

The Amplitwist

A Tempered Account

Flyxion

Independent Researcher

2026

Abstract

The amplitwist program proposes a common geometric language for transformations that combine changes of magnitude with changes of orientation. Its mathematical nucleus is elementary and exact: a strict amplitwist has the form $A = sR(\theta)$, composes according to $s_1R(\theta_1)s_2R(\theta_2) = s_1s_2R(\theta_1 + \theta_2)$, and excludes anisotropic amplification. The program subsequently extends this language through circuit theory, real representations of quantum mechanics, cortical dynamics, consciousness, cultural change, artificial-intelligence alignment, and, in later work, an ordered Lamphron–Amplitwist–Sheaf Morphism sequence. These extensions do not have uniform epistemic status.

This essay reconstructs the program while preserving those differences. It begins with the handwritten “ultimately equal” relation,

$$A \asymp B \iff \lim_{\varepsilon \rightarrow 0} \frac{A}{B} = 1,$$

and derives the strict amplitwist algebra. Circuit applications provide the strongest physical realization because standard impedance theory already decomposes transformations into magnitude and phase, while rectification, state compression, logical irreversibility, Landauer’s bound, and Johnson–Nyquist noise supply independently checkable results. A methodological boundary is then imposed: later domains receive no evidential credit merely from appearing after this derived core.

Beyond that boundary, the quantum flag-space construction remains an unverified reconstruction program. The cortical identification receives a specific negative result: $J = \mathcal{A}R$ with general symmetric $\mathcal{A}$ is not a strict amplitwist unless

$$\mathcal{A}_0 = \mathcal{A} - \frac{\text{tr } \mathcal{A}}{2} I = 0.$$

The consciousness proposal consequently remains downstream of an unmet operator-identification condition and an additional unproved coherence transition. By contrast, the predicted rotation of prefrontal population vectors during task switching is retained as a falsifiable but untested empirical claim. Cultural curvature and amplitwist-based AI alignment remain programmatic extensions.

Close examination of the later Lamphron–Amplitwist–Sheaf Morphism sequence upgrades both Lamphron and Sheaf Morphism from terminological to constraining

but unverified, while leaving the ordered composition itself at the level of taxonomy. A commuting toy realization shows that operator algebra alone does not force the order, and direct examination of the primary sources supplies no domain or precondition theorem that does so. The resulting account is neither a defense nor a rejection of the unifying program. It is an attempt to preserve derivations, expose type mismatches, distinguish predictions from interpretations, and convert unsupported transitions into explicit proof obligations.

Introduction

A mathematical idea can travel across domains in two importantly different ways. In the first, the same restricted structure is independently recovered in each new setting. The recurrence then becomes evidence for a genuine common formalism. In the second, terminology travels more easily than constraints. A transformation involving gain becomes an “amplitude,” a change of direction becomes a “twist,” local agreement becomes “coherence,” and the recurrence of the vocabulary creates an appearance of mathematical unity before the corresponding equivalences have been established.

The amplitwist program contains examples of both processes.

At its narrowest, an amplitwist is unproblematic. A transformation

$$A = sR(\theta) \tag{1}$$

uniformly scales and rotates the plane. This is the familiar geometry of multiplication by a nonzero complex number. Composition remains within the same class, amplitudes multiply, angles add, and the transformation preserves angles. In AC circuit theory, the same decomposition appears directly in complex impedance, where magnitude and phase already have independently established physical meanings.

The broader program asks whether this geometry is more than a convenient description of those familiar cases. It extends amplitwist language into real formulations of quantum mechanics, cortical dynamics, theories of conscious coherence, linguistic and cultural evolution, and artificial-intelligence alignment. A later research thread places Amplitwist between Lamphron and Sheaf Morphism in the proposed sequence

$$L \longrightarrow A \longrightarrow S. \tag{2}$$

The motivating intuition is powerful because the domains appear to share a recurrent pattern. Something has a magnitude. Something has an orientation. Transformations alter both. Local transformations compose. Some distinctions disappear under observation or admissibility. Other distinctions accumulate into obstruction, curvature, or irreversible history. It is therefore tempting to regard the recurrence of these motifs as evidence that a single operator has been found at multiple scales.

That conclusion requires more than recurrence.

The central problem is *credibility inheritance*. Suppose an essay first proves a mathematical theorem, then presents a correct application in electrical engineering, then introduces a neuroscientific model, and finally proposes consequences for consciousness and culture. Even when no explicit deduction connects the domains, their order can produce the impression that each later claim inherits some of the certainty of the preceding one. The rhetorical sequence

$$\text{mathematics} \rightarrow \text{physics} \rightarrow \text{neuroscience} \rightarrow \text{consciousness} \rightarrow \text{culture} \quad (3)$$

can begin to function as an implicit proof.

It is not one.

The appropriate response is not uniform skepticism. Treating every amplitwist claim as speculative would destroy distinctions just as surely as treating every claim as established. The composition theorem for $sR(\theta)$ does not become less true because a cortical application later fails a definitional test. A falsifiable neuroscientific prediction does not become worthless because the interpretation that motivated it remains conjectural. Conversely, an exact result about complex impedance does not become evidence for a theory of semantic change merely because both can be described using magnitude and orientation.

The required discipline is local calibration. Each claim should carry the status warranted by the argument available for that claim.

Five statuses will recur. A claim is *derived* when it follows from stated definitions or established theory. It is an *unverified extension* when a substantive formal proposal exists but the stronger cross-domain identification has not undergone the relevant independent test. It receives a *type-mismatch result* when a definite requirement has been examined and does not follow from the proposed formulation. A claim is *falsifiable but untested* when an observable, expected result, and meaningful failure condition can be specified even though the experiment has not been performed. Finally, a *programmatic extension* defines a research direction and possible formalism without yet establishing the distinctive constraint required for validation.

These statuses are not positions on a single scale of quality. They describe different logical situations.

The distinction is particularly important because the program itself begins with a notation for refusing premature identification. In the surviving handwritten source, the relation described as “ultimately equal” is represented here by $\asymp$ and defined by

$$A \asymp B \iff \lim_{\varepsilon \rightarrow 0} \frac{A}{B} = 1. \quad (4)$$

The two quantities need not be equal at finite scale. Their distinction disappears only under a specified limiting operation. The notation therefore contains, in miniature, the methodological principle needed to evaluate the later theory: resemblance does not authorize identification. The operation under which a distinction disappears must be stated.

The same requirement applies when amplitwist language crosses domains. If a corti-

cal transformation is claimed to instantiate the same operator as complex multiplication, the defining restriction of that operator must be checked. If cultural change is claimed to possess curvature, the state space, transport rule, and path dependence must be defined. If three operators are claimed to form a necessary causal sequence, some feature of their algebra, domains, or preconditions must make competing orders fail.

This essay therefore retains the domain progression of the amplitwist program while interrupting its rhetorical momentum. The first three sections establish the notational, algebraic, and circuit-theoretic core. After that point a hard methodological boundary is imposed. Quantum mechanics is treated as an unverified reconstruction program rather than as a consequence of the preceding circuit results. The cortical application is evaluated through a precise isotropy criterion. The consciousness proposal is then considered downstream of that result rather than downstream of the elementary amplitwist theorem. The population-vector prediction is separated from the interpretation that generated it so that its empirical content can be stated independently. Cultural and alignment applications are preserved as research programs without being promoted into established consequences.

The final section turns to the Lamphron–Amplitwist–Sheaf Morphism sequence. That sequence poses a different question from the earlier cross-domain applications. Its question is not simply whether amplitwist-like language can be useful in several domains. It asks whether relaxation, amplitwist transformation, and sheaf-theoretic coherence constitute a necessary ordered causal architecture. Keeping those questions separate prevents a negative result about composition order from being projected backward onto earlier amplitwist mathematics, while also preventing earlier mathematical successes from being projected forward onto the causal sequence.

The resulting account is intentionally uneven. Some claims survive unchanged. Some become conditional. Some acquire explicit counterconditions. Some remain open. Two operators initially suspected of being merely terminological turn out to contain genuine constraining mathematics. One proposed operator identification fails to follow from its own formal definition. A universal narrative becomes a collection of independently assessable claims.

That is not a defect in the resulting theory. It is the information produced by examining it critically.

The appropriate standard for a unifying mathematical program is not that every domain must receive the same verdict. It is that the same distinctions between definition, derivation, representation, prediction, and conjecture remain operative everywhere. A theory that claims to describe admissibility and the preservation of distinctions should be especially capable of preserving those distinctions in its own epistemology.

The amplitwist can therefore be examined most sympathetically by refusing to grant it more unity than has actually been derived.

1. Ultimately Equal: The Infinitesimal Seed

Status: derived notational foundation. The amplitwist program begins with a smaller mathematical idea: a relation intended to distinguish ordinary approximation from equality obtained in a limiting process. In the surviving handwritten source, the relation is described verbally as “ultimately equal.” Using $\asymp$ as the typographic stand-in for the handwritten operator, its definition is

$$A \asymp B \iff \lim_{\varepsilon \rightarrow 0} \frac{A}{B} = 1. \quad (5)$$

The qualification “ultimately” is essential. The relation does not assert that $A = B$ at finite ε , nor does it merely indicate that two quantities are informally close. It asserts that their multiplicative discrepancy disappears in the stated limit. For example, if

$$A(\varepsilon) = 1 + \varepsilon, \quad B(\varepsilon) = 1, \quad (6)$$

then $A(\varepsilon) \neq B(\varepsilon)$ for every $\varepsilon > 0$, while

$$\lim_{\varepsilon \rightarrow 0} \frac{A(\varepsilon)}{B(\varepsilon)} = \lim_{\varepsilon \rightarrow 0} (1 + \varepsilon) = 1. \quad (7)$$

Hence

$$A \asymp B. \quad (8)$$

The distinction matters because the notation arose in an explicitly geometric context. The surviving account of its transcription records that the handwritten sign was initially normalized into a conventional approximation or asymptotic symbol. The correction was that the operator had been chosen deliberately to mark a limiting relation rather than an unspecified resemblance. The accompanying explanation described finite measurable increments becoming their differential counterparts under infinitesimal magnification. An early informal account of this idea inflated it into the stronger claim that two “literal geometric shapes” perfectly align at the microscopic limit. That stronger language is unnecessary. The quotient-limit definition already supplies the precise mathematical content.

The same handwritten sheet records two elementary consequences of the definition. The first is multiplicative compatibility:

$$X \asymp Y \quad \text{and} \quad P \asymp Q \implies XP \asymp YQ. \quad (9)$$

This follows directly. Where the relevant quotients are defined,

$$\frac{XP}{YQ} = \frac{X}{Y} \frac{P}{Q}. \quad (10)$$

Therefore,

$$\lim_{\varepsilon \rightarrow 0} \frac{XP}{YQ} = \left(\lim_{\varepsilon \rightarrow 0} \frac{X}{Y} \right) \left(\lim_{\varepsilon \rightarrow 0} \frac{P}{Q} \right) \quad (11)$$

$$= 1 \cdot 1 \quad (12)$$

$$= 1, \quad (13)$$

and consequently $XP \asymp YQ$.

