichende Betrachtungen über neuere geometrische Forschungen* [Erlangen Program]. Deichert. (English translation: M. W. Haskell, 1893, *Bulletin of the New York Mathematical Society*, 2, 215–249.)

Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis: Connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2, 4.

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211–240.

Levin, M. (2022). Technological Approach to Mind Everywhere: An experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience*, 16, 768201.

Li, K., Hopkins, A. K., Bau, D., Viégas, F., Pfister, H., & Wattenberg, M. (2023). Emergent world representations: Exploring a sequence model trained on a synthetic task. In *International Conference on Learning Representations*. *arXiv:2210.13382*.

Lobachevsky, N. I. (1829). O nachalakh geometrii [On the principles of geometry]. *Kazan Messenger*.

Mac Lane, S., & Moerdijk, I. (1992). *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer.

McGhee, G. R. (1999). *Theoretical Morphology: The Concept and Its Applications*. Columbia University Press.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv:1301.3781*.

Mitteroecker, P., & Huttegger, S. M. (2009). The concept of morphospaces in evolutionary and developmental biology: Mathematics and metaphors. *Biological Theory*, 4(1), 54–67.

Nanda, N., Chan, L., Lieberum, T., Smith, J., & Steinhardt, J. (2023). Progress measures for grokking via mechanistic interpretability. In *International Conference on Learning Representations*. *arXiv:2301.05217*.

Pauling, L. (1935). The structure and entropy of ice and of other crystals with some randomness of atomic arrangement. *Journal of the American Chemical Society*, 57(12), 2680–2684.

Poincaré, H. (1895). Analysis Situs. *Journal de l'École Polytechnique*, (2)1, 1–123.

Raup, D. M. (1966). Geometric analysis of shell coiling: General problems. *Journal of Paleontology*, 40(5), 1178–1190.

Riemann, B. (1868). Über die Hypothesen, welche der Geometrie zu Grunde liegen. *Abhandlungen der Königlichen Gesellschaft der Wissenschaften zu Göttingen*, 13, 133–150. (Lecture delivered 1854.)

Saccheri, G. (1733). *Euclides ab omni naevo vindicatus*. Paolo Antonio Montano.

Torretti, R. (1978). *Philosophy of Geometry from Riemann to Poincaré*. D. Reidel.

Tymoczko, D. (2006). The geometry of musical chords. *Science*, 313(5783), 72–74.

Wannier, G. H. (1950). Antiferromagnetism. The triangular Ising net. *Physical Review*, 79(2), 357–364.