The second recorded rule is

$$A \asymp BC \iff \frac{A}{B} \asymp C. \quad (14)$$

Again, this is not an additional axiom. Both sides reduce to the same limiting quotient:

$$\frac{A}{BC} = \frac{A/B}{C}, \quad (15)$$

subject to the ordinary nonvanishing conditions required for the divisions to be defined. The handwritten page therefore already contains a small calculus for the relation. Products preserve ultimate equality, and multiplicative factors may be transferred across the relation by division.

A second handwritten construction gives the notation its infinitesimal geometric motivation. Let

$$T = \tan \theta. \quad (16)$$

For a sufficiently small angular displacement $\delta\theta$ at radius L , in the spirit of the geometric approach to calculus developed by Needham (1997; 2021), the corresponding arc displacement satisfies

$$\delta s \asymp L \delta\theta. \quad (17)$$

At finite scale, the relevant arc, chord, and tangent displacements remain geometrically distinct. Their differences become negligible relative to the quantities themselves in the limiting construction. The notation is therefore useful precisely because it does not require premature identification of finite objects that become equivalent only after passage to the limit.

For the tangent construction,

$$T = \tan \theta, \quad (18)$$

so a small change in angle gives

$$\delta T \asymp \sec^2 \theta \delta\theta. \quad (19)$$

Passing to the differential limit yields

$$\frac{dT}{d\theta} = \sec^2 \theta = 1 + \tan^2 \theta = 1 + T^2. \quad (20)$$

The status of these claims is *derived*. Once $\approx$ is defined by the quotient limit, its multiplicative rules follow from ordinary limit laws, and the tangent example belongs to standard differential geometry and calculus. No cross-domain interpretation is needed to establish them.

What this foundation contributes to the later amplitwist program is therefore methodological rather than cosmological. It establishes a disciplined distinction between exact equality at finite scale and equivalence that emerges under a specified limiting operation. It also establishes an early preference for interpreting algebraic notation through the geometric transformation it compresses. The latter preference becomes central when complex multiplication is interpreted not merely as arithmetic on ordered pairs, but as the composition of two explicit geometric actions: scaling and rotation.

2. The Amplitwist as Derivable Structure

Status: derived. The amplitwist proper begins with a familiar fact about complex multiplication. A nonzero complex number can be written in polar form as

$$z = se^{i\theta}, \quad s > 0, \quad (21)$$

and multiplication by z performs two geometric operations on the complex plane. The modulus s uniformly scales magnitude, while the argument θ rotates orientation. In real matrix notation the same transformation is

$$A = sR(\theta), \quad (22)$$

where

$$R(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}. \quad (23)$$

The term *amplitwist* is due to Needham (1997), who uses it to name this decomposition into amplitude and twist. At this level, the terminology does not introduce a new mathematical object. It makes explicit the geometric content already present in multiplication by a nonzero complex number.

A useful physical intuition for the operation comes from a familiar two-finger gesture on a touchscreen map. Suppose a map is both too small and incorrectly oriented. A two-finger gesture can enlarge the map while simultaneously rotating it. Pulling the fingers apart supplies the scaling component, while turning them supplies the rotational component. The combined gesture is therefore an informal physical realization of

$$sR(\theta). \quad (24)$$

The analogy should not carry more weight than this. It illustrates the operator but does not establish anything about the universality of that operator.

The connection with the imaginary unit is similarly geometric. Multiplication by i corresponds to the matrix

$$J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = R\left(\frac{\pi}{2}\right). \quad (25)$$

Consequently,

$$J^2 = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} = -I. \quad (26)$$

Thus the familiar identity

$$i^2 = -1 \quad (27)$$

has an immediate geometric interpretation. One multiplication by i produces a quarter-turn. A second produces another quarter-turn. Their composition is a half-turn, which sends every vector to its negative. Beginning at 1, one rotation reaches i , and a second reaches -1 . That account is mathematically sound once any rhetoric about i being a “hack” or a non-number is removed.

The defining algebraic property of the amplitwist is closure under composition. Let

$$A_1 = s_1 R(\theta_1) \quad (28)$$

and

$$A_2 = s_2 R(\theta_2). \quad (29)$$

Then

$$A_1 A_2 = [s_1 R(\theta_1)] [s_2 R(\theta_2)] \quad (30)$$

$$= s_1 s_2 R(\theta_1) R(\theta_2). \quad (31)$$

It remains only to evaluate the product of the rotations. Direct multiplication gives

$$R(\theta_1)R(\theta_2) = \begin{pmatrix} \cos \theta_1 & -\sin \theta_1 \\ \sin \theta_1 & \cos \theta_1 \end{pmatrix} \begin{pmatrix} \cos \theta_2 & -\sin \theta_2 \\ \sin \theta_2 & \cos \theta_2 \end{pmatrix} \quad (32)$$

$$= \begin{pmatrix} \cos \theta_1 \cos \theta_2 - \sin \theta_1 \sin \theta_2 & -\cos \theta_1 \sin \theta_2 - \sin \theta_1 \cos \theta_2 \\ \sin \theta_1 \cos \theta_2 + \cos \theta_1 \sin \theta_2 & -\sin \theta_1 \sin \theta_2 + \cos \theta_1 \cos \theta_2 \end{pmatrix}. \quad (33)$$

Using the angle-addition identities,

$$\cos(\theta_1 + \theta_2) = \cos \theta_1 \cos \theta_2 - \sin \theta_1 \sin \theta_2 \quad (34)$$

and

$$\sin(\theta_1 + \theta_2) = \sin \theta_1 \cos \theta_2 + \cos \theta_1 \sin \theta_2, \quad (35)$$

this becomes

$$R(\theta_1)R(\theta_2) = R(\theta_1 + \theta_2). \quad (36)$$

Therefore

$$\boxed{A_1 A_2 = s_1 s_2 R(\theta_1 + \theta_2)}. \quad (37)$$

This is the basic composition theorem. Amplitudes multiply and twists add. The result is again an amplitwist, so the class is closed under composition.

The remaining group properties follow immediately. The identity is

$$I = 1 \cdot R(0), \quad (38)$$

and every amplitwist has the inverse

$$A^{-1} = \frac{1}{s} R(-\theta), \quad (39)$$

because

$$AA^{-1} = sR(\theta) \frac{1}{s} R(-\theta) \quad (40)$$

$$= R(\theta - \theta) \quad (41)$$

$$= R(0) \quad (42)$$

$$= I. \quad (43)$$

Associativity is inherited from matrix multiplication. The strict amplitwists therefore form a group under composition. In two dimensions this group is equivalent to the multiplicative group of nonzero complex numbers represented geometrically as uniform scal-

ings and rotations.

The word *uniform* is essential. An amplitwist does not permit arbitrary deformation. If a vector is transformed by

$$A = sR(\theta), \quad (44)$$

then every direction is scaled by the same factor s . Consequently, for any vectors u and v ,

$$\|Au\| = s\|u\|, \quad \|Av\| = s\|v\|. \quad (45)$$

Moreover,

$$(Au) \cdot (Av) = s^2(Ru) \cdot (Rv) \quad (46)$$

$$= s^2(u \cdot v), \quad (47)$$

because rotations preserve inner products. Hence

$$\frac{(Au) \cdot (Av)}{\|Au\|\|Av\|} = \frac{u \cdot v}{\|u\|\|v\|}. \quad (48)$$

Angles are preserved. The strict amplitwist is therefore conformal: it can change size and orientation, but it cannot stretch one direction more strongly than another.

This restriction provides an important diagnostic for later applications. Let $\mathcal{A}$ be a real symmetric 2×2 amplification matrix. It can be decomposed into an isotropic part and a traceless deviation,

$$\mathcal{A} = \frac{\text{tr } \mathcal{A}}{2} I + \mathcal{A}_0, \quad (49)$$

where

$$\boxed{\mathcal{A}_0 = \mathcal{A} - \frac{\text{tr } \mathcal{A}}{2} I}. \quad (50)$$

By construction,

$$\text{tr } \mathcal{A}_0 = 0. \quad (51)$$

The condition

$$\mathcal{A}_0 = 0 \quad (52)$$

is equivalent to

$$\mathcal{A} = \frac{\text{tr } \mathcal{A}}{2} I = sI. \quad (53)$$

Thus

$$\boxed{\|\mathcal{A}_0\| = 0 \iff \mathcal{A} = sI} \quad (54)$$

for any matrix norm. This gives a quantitative measure of departure from strict amplitwist structure. If $\|\mathcal{A}_0\| > 0$, the amplification is anisotropic.

For example,

$$\mathcal{A} = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} \quad (55)$$

has

$$\frac{\text{tr } \mathcal{A}}{2} I = \begin{pmatrix} 3/2 & 0 \\ 0 & 3/2 \end{pmatrix}, \quad (56)$$

and therefore

$$\mathcal{A}_0 = \begin{pmatrix} 1/2 & 0 \\ 0 & -1/2 \end{pmatrix} \neq 0. \quad (57)$$

This transformation doubles lengths along one axis while leaving lengths along the other unchanged. It is not uniform scaling and hence is not a strict amplitwist. Appending a rotation,

$$J = \mathcal{A}R(\theta), \quad (58)$$

does not repair the mismatch. The anisotropic component remains present.

Introducing $\mathcal{A}_0$ here is not merely preparation for a later criticism. It makes explicit what the amplitwist definition forbids. A useful mathematical concept must exclude possibilities as well as describe them. The strict amplitwist excludes anisotropic amplification. Consequently, when the language is extended to another domain, the statement that some transformation “scales and rotates” is insufficient by itself. The scaling must be isotropic if the term is being used in its strict sense.

The epistemic status of the present section is therefore *derived*. The scale-rotation decomposition, the composition law, the inverse, conformality, and the traceless-deviation criterion follow from elementary linear algebra. They do not depend on whether amplitwists later prove useful in quantum mechanics, neuroscience, linguistics, or artificial intelligence.

The elementary result supports a narrower and stronger conclusion than a continuous ascent toward a universal geometric ontology. The amplitwist is a well-defined transformation class with a simple composition law and a checkable exclusion criterion. Whether a physical or representational system belongs to that class must be established separately in each domain.

Electrical circuit theory provides the first case in which that identification can be made

without weakening the definition.

3. Circuits: The Strongest Physical Realization

Status: derived and directly checkable against standard circuit theory. Electrical circuits provide the strongest physical realization of the amplitwist language because the relevant decomposition into magnitude and phase is already native to alternating-current circuit theory. No new physical mechanism is required to obtain it. The amplitwist terminology supplies a geometric reading of an established mathematical representation.

For a sinusoidal signal at angular frequency ω , impedance can be written in polar form as

$$Z(j\omega) = |Z(\omega)|e^{j\theta(\omega)}. \quad (59)$$

If voltage and current phasors satisfy

$$V = ZI, \quad (60)$$

then multiplication by Z changes the magnitude of I by the factor $|Z|$ and changes its phase by θ . Under the standard identification of complex multiplication with planar rotation and scaling, this is precisely an operator of the form

$$A = sR(\theta). \quad (61)$$

The amplitude is $s = |Z|$, and the twist is the phase angle θ . In physical language, impedance has a magnitude that scales a sinusoidal response and a phase that shifts its timing relative to the driving signal.

Calling impedance an amplitwist therefore does not establish a new theorem about circuits. It identifies an existing piece of AC theory with the geometric structure derived in the preceding section. This is precisely why the example is strong. The identification does not require the amplitwist definition to be weakened in order to accommodate the application.

Resonance provides an equally direct example. For a series RLC circuit,

$$Z(\omega) = R + j \left(\omega L - \frac{1}{\omega C} \right). \quad (62)$$

The resonant frequency satisfies

$$\omega_0 L = \frac{1}{\omega_0 C}, \quad (63)$$

and therefore

$$\omega_0 = \frac{1}{\sqrt{LC}}. \quad (64)$$

At this frequency,

$$Z(\omega_0) = R. \quad (65)$$

The impedance is purely real, so

$$\arg Z(\omega_0) = 0. \quad (66)$$

In amplitwist notation, the transformation has reduced from

$$A = sR(\theta) \quad (67)$$

to

$$A = sR(0) = sI. \quad (68)$$

Resonance can therefore be described as a zero-twist condition: the simultaneous scale-and-rotate action becomes, at ω_0 , a pure scaling with no rotation. This is a faithful geometric reinterpretation of the phase condition at resonance.

It is tempting to describe resonance further as a point of “minimum geometric resistance to continuation,” with energy preferentially continuing because twist has vanished. That language should be separated from the result actually established. Standard circuit theory establishes cancellation of the reactive components and the corresponding zero phase angle for the series RLC impedance. It does not, from that fact alone, establish a general principle according to which vanishing geometric twist minimizes resistance to continuation. The zero-twist characterization survives. The broader continuation claim would require an additional argument.

A related amplitude-phase condition occurs in feedback oscillators. In its idealized form, the Barkhausen criterion (Sedra and Smith, 2015) is

$$|A\beta| = 1 \quad (69)$$

together with

$$\arg(A\beta) = 0 \pmod{2\pi}. \quad (70)$$

The first condition concerns amplitude, while the second concerns phase. The structural resemblance to an amplitwist fixed condition is therefore real. It is also structurally related to the onset of oscillation in dynamical systems, including systems undergoing Hopf bifurcation. The correspondence should not be promoted to an identity without the required dynamical hypotheses. The useful point is narrower: gain and rotation appear as independently constrained quantities in both descriptions.

The diode introduces a different operation. An ideal diode is not an invertible scale-and-rotate transformation. It distinguishes between two regions of its input space, admitting one and refusing the other: forward bias is evaluated and admitted, while reverse

bias is collapsed to a null output. Whatever terminology is attached to this distinction, the associated irreversibility can be proved directly.

Consider a rectifying map f for which two distinct input values are identified. In the symmetric formulation used by the rectification argument, suppose that for some $V \neq 0$,

$$f(V) = f(-V). \quad (71)$$

Assume, for contradiction, that f is invertible. Every invertible function is injective, so

$$f(x) = f(y) \implies x = y. \quad (72)$$

Applying injectivity to the rectification condition gives

$$V = -V. \quad (73)$$

Hence

$$2V = 0, \quad (74)$$

and therefore

$$V = 0, \quad (75)$$

contradicting the choice $V \neq 0$. Thus

$$\boxed{f(V) = f(-V), V \neq 0 \implies f \text{ is not invertible.}} \quad (76)$$

The same conclusion is even more immediate for the idealized blocking operation of an ideal diode. If every reverse-biased voltage is mapped to zero, then, for example,

$$f(-5) = f(-50) = f(-5000) = 0. \quad (77)$$

Distinct inputs have the same output, so the map is non-injective and therefore non-invertible. This is exactly why the original reverse voltage cannot be reconstructed from the blocked output.

The conclusion should be stated at the appropriate level. Rectification requires a loss of information whenever the rectifying map identifies inputs that were previously distinguishable. The irreversibility follows from elementary function theory.

A capacitor provides a complementary example in which information about a trajectory is compressed into a state variable. For an ideal capacitor,

$$i(t) = C \frac{dv_C}{dt}. \quad (78)$$

Integrating from an initial time t_i to t_0 gives

$$v_C(t_0) = v_C(t_i) + \frac{1}{C} \int_{t_i}^{t_0} i(\tau) d\tau. \quad (79)$$

The voltage at t_0 therefore depends on an integral over the preceding current history. It does not uniquely encode that history.

Let

$$\mathcal{F}[i] = \int_{t_i}^{t_0} i(\tau) d\tau. \quad (80)$$

The functional $\mathcal{F}$ is non-injective. There exist distinct histories i_1 and i_2 such that

$$i_1 \neq i_2 \quad (81)$$

while

$$\mathcal{F}[i_1] = \mathcal{F}[i_2]. \quad (82)$$

A slow, steady current over an extended interval and a large brief pulse followed by a long period of silence can be chosen to deposit the same total charge. Their temporal histories are radically different, yet their final capacitor voltage can be identical.

This motivates a more general statement:

A state variable can be a lossy compression of history.

The word “lossy” is literal here. Knowing $v_C(t_0)$ does not supply enough information to invert the integral and recover a unique $i(t)$. The state preserves what the model requires for subsequent evolution while discarding distinctions among multiple possible histories.

Digital logic makes this information loss countable. Consider a two-input Boolean gate with independent uniformly distributed inputs. The input pair (X_1, X_2) has entropy

$$H(X_1, X_2) = 2 \text{ bits}. \quad (83)$$

For an AND gate, the output is 1 only for the input $(1, 1)$. Thus

$$P(Y = 1) = \frac{1}{4}, \quad P(Y = 0) = \frac{3}{4}. \quad (84)$$

Its output entropy is

$$H(Y) = -\frac{1}{4} \log_2 \frac{1}{4} - \frac{3}{4} \log_2 \frac{3}{4} \quad (85)$$

$$\approx 0.811 \text{ bits}. \quad (86)$$

The information about the input pair that cannot be recovered from the output is therefore

$$H(X_1, X_2) - H(Y) \approx 2 - 0.811 = 1.189 \text{ bits.} \quad (87)$$

The same value holds for OR by symmetry.

For XOR,

$$P(Y = 0) = P(Y = 1) = \frac{1}{2}, \quad (88)$$

so

$$H(Y) = 1 \text{ bit}, \quad (89)$$

and therefore

$$H(X_1, X_2) - H(Y) = 1 \text{ bit.} \quad (90)$$

These figures, approximately 1.189 bits for AND and exactly 1 bit for XOR under uniform independent inputs, follow directly from the truth tables and Shannon entropy (Shannon, 1948).

The thermodynamic significance of logical irreversibility is supplied by Landauer's principle (Landauer, 1961). For the erasure of one bit under the standard ideal assumptions, the minimum heat dissipation is

$$Q_{\min} = k_B T \ln 2. \quad (91)$$

More generally, when a computation maps an input variable X to an output Y , the information discarded by the map can be expressed through conditional entropy, giving an idealized lower bound of the form

$$Q_{\min} \geq k_B T \ln 2 H(X | Y), \quad (92)$$

when the entropy is measured in bits and the appropriate thermodynamic assumptions hold.

This result requires a correction to a common piece of popular rhetoric. The warmth of a laptop cannot simply be identified with Landauer heat from logically erased bits. Real processors dissipate energy through many mechanisms and normally operate far above the Landauer limit. The rigorous claim is that logically irreversible information loss has a nonzero thermodynamic floor, not that the macroscopic temperature of a particular computer directly measures the number of erased logical distinctions.

Johnson–Nyquist noise (Johnson, 1928; Nyquist, 1928) supplies another exact thermodynamic relation. For a resistor R at temperature T , the mean-square thermal-noise voltage over bandwidth Δf is

$$\overline{v_n^2} = 4k_B T R \Delta f. \quad (93)$$

The same resistance that participates in dissipation therefore also sets the scale of equilibrium electrical fluctuations. This is an instance of the broader fluctuation-dissipation relation. It is legitimate to place thermal noise and dissipation in the same physical discussion because they are connected by established statistical physics.

It would be too strong, however, to infer that Johnson-Nyquist noise is itself the direct exhaust of a particular preceding logical erasure. The established relation is subtler. Logical irreversibility supplies a lower bound on thermodynamic cost, while Johnson-Nyquist noise describes equilibrium fluctuations associated with a dissipative electrical element. Their common appearance in thermodynamic information processing is substantive, but it does not make them event-by-event equivalents.

The epistemic status of this section is therefore strong but differentiated. Impedance as magnitude plus phase, resonance as cancellation of reactance and vanishing impedance phase, the non-invertibility of many-to-one rectification, the non-injectivity of capacitor-history compression, the information loss of Boolean gates, Landauer's bound, and Johnson-Nyquist noise are all directly derivable or established results. The amplitwist language provides a coherent geometric interpretation of part of this structure without changing the underlying circuit theory.

This is also the point at which the argument must stop accumulating authority by sequence. The fact that impedance exactly realizes scale plus phase does not imply that a cortical column does. The fact that a diode is non-injective does not imply that semantic change is governed by the same operator family. The fact that information erasure has a thermodynamic cost does not establish that cognitive uncertainty is Johnson-Nyquist noise in another substrate.

The mathematical and circuit results survive because each can be checked independently. The later claims must meet that same standard on their own terms.

4. The Boundary

Status: methodological statement, not a claims section. The argument reaches an epistemic boundary at this point. The preceding sections concern claims that can be established without accepting the broader continuation-geometry program. The relation $\approx$ is defined by a limit and its elementary calculus follows from that definition. The strict amplitwist $A = sR(\theta)$ has an explicit composition law and a checkable conformality condition. In circuit theory, complex impedance already decomposes into magnitude and phase, resonance supplies a zero-phase case, rectification supplies a direct non-injectivity result, and the information-theoretic and thermodynamic claims can be evaluated independently.

Nothing in those derivations licenses automatic transfer of the same status to another domain.

This boundary is necessary because the broader program is easily read as a cascade. Once complex multiplication has been interpreted geometrically, the same language can be carried into quantum mechanics. Once impedance has been recognized as magnitude plus

phase, the circuit result can be carried into cortex. Cortical scaling and rotation can then be used to motivate a theory of conscious coherence, which in turn can be expanded into linguistic change, cultural evolution, and artificial-intelligence alignment. A natural closing move is to present this succession as evidence that the same geometric engine operates from quantum systems through circuits and neural columns to language and culture.

That succession is a useful map of the research program. It is not itself an argument for the equivalence of the domains.

From this section onward, each application must therefore establish its own relation to the strict amplitwist. The relevant question is not merely whether the system can be described using words such as *scale*, *rotation*, *phase*, *flow*, or *twist*. Such descriptions are available for an enormous class of dynamical systems. The stronger question is whether the amplitwist identification imposes a restriction that a more generic transformation would not impose.

For a two-dimensional linear transformation, the distinction can already be stated precisely. If the proposed operator has the form

$$A = sR(\theta), \quad (94)$$

then its scaling is isotropic and its traceless amplification component vanishes:

$$\mathcal{A}_0 = \mathcal{A} - \frac{\text{tr } \mathcal{A}}{2} I = 0. \quad (95)$$

If the application instead requires a general transformation

$$J = \mathcal{A}R \quad (96)$$

with

$$\mathcal{A}_0 \neq 0, \quad (97)$$

then the strict amplitwist classification has failed unless a further argument shows why the anisotropic component may be removed, neglected, or quotiented without changing the phenomenon being modeled. Calling both objects “scale and twist” does not eliminate the difference.

This supplies a general test for the sections that follow:

Does modeling the phenomenon as an amplitwist forbid anything that a more generic transformation of the same type would permit?

If the answer is yes and the restriction is independently supported, the amplitwist language has acquired explanatory content in the new domain. If the answer is unknown because the necessary derivation or empirical comparison has not been performed, the claim is an *unverified extension*. If the defining restriction has been examined and does not follow from the proposed model, the appropriate status is a *type-mismatch result*. A prediction may instead be *falsifiable but untested* when the relevant observable and possible

failure condition can be specified before measurement. A proposed formalism can remain a *programmable extension* when it defines a research direction without yet establishing the distinctive constraint required for validation.

These categories are deliberately asymmetric. “Unverified” does not mean false. A “type-mismatch result” does not mean that every weaker claim in the same domain has been refuted. “Falsifiable but untested” does not mean evidentially worthless. Nor does a successful derivation in one domain raise the prior status of a later application merely because the same terminology appears in both.

The quantum-mechanical extension illustrates the first of these cases. It contains explicit mathematics, including a real matrix representation of the complex quarter-turn and a proposed quotient construction for composite systems. The question at issue is therefore not whether the proposal contains mathematics. It plainly does. The question is whether the stronger reconstruction claims have received the independent derivation required to place them alongside the results of the preceding sections.

They have not yet received that status here. The quantum construction should therefore be presented as a mathematical program, not as credibility inherited from the amplitwist algebra that precedes it.

5. Quantum Mechanics: An Unverified Program

Status: unverified extension. The quantum-mechanical extension (Flyxion, 2026d) begins from a narrower observation than is sometimes assigned to it. Standard quantum mechanics makes essential use of complex vector spaces, and multiplication by the imaginary unit can be represented geometrically as a quarter-turn on a real two-dimensional auxiliary space. Making that rotation explicit can expose geometric structure that complex notation compresses. The stronger proposal is that this real representation, together with an admissibility quotient for composite systems, provides an explanatory reconstruction of quantum mechanics in which physically irrelevant representational distinctions are separated from invariant structure.

The distinction between these two claims is important. The first is elementary linear algebra. The second is a reconstruction program whose general claims have not undergone the same level of independent derivation applied elsewhere in this essay.

The motivating criticism is not that conventional complex quantum mechanics is incorrect. It is that a formally successful representation can become epistemically opaque when the operations encoded by its notation cease to be geometrically visible. Complex notation produces correct predictions without necessarily making its proposed generative mechanism explicit. The real flag-space construction is offered as a candidate for making that mechanism explicit.

The elementary construction begins with the matrix

$$J_F = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}. \quad (98)$$

It satisfies

$$J_F^2 = -I. \quad (99)$$

Consequently, the algebra generated by I and J_F reproduces ordinary complex multiplication. For real a, b, c, d ,

$$(aI + bJ_F)(cI + dJ_F) = acI + adJ_F + bcJ_F + bdJ_F^2 \quad (100)$$

$$= (ac - bd)I + (ad + bc)J_F. \quad (101)$$

This is exactly the multiplication rule

$$(a + ib)(c + id) = (ac - bd) + i(ad + bc). \quad (102)$$

At this level there is no speculative step. The complex quarter-turn has simply been represented by an explicit real matrix.

An expanded flag-space representation makes this replacement intuitive. A conventional complex amplitude stores real and imaginary components within a single complex coordinate. In the expanded representation, an auxiliary two-dimensional flag records the rotational degree of freedom explicitly. One may picture a flagpole with an attached flag: the pole supplies the underlying position while the flag orientation makes the phase degree of freedom geometrically visible. Multiplication by i is then represented not by an opaque symbolic operation but by a quarter-turn of the flag.

Schematically, a complex state may be mapped into a real flagged representation by

$$S(|\psi\rangle) = \text{Re } |\psi\rangle \otimes |0\rangle_F + \text{Im } |\psi\rangle \otimes |1\rangle_F. \quad (103)$$

A complex phase transformation then appears as a real rotation of the flag degree of freedom. The proposal is not that the predictions of complex quantum mechanics are discarded. Rather, the complex representation is expanded into a real one in which the rotational action formerly encoded by i becomes explicit.

This is the sense in which the construction can be described as “uncompressing” complex geometry. The terminology should be understood representationally. Replacing a complex coordinate by additional real coordinates does not demonstrate that complex numbers were physically unreal, nor does it establish that quantum phenomena are caused by literal spatial rotation. It demonstrates that the relevant complex algebra admits a real representation in which the quarter-turn structure can be displayed explicitly.

Composite systems expose the central difficulty with this expansion.

Suppose a complex composite state has four complex coordinates. Its ordinary repre-

sensation then contains eight real parameters. A naive local flag construction can instead produce a sixteen-dimensional real representation. The problem is not merely that the real description has become inefficient. The local construction can distinguish vectors that correspond to the same global complex state.

Consider the state

$$i|00\rangle. \quad (104)$$

The phase i may be attached to the first local factor,

$$(i|0\rangle) \otimes |0\rangle, \quad (105)$$

or to the second,

$$|0\rangle \otimes (i|0\rangle). \quad (106)$$

These are the same vector in the ordinary complex tensor product:

$$(i|0\rangle) \otimes |0\rangle = |0\rangle \otimes (i|0\rangle) = i|00\rangle. \quad (107)$$

A naive procedure that realifies each local factor independently, however, can map these decompositions to different vectors in the enlarged real tensor-product space. The local construction therefore remembers a distinction that the original global state did not possess. The proposed construction identifies the resulting space as sixteen-dimensional, compared with the eight-dimensional real space obtained by applying the global map directly. The difficulty is consequently both dimensional and representational: the naive construction depends on an arbitrary choice of which factor “carries” the phase.

The proposed solution is an admissibility quotient.

Let X denote the raw local flag space, and let $\mathcal{M}$ denote the class of measurements regarded as physically admissible. Define an equivalence relation by

$$x \sim_A y \iff M(x) = M(y) \text{ for every } M \in \mathcal{M}, \quad (108)$$

where equality here means equality of the measurement statistics generated by the two representatives.

Reflexivity and symmetry are immediate. Transitivity follows because if

$$x \sim_A y \quad (109)$$

and

$$y \sim_A z, \quad (110)$$

then every admissible measurement agrees on x and y , and on y and z , hence also on x and z . Thus $\sim_A$ is an equivalence relation.

The physical representation is then taken not to be the raw space X , but the quotient

$$X/\sim_A. \quad (111)$$

Representations that remain distinct as vectors but cannot be distinguished by any admissible measurement occupy the same equivalence class.

This is the mathematical content behind the phrase *witness invariance*. If two states look different in the flag-space equations but yield identical results to every physical detector, the representational distinction should not be promoted into a physical distinction.

The quotient should not be described as simply deleting the extra degrees of freedom. Quotienting reorganizes them. If a group G acts on a representation space X , then the quotient map

$$\pi : X \longrightarrow X/G \quad (112)$$

identifies every orbit

$$Gx = \{gx : g \in G\} \quad (113)$$

with a single quotient point $[x]$. The representatives remain present in the fiber

$$\pi^{-1}([x]) = Gx. \quad (114)$$

The ambiguity has therefore moved from the choice of physical state to the choice of representative of that state. Quotienting reorganizes the ambiguity rather than annihilating it.

This gives the $16 \rightarrow 8$ discussion a more precise interpretation than a purely visual account of a map with sixteen roads and eight destinations would suggest. The raw local representation contains distinctions generated by its representational construction. The admissibility quotient proposes to identify those distinctions when they have no effect on the specified measurement statistics.

That proposal is mathematically meaningful. Its status here nevertheless remains *unverified*.

The reason is not that the construction lacks formal content. The quarter-turn matrix is explicit, the local-factor ambiguity is explicit, and the admissibility equivalence can be defined precisely. The unresolved issue is the scope of the reconstruction. Establishing that a particular redundancy can be removed while preserving a specified measurement class is not by itself a derivation that the resulting quotient recovers the complete structure and predictions of arbitrary composite quantum systems.

The stronger claims therefore require independent verification. One must specify the admissible measurement algebra, establish that the quotient preserves the complete set of physical predictions claimed for it, demonstrate that no physically distinguishable states are inadvertently identified, and show that the construction behaves correctly under general multipartite composition.

It would also exceed what has been earned to conclude, from the quotient construction alone, that quantum weirdness is simply a consequence of forgetting that algebra is compressed geometry, or that the universe is thereby shown to be executing precise geometric rotations. A real representation of complex quantum mechanics can make rotational structure explicit without proving that rotation is ontologically prior, that every specifically quantum phenomenon reduces to an amplitwist, or that the quotient has identified the unique underlying ontology.

The appropriate conclusion is narrower.

Complex multiplication admits an exact real matrix representation in which multiplication by i appears as a quarter-turn. Extending that representation locally to composite systems introduces redundancy and representation dependence. An admissibility quotient provides a principled proposal for identifying representations that no allowed measurement distinguishes. These facts establish a concrete mathematical program.

They do not yet establish the program's strongest reconstruction claims.

The status of the quantum extension is therefore *unverified extension*: presented rather than endorsed or rejected. Unlike the material before the boundary, its central claim is not being imported into this essay as a settled result. Unlike the cortical application considered next, it has also not received a negative verdict. Its outstanding obligations are derivational.

The distinction matters. An unverified mathematical program should not be weakened merely because it remains unverified, but neither should the exactness of its elementary quarter-turn construction be allowed to confer certainty on the larger reconstruction by association. The appropriate next step is independent derivation of the quotient's claimed physical equivalence for the general systems to which it is meant to apply.

6. Cortical Columns: A Type Mismatch

Status: type-mismatch result. The cortical extension (Flyxion, 2026b) makes a substantially stronger identification than the quantum section. It does not merely propose that amplitwist geometry might provide a useful representation of neural dynamics. The claim is direct: cortical columns are local amplitwist operators. On this account, a cortical column acts on a neural representation in the same structural manner that a complex impedance acts on an electrical signal, scaling some component of the representation while rotating it through a representational space.

The associated interpretation divides the cognitive description among three fields. A scalar field Φ represents semantic salience or intensity. A vector field $\mathbf{v}$ represents the directional flow of attention or conceptual motion. An entropy field S represents uncertainty or dispersion. Increased semantic salience maps onto amplitude and changes in attentional direction map onto twist, while S modulates the stability of the resulting dynamics.

An illustrative example considers perception of an apple. Recognition increases the salience associated with the concept, while hunger changes the direction of attention from perceptual recognition toward eating. In an informal presentation, this transition might

be described as the cortical representation being physically rotated from one conceptual configuration toward another.

The example is suggestive, but it does not decide the mathematical classification. Neural gain and representational rotation are not sufficient to establish that the underlying operator is a strict amplitwist. The relevant question is the operator actually supplied by the cortical formulation rather than the verbal analogy.

The relevant local Jacobian is written

$$J = \nabla \mathbf{v} = \mathcal{A}R, \quad (115)$$

where R is rotational and $\mathcal{A}$ is a symmetric amplification matrix.

This differs from the strict amplitwist derived in Section 2,

$$A = sR(\theta), \quad (116)$$

where s is a scalar.

The difference is not cosmetic. A scalar multiple of the identity performs isotropic amplification. A general symmetric matrix need not.

The criterion introduced earlier makes the distinction explicit. Decompose $\mathcal{A}$ as

$$\mathcal{A} = \frac{\text{tr } \mathcal{A}}{2} I + \mathcal{A}_0, \quad (117)$$

where

$$\mathcal{A}_0 = \mathcal{A} - \frac{\text{tr } \mathcal{A}}{2} I. \quad (118)$$

The traceless component $\mathcal{A}_0$ measures the departure from isotropic amplification. If

$$\mathcal{A}_0 = 0, \quad (119)$$

then

$$\mathcal{A} = \frac{\text{tr } \mathcal{A}}{2} I = sI \quad (120)$$

for some scalar s . Consequently,

$$J = \mathcal{A}R = sIR = sR, \quad (121)$$

and the cortical operator is a strict amplitwist.

Conversely, if

$$\mathcal{A}_0 \neq 0, \quad (122)$$

then $\mathcal{A}$ contains an anisotropic component. The transformation may amplify different directions by different factors. It therefore belongs to the broader class of a rotation com-

posed with symmetric deformation, not necessarily to the strict conformal class $sR(\theta)$.

The resulting test is exact:

$$\boxed{J = \mathcal{A}R \text{ is a strict amplitwist} \iff \mathcal{A}_0 = 0} \quad (123)$$

under the stated symmetric-amplification decomposition. Equivalently,

$$\boxed{\|\mathcal{A}_0\| = 0} \quad (124)$$

for any matrix norm.

A simple counterexample demonstrates why the distinction matters. Let

$$\mathcal{A} = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}. \quad (125)$$

Then

$$\frac{\text{tr } \mathcal{A}}{2} = \frac{3}{2}, \quad (126)$$

so

$$\mathcal{A}_0 = \begin{pmatrix} 1/2 & 0 \\ 0 & -1/2 \end{pmatrix}. \quad (127)$$

Hence

$$\mathcal{A}_0 \neq 0. \quad (128)$$

The transformation doubles one principal direction while leaving the other unchanged. No scalar s can represent both amplification factors simultaneously. Appending a rotation does not remove this anisotropy:

$$J = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix} R(\theta) \quad (129)$$

remains outside the strict amplitwist class.

The same point can be expressed through singular values. For

$$A = sR(\theta), \quad (130)$$

both singular values equal s :

$$\sigma_1(A) = \sigma_2(A) = s. \quad (131)$$

A general symmetric amplification matrix permits

$$\sigma_1(\mathcal{A}) \neq \sigma_2(\mathcal{A}). \quad (132)$$

Unequal singular values provide another direct signature of anisotropic deformation. The amplitwist restriction is therefore empirically and mathematically substantive. It excludes transformations that a generic symmetric amplification followed by rotation permits.

This is the basis of the type-mismatch result established above.

The result is not that cortical columns have been shown *not* to behave as amplitwists. Such a conclusion would require evidence that the relevant cortical transformations have

$$\|\mathcal{A}_0\| > 0 \quad (133)$$

in the regime under consideration. The present result is instead that the mathematical formulation offered in support of the cortical claim does not establish the claim. Writing

$$J = \mathcal{A}R \quad (134)$$

with general symmetric $\mathcal{A}$ permits a strictly larger transformation class than

$$J = sR. \quad (135)$$

The strict amplitwist is recovered only as the isotropic special case.

This makes the problem a *type mismatch* rather than a generic shortage of evidence. Membership is asserted in one transformation class while the supporting formalism specifies membership only in a broader class. Shared vocabulary about amplification and rotation can obscure that distinction. After decomposition by $\mathcal{A}_0$, the missing condition becomes checkable.

The result also illustrates why the boundary in the preceding section is necessary. In circuit theory, complex impedance already has a scalar magnitude and a phase:

$$Z = |Z|e^{j\theta}. \quad (136)$$

Its amplitude is therefore intrinsically isotropic in the two-dimensional phasor representation. No corresponding result follows merely from saying that neural dynamics contain gain and rotation. The circuit case cannot lend its scalar-amplitude property to the cortical case by analogy.

The three-field interpretation remains meaningful at a weaker level. A scalar field can represent changing salience, a vector field can represent directional dynamics, and an entropy field can parameterize uncertainty or dispersion. Neural population dynamics can likewise contain rotations, contractions, expansions, and more general deformations. None of those possibilities is excluded by the present result. What fails is the stronger inference

$$\text{amplification} + \text{rotation} \implies \text{strict amplitwist}. \quad (137)$$

The correct implication is

$$\text{isotropic amplification} + \text{rotation} \implies \text{strict amplitwist.} \quad (138)$$

The outstanding proof obligation is therefore unusually compact:

$$\boxed{\mathcal{A}_0 = 0.} \quad (139)$$

There are several ways such an obligation could eventually be met. A theoretical model could derive isotropy from the governing neural dynamics. An empirical reconstruction of the local Jacobian could find approximately equal singular values within a specified tolerance. A weaker effective theory could show that anisotropic components average or renormalize away at the scale at which the amplitwist description is intended to operate. Any such route would require its own argument. None follows merely from the present notation.

The correspondence between impedance and cortical columns, if drawn too quickly, risks moving from the observation that impedance scales and phase-shifts electrical signals to the stronger statement that cortical columns perform the exact same geometric operations in the brain, with semantic salience as amplitude, attentional motion as twist, and cognitive uncertainty as an analogue of thermal noise. The first two correspondences can motivate a model. They do not establish equality of operator class. The comparison between cognitive uncertainty and Johnson-Nyquist noise is weaker still unless a physical relation connecting the corresponding stochastic processes is supplied.

The appropriate status is consequently a *type-mismatch result*. This is stronger than “unverified” because a specific defining condition has been examined and found not to follow from the formulation. It is narrower than empirical falsification because it has not been established that real cortical dynamics violate that condition.

This distinction is consequential for the next extension. A proposed theory of consciousness proceeds from cortical amplitwists to a claim that local amplitwist operators synchronize into a coherent field configuration. If the amplitwist classification of the local operators remains unestablished, that later theory cannot use the classification as an established premise.

The consciousness proposal must therefore be evaluated downstream of this unmet condition rather than downstream of the derived amplitwist algebra of Sections 2 and 3.

7. Consciousness: A Conjecture Downstream of an Unmet Precondition

Status: programmatic extension downstream of a type-mismatch result. The cortical amplitwist proposal leads directly to the most ambitious claim considered in this essay: consciousness is identified with the formation of a coherent region in the interacting fields Φ , $\mathbf{v}$, and S . Such a region may be called a *coherence tile*, with its formation described as

a critical phase transition. As uncertainty falls, the scalar field of semantic salience and the vector field of attentional motion are said to become sufficiently aligned for local amplitwist operators to synchronize. The resulting stable configuration is then identified with a conscious state.

The claim can be represented schematically as

$$(\Phi, \mathbf{v}, S) \longrightarrow \text{local operator alignment} \longrightarrow \text{coherence tile} \longrightarrow \text{conscious state.} \quad (140)$$

Here Φ denotes the scalar field associated with semantic salience, $\mathbf{v}$ denotes the vector field associated with attentional or conceptual motion, and S denotes the entropy field associated with uncertainty or dispersion. The proposed transition occurs when these fields enter a sufficiently coherent regime.

This is a substantive hypothesis. It is also much stronger than the mathematical results established for the amplitwist itself.

The distinction matters because the cortical and consciousness claims are easily treated as a continuous derivation. Once cortical columns have been identified as local amplitwists, synchronization among those amplitwists can be presented as the mechanism by which a conscious state appears: when the entropy gradient becomes sufficiently small, semantic intensity and attentional flow are said to align, after which the amplitwist operators in the region synchronize into a stable conscious configuration.

The first step in this chain has not been earned in the strict mathematical sense (Section 6). The cortical formulation

$$J = \nabla \mathbf{v} = \mathcal{A}R \quad (141)$$

does not imply

$$J = sR \quad (142)$$

unless

$$\mathcal{A}_0 = \mathcal{A} - \frac{\text{tr } \mathcal{A}}{2} I = 0. \quad (143)$$

Consequently, a theory in which conscious coherence is specifically the synchronization of cortical *amplitwists* inherits an unmet precondition:

$$\boxed{\|\mathcal{A}_0\| = 0} \quad (144)$$

for the local transformations to which the amplitwist classification is being applied.

This does not refute the consciousness hypothesis.

The logical distinction is important. Suppose the proposed argument has the form

$$P \implies Q, \quad (145)$$

where P is the claim that the relevant cortical operators are strict amplitwists and Q is the claim that their collective coherence participates in the proposed conscious state. Failure to establish P does not establish

$$\neg Q. \quad (146)$$

It establishes only that Q cannot presently inherit support through that route.

This changes the status of the consciousness claim without supplying a countertheory of consciousness. A coherent transition in cortical dynamics might exist even if the local transformations are anisotropic. A phase transition involving Φ , $\mathbf{v}$, and S might be definable without strict amplitwists. Neural population dynamics might exhibit collective synchronization without satisfying the conformal condition of Section 2. Those possibilities remain open. What cannot presently be claimed is that the strict amplitwist algebra has already supplied the local operator from which the consciousness result follows.

There is a second and independent gap. Even if the cortical obligation were discharged and

$$\mathcal{A}_0 = 0 \quad (147)$$

were established for the relevant local dynamics, the transition from amplitwist coherence to consciousness would still require its own derivation.

The phrase *critical phase transition* has mathematical content only after an order parameter, control parameter, and transition criterion have been specified. A schematic version might involve a coherence quantity

$$C = C(\Phi, \mathbf{v}, S) \quad (148)$$

and a critical value C_c such that a qualitative change occurs when

$$C \geq C_c. \quad (149)$$

Alternatively, if entropy or an entropy gradient supplies the control parameter, one would require a defined quantity such as

$$\Gamma_S = \|\nabla S\| \quad (150)$$

and a demonstrated transition at some critical regime

$$\Gamma_S \leq \Gamma_c. \quad (151)$$

These expressions illustrate the form of the missing obligation. They are not results already established in the surrounding program. The claim that a critical transition occurs when the relevant fields align has been asserted, but a derived order parameter, a critical exponent, a threshold theorem, or an independently established correspondence between such a transition and conscious experience has not been supplied.

The word *coherence* requires the same discipline. It could denote phase synchronization, alignment of vector fields, low variance among local operators, increased mutual information, topological consistency, or several other mathematically inequivalent conditions. A theory cannot obtain predictive force merely from the intuitive desirability of coherence. The particular form of coherence must constrain the allowed states.

The strongest form of the underlying claim goes further still: that thoughts are literal geometric flows over a representational manifold, rather than a metaphor for such flows. Such a statement can be read in at least two ways.

In the weaker sense, neural population states can be represented as trajectories on a manifold. This is a standard and empirically tractable mathematical strategy. If neural activity is embedded in a lower-dimensional state space,

$$x(t) \in \mathcal{M}, \quad (152)$$

then cognition can be studied through the trajectory

$$t \mapsto x(t). \quad (153)$$

Calling that trajectory geometric does not require any controversial ontology.

The stronger reading asserts that the representational geometry is itself the physical mechanism of thought, rather than a useful coordinate description of neural dynamics. That conclusion requires considerably more. A successful manifold representation does not by itself establish that the manifold is ontologically fundamental. Different coordinate systems, embeddings, or latent-variable models may represent the same physical dynamics. The distinction between representation and invariant structure that motivated the quantum section therefore applies here as well.

The consciousness proposal is consequently downstream of two separate proof obligations:

$$\boxed{\text{cortical amplitwist identification}} \quad (154)$$

and

$$\boxed{\text{coherence transition} \implies \text{conscious state.}} \quad (155)$$

The first currently fails to follow from the stated cortical operator because a general symmetric $\mathcal{A}$ permits anisotropic amplification. The second has not been derived from the amplitwist algebra even conditionally.

These obligations should not be collapsed into a generic statement that consciousness is “speculative.” Doing so would obscure where the argument actually stands. The amplitwist algebra itself is not speculative. The cortical identification has received a specific negative result. The coherence-tile proposal is a further conjectural construction that depends partly on that unresolved identification and partly on an additional transition claim of its own.

This local status assignment also prevents a converse error. The failure of the strict

cortical classification does not make every empirical consequence associated with the larger proposal worthless. A theory can contain an underived interpretation while still generating a discriminating prediction.

One such prediction concerns task switching: that it should produce a characteristic rotational trajectory in neural population space. Unlike the general assertion that consciousness is a coherence tile, this claim can be translated into an observable, compared against alternative transformation classes, and exposed to empirical failure.

That prediction deserves to be considered separately from the consciousness interpretation precisely because its epistemic status is better defined.

8. A Genuine Prediction: Population-Vector Rotation

Status: falsifiable but untested. The strongest empirical consequence of the cortical amplitwist proposal is narrower than the consciousness interpretation and, for that reason, more scientifically tractable. Rotational structure in neural population dynamics has been documented in a related but distinct context, motor cortex activity during reaching (Churchland et al., 2012); the claim considered here, that prefrontal population activity should exhibit comparable rotation-like trajectories specifically during task switching, is not drawn from that literature and should be understood as an independent proposal awaiting its own test. Neural population activity in prefrontal cortex is predicted to exhibit rotation-like trajectories during task switching. Rather than treating the geometric language only as an interpretation imposed after observing neural dynamics, this claim identifies a pattern that should be present in data if the proposed operator structure is correct.

The observable is a time-resolved neural population state. Let

$$x(t) \in \mathbb{R}^n \quad (156)$$

denote the activity of a recorded neural population at time t , or a latent representation obtained from that population by a specified dimensionality-reduction procedure. During a task switch, the resulting trajectory is

$$t \mapsto x(t). \quad (157)$$

The prediction is not merely that $x(t)$ changes. Any successful account of task switching must allow neural activity to change. Nor is the prediction satisfied merely because a two-dimensional projection of the trajectory happens to look curved. The amplitwist hypothesis becomes distinctive only if the transformation relating relevant population states is well approximated by the restricted operator class

$$J_t = s_t R(\theta_t), \quad (158)$$

or by the corresponding higher-dimensional generalization if the hypothesis is explicitly extended beyond the two-dimensional case.

For the two-dimensional formulation considered here, suppose the locally fitted transformation is written in polar form as

$$J_t = \mathcal{A}_t R_t, \quad (159)$$

with $\mathcal{A}_t$ symmetric. The strict amplitwist prediction then requires

$$\mathcal{A}_t \approx s_t I. \quad (160)$$

Using the traceless-deviation criterion,

$$\mathcal{A}_{0,t} = \mathcal{A}_t - \frac{\text{tr } \mathcal{A}_t}{2} I, \quad (161)$$

the expected result is

$$\boxed{\|\mathcal{A}_{0,t}\| \approx 0} \quad (162)$$

over the portion of the trajectory to which the amplitwist description is claimed to apply.

This condition separates the prediction from the weaker observation that neural trajectories can rotate. A general linear dynamical system may exhibit rotational components while also undergoing shear or anisotropic expansion and contraction. Such a system can be written schematically as

$$J_t = \mathcal{A}_t R_t \quad (163)$$

without satisfying

$$\mathcal{A}_t = s_t I. \quad (164)$$

The existence of a rotational component would therefore support only the weaker claim that task switching involves representational rotation. Evidence for a *strict amplitwist* requires the additional isotropy constraint.

The same distinction can be formulated through singular values. If

$$J_t = s_t R_t, \quad (165)$$

then

$$\sigma_1(J_t) = \sigma_2(J_t) = s_t. \quad (166)$$

A reproducible difference

$$\sigma_1(J_t) \neq \sigma_2(J_t) \quad (167)$$

would indicate anisotropic deformation. The ratio

$$\rho_t = \frac{\sigma_{\max}(J_t)}{\sigma_{\min}(J_t)} \quad (168)$$

therefore supplies another possible diagnostic. A strict amplitwist predicts

$$\rho_t = 1, \quad (169)$$

with empirical data permitting some prespecified tolerance around unity.

The observable and expected result can thus be stated with reasonable precision:

During task switching, prefrontal neural population states should follow a reproducible rotation-like trajectory whose locally fitted transformation is approximately conformal, with negligible anisotropic amplification relative to an explicitly specified measurement tolerance.

This formulation is stronger than a purely visual description because it specifies what distinguishes amplitwist dynamics from generic rotational dynamics.

It also makes the hypothesis vulnerable to several kinds of negative evidence.

Most directly, the hypothesis would be disfavored if task-switch trajectories reproducibly require

$$\|\mathcal{A}_{0,t}\| \not\approx 0. \quad (170)$$

Likewise, systematic inequality of the relevant singular values would count against the strict model:

$$\sigma_1(J_t) \not\approx \sigma_2(J_t). \quad (171)$$

A second negative result would occur if a rotational trajectory appeared only under a specially chosen dimensionality reduction and disappeared under other reasonable representations of the same neural data. In that case, the apparent rotation could be an artifact of coordinates rather than an invariant feature of the underlying dynamics.

A third test is comparative. Let $\mathcal{M}_A$ denote the restricted amplitwist model and $\mathcal{M}_G$ a more general linear model. The latter permits anisotropic deformation and therefore contains a larger class of transformations. Because

$$\mathcal{M}_A \subset \mathcal{M}_G, \quad (172)$$

the unrestricted model can generally fit training data at least as well as the restricted one. The meaningful question is whether the amplitwist restriction retains predictive adequacy on held-out data while using fewer degrees of freedom. If the general model produces a substantial and reproducible improvement in out-of-sample prediction because anisotropic components are required, then the amplitwist restriction has failed to capture the relevant dynamics.

This comparison expresses the methodological criterion introduced at the boundary. The amplitwist hypothesis earns empirical content precisely because it forbids some transformations that the generic model allows.

A complete experimental protocol is not supplied here. Identifying prefrontal population vectors during task switching and predicting rotation-like trajectories does not by itself settle the recording modality, sampling requirements, latent-space construction, dimensionality, statistical threshold, admissible tolerance for $\|\mathcal{A}_0\|$, model-selection criterion, or required effect size. Possible observational routes through neural population codes and other neurophysiological measurements do not amount to a preregistered experiment.

Those missing specifications should remain missing here rather than being supplied retrospectively. In particular, no numerical threshold of the form

$$\|\mathcal{A}_{0,t}\| < \varepsilon_A \quad (173)$$

can be justified without an independently derived or empirically calibrated value of ε_A . Similarly, no claim is made here about how much improvement by a general deformation model would be sufficient to reject the restricted amplitwist description.

The prediction is nevertheless genuinely falsifiable at the structural level. Its three essential components can now be separated.

The observable is the trajectory of a neural population representation during task switching.

The expected result is a reproducible rotational trajectory whose local transformation is approximately a uniform scaling composed with rotation.

Evidence against the hypothesis would include robust anisotropic deformation, absence of the predicted rotational structure, dependence of the apparent rotation on arbitrary representational choices, or consistent predictive superiority of a more general transformation class specifically because the amplitwist restriction is violated.

This status is epistemically stronger than the surrounding consciousness claim in one important respect. The statement that consciousness emerges as a coherence tile can presently accommodate a wide range of possible neural dynamics because its transition criterion remains incompletely specified. The population-vector prediction instead exposes a particular geometric expectation to observation.

The distinction also prevents a possible overcorrection following the cortical result. Section 6 shows that the existing formulation

$$J = \mathcal{A}R \quad (174)$$

does not establish that cortical dynamics are strict amplitwists. That negative result does not make the amplitwist hypothesis empirically empty. It identifies exactly what must now be measured. The previously missing condition

$$\mathcal{A}_0 = 0 \quad (175)$$

becomes an empirical target rather than an implicit assumption.

The status of this section is therefore *falsifiable but untested*. Moving beyond retrospective interpretation to make a claim about what neural data should contain is a genuine strength, sharpened here by distinguishing generic rotation from strict amplitwist structure and by specifying what kind of observation would count against the latter.

Whether the prediction survives such a test remains an empirical question. Its value at present lies precisely in the fact that the answer is not guaranteed.

9. Culture and Alignment: Programmatic Extensions

Status: programmatic extensions. The final domain expansion considered here (Flyxion, 2026a) moves from neural dynamics to cultural change. The proposed bridge is recursive transformation. A representation is altered at one level, inherited by another, and altered again. This process is organized into three layers, denoted R_1 , R_2 , and R_3 , corresponding respectively to phonetic drift, grammaticalization, and semantic bleaching. The amplitwist language is then used to describe changes in the magnitude and orientation of cultural representations as they propagate through the cascade.

This proposal differs epistemically from the circuit application. In circuit theory, magnitude and phase are already independently defined physical quantities, and complex impedance supplies the scale-and-rotate structure directly. In the cultural case, the amplitwist representation is proposed as a model of change. The question is therefore whether the restriction to amplitwist geometry captures something that a generic dynamical description of linguistic change would not.

The first layer of the proposed cascade is phonetic drift,

$$R_1 : x_0 \mapsto x_1. \quad (176)$$

A linguistic form is transformed as its pronunciation changes across speakers, generations, or communities, associated here with a vowel-shift generator. The second layer is grammaticalization,

$$R_2 : x_1 \mapsto x_2, \quad (177)$$

where a lexical or aspectual form becomes increasingly grammatical and may acquire functions such as tense or aspect marking. The third layer is semantic bleaching,

$$R_3 : x_2 \mapsto x_3, \quad (178)$$

where repeated use weakens, broadens, or otherwise reorganizes some component of the expression's earlier lexical content.

The full cascade is therefore

$$x_3 = R_3 R_2 R_1 x_0. \quad (179)$$

At this level, the composition is simply a formal representation of sequential change. Nothing controversial follows from representing successive transformations by composable maps. The specifically amplitwist claim begins only when the R_k are assigned restricted scale-and-rotation structure and that structure is taken to explain measurable properties of cultural evolution.

The semantic history of *terrible* supplies a clear example. The earlier semantic force of the word was closely associated with terror or the capacity to inspire terror. In ordinary modern usage, the same lexical form can occur in expressions such as “a terrible sandwich,” where the judgment is strongly negative but no literal terror is involved. This development can be interpreted geometrically as a reduction in semantic amplitude.

That description is intuitively attractive. If a semantic representation assigns a coordinate or magnitude to terror-related content, one can imagine the corresponding component decreasing over historical time:

$$a_{\text{terror}}(t_2) < a_{\text{terror}}(t_1). \quad (180)$$

If other semantic associations simultaneously change, the trajectory could also be represented as a rotation through an embedding space. In a sufficiently low-dimensional idealization one might write

$$x(t_2) \approx sR(\theta)x(t_1). \quad (181)$$

The existence of such a representation, however, does not yet establish that semantic bleaching *is* an amplitwist. A generic transformation

$$x(t_2) = F(x(t_1)) \quad (182)$$

can also represent semantic change. Even within a linear approximation,

$$x(t_2) = Mx(t_1), \quad (183)$$

the matrix M need not satisfy the restrictive form

$$M = sR(\theta). \quad (184)$$

The same criterion used for cortex therefore reappears conceptually, although it has not been empirically examined here for the cultural case. To earn the strict amplitwist classification, the model would have to show that the relevant deformation is approximately conformal, or otherwise state explicitly that a broader definition of amplitwist is being proposed.

This matters because semantic change is naturally capable of anisotropy. One semantic dimension may weaken while another strengthens. Associations can split, merge, disappear, or emerge. A word may develop multiple senses whose frequencies evolve differently. Such changes are not automatically representable by uniform scaling plus rotation. If an

amplitwist model nevertheless captures them, that success would be informative precisely because the model is more restrictive than a generic deformation.

The proposed quantity called *cultural curvature* is intended to capture the cumulative effect of successive transformations, as a potentially measurable feature of linguistic and cultural trajectories rather than merely a metaphor. The natural mathematical motivation is that successive local transformations need not compose trivially. A representation can return near a previous location while carrying a residual change induced by the path through which it arrived there. In geometric language, such path dependence motivates curvature or holonomy-like quantities.

For cultural curvature to become an empirical quantity, however, several ingredients must be fixed. There must be a representation space,

$$\mathcal{C}, \tag{185}$$

a rule for mapping historical observations into states

$$x_t \in \mathcal{C}, \tag{186}$$

and a defined transformation or connection relating nearby states. Only then can a curvature-like observable be evaluated independently of arbitrary coordinate choices.

A validated measurement procedure satisfying these requirements is not supplied here. Consequently, *cultural curvature* should be treated as the name of a proposed observable, not as an already measured property of cultural systems.

The same qualification applies to *cognitive temperature*. The terminology suggests a scalar measure of dispersion, uncertainty, or volatility within a representational field. If a probability distribution $p(x)$ over cultural or cognitive states were specified, one could define quantities related to entropy,

$$S[p] = - \sum_x p(x) \log p(x), \tag{187}$$

or introduce a temperature-like parameter in a statistical model. But a variable becomes physically or empirically temperature-like only through the relations it satisfies, not through the choice of name. The cultural extension therefore requires an operational definition connecting the proposed temperature to observable distributions or transition rates.

Illustrating semantic bleaching with *terrible* does not, by itself, establish amplitwist geometry as a rigorous mathematical lingua franca for the humanities. It demonstrates that semantic change can be described suggestively in terms of changing magnitude and direction. Whether that description yields a distinctive quantitative theory remains an open research problem.

The artificial-intelligence extension takes the same programmatic structure and turns it into a proposed optimization objective. An amplitwist alignment loss can be written schematically as

$$\mathcal{L}_A = \sum_{k=1}^N \left\| A_{\text{LLM}}^{(k)}(x) - A_{\text{human}}^{(k)}(x) \right\|^2. \quad (188)$$

Here $A_{\text{LLM}}^{(k)}(x)$ denotes a proposed amplitwist representation associated with the model at layer or stage k , while $A_{\text{human}}^{(k)}(x)$ denotes the corresponding human target. The objective penalizes geometric divergence between the two.

As an optimization expression, the construction is straightforward. If two objects occupy a common normed representation space, then their discrepancy can be minimized:

$$\min_{\theta} \sum_{k=1}^N \left\| A_{\theta}^{(k)}(x) - A_{\text{target}}^{(k)}(x) \right\|^2. \quad (189)$$

The specifically amplitwist content lies elsewhere. One must establish what the operators $A_{\text{LLM}}^{(k)}$ and $A_{\text{human}}^{(k)}$ are, how they are estimated, why they inhabit a common comparable space, and what restriction makes them amplitwists rather than arbitrary latent representations.

Without those conditions, the objective reduces to generic representation matching. A sufficiently flexible embedding could always provide vectors whose Euclidean distance can be minimized. The amplitwist proposal becomes distinctive only if the restricted geometry contributes something independently testable.

For example, suppose the learned transformation at layer k is represented by

$$M^{(k)}. \quad (190)$$

A strict amplitwist account would require, in two dimensions,

$$M^{(k)} = s_k R(\theta_k), \quad (191)$$

rather than an arbitrary matrix. The corresponding restriction removes shear and anisotropic scaling from the hypothesis class. An alignment experiment could then compare the restricted model

$$\mathcal{M}_A = \{sR(\theta)\} \quad (192)$$

against a broader model

$$\mathcal{M}_G = \{M : M \in GL(2, \mathbb{R})\}, \quad (193)$$

or against another explicitly specified alternative appropriate to the representation space.

If the amplitwist-restricted objective predicted human judgments, behavioral responses, or transfer properties as well as or better than the generic alternative despite its smaller transformation class, the restriction would acquire empirical significance. If successful alignment required persistent shear or anisotropic deformation, the strict amplitwist hy-

pothesis would instead be disfavored.

No such comparison is reported here.

This point is essential to the status assignment. The cultural cascade, cultural curvature, cognitive temperature, and amplitwist alignment loss have not been subjected to the same scrutiny as the cortical operator or the Lamphron–Amplitwist–Sheaf sequence considered next. Absence of a negative verdict therefore cannot be interpreted as validation. These proposals remain *programmatic extensions*.

That status also prevents an unwarranted negative conclusion. Nothing in the cortical type mismatch proves that amplitwist geometry cannot become useful in historical linguistics or representation alignment. The cultural and AI applications make different claims about different objects. They must therefore be tested independently rather than inheriting either the success of circuit theory or the failure of the cortical definition.

The appropriate next step is the same one demanded by the methodological boundary in Section 4. Each proposal must identify what amplitwist structure forbids.

For semantic change, that means specifying a representation in which scale and rotation have independently interpretable meanings, then comparing the restricted transformation against generic deformation.

For cultural curvature, it means defining the state space, transport rule, and path-dependent quantity sufficiently precisely that different historical trajectories can yield measurably different curvature.

For cognitive temperature, it means identifying an observable distribution or dynamical law from which the proposed scalar can be estimated.

For AI alignment, it means defining the human and machine operators independently of the loss that is intended to compare them, then showing that amplitwist agreement predicts something not already captured by generic representation similarity.

Until those obligations are met, the cultural and alignment material should be read as a research program. Its contribution is to propose a common geometric vocabulary and a set of possible observables. Its unresolved question is whether that vocabulary compresses genuine shared structure or merely redescribes heterogeneous transformations in a common notation.

That question cannot be answered by extending the cascade one domain further. It requires independent examination of exactly the kind applied to the three-operator sequence considered next, which makes a different and more specific claim than that amplitwist language can merely be proposed across domains.

10. Epilogue: Lamphron, Amplitwist, and Sheaf Morphism

Status: component-level upgrades, with a taxonomy verdict for the ordered composition. The Lamphron–Amplitwist–Sheaf Morphism sequence (Flyxion, 2026c,e) extends the amplitwist program in its most demanding form. The preceding sections ask whether a common geometric vocabulary of scaling, rotation, admissibility, and coherence can be ex-

tended across mathematics, quantum representation, circuits, cortex, culture, and artificial intelligence. This sequence makes a different and stronger claim: three named operators form an ordered causal architecture,

$$L \longrightarrow A \longrightarrow S, \quad (194)$$

schematically,

$$\text{relax} \longrightarrow \text{transform} \longrightarrow \text{cohere}. \quad (195)$$

Examining this sequence does not retroactively decide every amplitwist application discussed above. Conversely, the exact amplitwist algebra and the successful circuit realization do not establish the causal ordering of L , A , and S . The two questions must remain separate.

Closer examination changes the status of the operators themselves. In particular, both Lamphron and Sheaf Morphism can be upgraded from an initial *terminological* classification to *constraining but unverified*. Those upgrades matter because a stronger verdict on the composition should not erase genuine mathematical content found in its components.

For Lamphron, the initial concern is that the term might merely rename ordinary relaxation or diffusion. If the proposed operator did nothing more than smooth a field according to a generic dissipative equation, then the new terminology would add no distinctive constraint. Closer inspection shows that this characterization is too weak. The Lamphron formulation contains a nonlinear entropy-gradient coupling, so its dynamics are not exhausted by generic diffusion.

Schematically, the relevant distinction is between a simple dissipative evolution such as

$$\partial_t u = D \nabla^2 u \quad (196)$$

and an evolution in which the relaxation dynamics are coupled nonlinearly to an entropy field or its gradient,

$$\partial_t u = D \nabla^2 u + \mathcal{N}(u, \nabla S, \dots), \quad (197)$$

with

$$\mathcal{N} \neq 0. \quad (198)$$

The precise significance of the additional term is that it restricts the allowed dynamics. Lamphron therefore cannot be dismissed as a new label attached to unconstrained diffusion.

That result justifies the classification

$$\boxed{\text{Lamphron: constraining but unverified.}} \quad (199)$$

The qualification *unverified* remains essential. The stronger claims attached to Lam-

phron require more than the existence of a nonlinear coupling. In particular, the proposed route from the Lamphron dynamics to the relevant Beltrami or admissibility condition has not been established merely by writing the coupled evolution. The required convergence claim and its associated parameter bounds remain proof obligations. There is real mathematical structure here without the theorem needed to support the full causal role assigned to the operator.

Sheaf Morphism receives a parallel upgrade for a different reason. The initial terminological concern is that “sheaf” might function metaphorically as a synonym for local-to-global coherence. The underlying material instead contains genuine presheaf structure, restriction maps, and cohomological obstruction language. The relevant mathematical idea can be stated independently of the surrounding terminology.

Let $\mathcal{F}$ be a presheaf on a space X . For an inclusion

$$V \subseteq U, \quad (200)$$

there is a restriction map

$$\rho_{UV} : \mathcal{F}(U) \longrightarrow \mathcal{F}(V). \quad (201)$$

Suppose a collection of local sections

$$s_i \in \mathcal{F}(U_i) \quad (202)$$

is given on a cover $\{U_i\}$ of X . A global-coherence claim asks whether there exists

$$s \in \mathcal{F}(X) \quad (203)$$

such that

$$s|_{U_i} = s_i \quad (204)$$

for every i .

Local agreement does not automatically guarantee such a global section. Cohomological data can record an obstruction to gluing (Hartshorne, 1977). In the formulation used here, the first cohomology group

$$H^1 \quad (205)$$

plays this obstruction role. A nonvanishing class can therefore witness failure of global assembly even when substantial local compatibility is present.

The important point is not that every gluing problem is exhausted by the slogan

$$H^1 = 0. \quad (206)$$

The relevant presheaf, cover, coefficient structure, and obstruction theory must be spec-

ified before such a condition has the intended force. The point is that the proposal contains enough genuine local-to-global mathematics that “Sheaf Morphism” is not merely decorative vocabulary.

Its status is consequently

$$\boxed{\text{Sheaf Morphism: constraining but unverified.}} \quad (207)$$

The remaining verification obligation is application-specific. The actual local data and restriction maps must be defined, the relevant obstruction class must be identified, and the claimed transition to global coherence must follow from the stated hypotheses. The presence of sheaf-theoretic language alone does not discharge those obligations, but the presence of a real obstruction structure is enough to distinguish the proposal from a purely terminological analogy.

Amplitwist occupies a different position in this examination. Its elementary mathematics is derived (Section 2), but the cortical application used in the three-operator sequence remains terminological in the specific sense examined here. The relevant operator is

$$J = \mathcal{A}R \quad (208)$$

with general symmetric $\mathcal{A}$, while the strict amplitwist requires

$$J = sR. \quad (209)$$

The missing condition is therefore

$$\boxed{\mathcal{A}_0 = \mathcal{A} - \frac{\text{tr } \mathcal{A}}{2}I = 0.} \quad (210)$$

Thus the component ledger is deliberately nonuniform:

$$\begin{aligned} L & : \text{constraining but unverified,} \\ A_{\text{cortical}} & : \text{terminological pending } \mathcal{A}_0 = 0, \\ S & : \text{constraining but unverified.} \end{aligned} \quad (211)$$

This component-level result does not, however, establish the arrows

$$L \longrightarrow A \longrightarrow S. \quad (212)$$

The ordering is treated here as a separate claim and subjected to two independent tests.

The first is algebraic. Consider a toy realization of the three operators on two overlapping local patches, in which each operator is instantiated in a simple form preserving its proposed qualitative role: L as ordinary diffusion between the two local values, A as a shared rotation-scaling matrix acting identically on each patch, and S as projection onto their common average. The question is whether composition order itself forces the advertised sequence. It does not. In this toy realization the operators commute:

$$LA = AL, \quad (213)$$

$$AS = SA, \quad (214)$$

and

$$LS = SL. \quad (215)$$

Consequently,

$$LAS = LSA = ALS = ASL = SLA = SAL. \quad (216)$$

This result is decisive against one possible basis for the ordering. If every permutation gives the same result in a legitimate simple realization, then the causal significance of

$$L \rightarrow A \rightarrow S \quad (217)$$

cannot be derived merely from the algebra of the three operators.

The toy result does not prove that every richer realization must commute. Nonlinearity, state dependence, changing domains, boundary conditions, or history dependence could break commutativity. Its significance is instead diagnostic. Order dependence is not automatic. If the proposed causal sequence is to be necessary, the theory must contain some additional structure that makes the wrong order fail.

That observation motivates a second, independent test for such structure.

A genuine causal ordering could survive algebraic commutativity if the operators have asymmetric domains or preconditions. For example, Lamphron might produce states on which Amplitwist is defined,

$$\text{Im}(L) \subseteq \text{Dom}(A), \quad (218)$$

while the reverse ordering fails because

$$\text{Im}(A) \not\subseteq \text{Dom}(L) \quad (219)$$

for the relevant states. Likewise, Amplitwist might establish a condition required before a Sheaf Morphism can glue local data coherently.

More generally, one could have predicates

$$P_A(x) \quad (220)$$

and

$$P_S(x) \quad (221)$$

such that

$$L(x) \implies P_A, \quad (222)$$

and

$$A(L(x)) \implies P_S, \quad (223)$$

while the reversed compositions fail to satisfy the corresponding prerequisites. Such implications would supply a genuine gating mechanism even if simplified linear versions of the operators happened to commute.

No established domain or precondition implications of this kind are currently available. Relaxation, transformation, and coherence are repeatedly presented in the order

$$L \rightarrow A \rightarrow S, \quad (224)$$

but presentation in that order is not itself evidence that the order is necessary. No independently established result shows that Lamphron creates a condition without which Amplitwist cannot act, or that Amplitwist creates a condition without which the Sheaf stage cannot glue. This absence is independent of the toy-system calculation.

The two findings therefore reinforce one another.

The toy system establishes

$$\boxed{\text{order is not forced by the simplest operator algebra.}} \quad (225)$$

Direct examination of the primary sources establishes that

$$\boxed{\text{no independent precondition theorem currently forces the order.}} \quad (226)$$

Together they support a central verdict:

$$\boxed{L \rightarrow A \rightarrow S \text{ is presently a taxonomy, not yet a causal theory.}} \quad (227)$$

The word *taxonomy* is not dismissive. The sequence distinguishes three conceptually different operations: relaxation, transformation, and local-to-global coherence. That classification may be useful even if the operations can occur in different orders. A taxonomy becomes a causal theory only when the arrows acquire content beyond temporal or expository sequencing.

This converts the sequence into a set of explicit proof obligations.

For Lamphron, the nonlinear entropy-gradient dynamics must be connected to the downstream condition claimed for them. It is not sufficient that the operator relaxes a field in a distinctive way. The relevant convergence or admissibility theorem must follow.

For Amplitwist, the operator used in the causal chain must satisfy the defining restriction of the amplitwist class. In the cortical formulation this requires

$$\|\mathcal{A}_0\| = 0 \quad (228)$$

or an independently justified approximation under which the anisotropic component can be neglected.

For Sheaf Morphism, the local data, restriction maps, and gluing problem must be instantiated for the actual system. If vanishing of a first cohomological obstruction is the proposed criterion, then the relevant statement must be established in that system:

$$H^1 = 0, \quad (229)$$

with the precise cohomological object and hypotheses specified.

Finally, and separately, the arrows themselves require a gating theorem. A successful theory must establish something of the form

$$\text{Im}(L) \subseteq \text{Dom}(A), \quad (230)$$

$$\text{Im}(A \circ L) \subseteq \text{Dom}(S), \quad (231)$$

together with a meaningful failure of at least one competing order, or an equivalent state-dependent condition that makes the proposed ordering necessary.

Without this final obligation, proving each operator individually would still not prove

$$L \rightarrow A \rightarrow S. \quad (232)$$

Three valid operators do not automatically constitute a valid causal chain.

This closer examination clarifies rather than collapses the amplitwist program. The elementary amplitwist remains a derived mathematical object. Its circuit realization remains strong. The quantum reconstruction remains an unverified mathematical extension. The cortical identification has a precise unresolved type condition. The consciousness claim remains downstream of that condition and of a further unproved transition. The population-vector proposal remains falsifiable and untested. The cultural and alignment applications remain programmatic. Lamphron and Sheaf Morphism, on closer inspection, contain more constraining mathematics than an initial terminological reading would suggest. The ordered composition of the three, however, remains unsupported as a necessary causal architecture.

The resulting picture is deliberately uneven because the evidence is uneven.

This is the central methodological consequence of the discipline applied throughout. Mathematical vocabulary should neither be protected because it belongs to an ambitious unifying program nor rejected merely because some of its extensions outrun their derivations. The appropriate unit of evaluation is the claim. A composition proof remains a composition proof even if a later neuroscientific identification fails. A nonlinear coupling remains mathematically substantive even if its convergence theorem is missing. A cohomology

logical obstruction remains a genuine constraint even if its application has not been verified. A falsifiable prediction remains worth testing even when the interpretation that generated it is underdetermined.

The amplitwist program is therefore strongest when its own geometric discipline is applied to its epistemology. Distinctions should be preserved until an explicit operation licenses their collapse. The relation $\asymp$ began with exactly that lesson: two objects may become equivalent under a specified limiting condition without having been identical beforehand.

The same discipline applies to theories. Analogy is not derivation. A proposed operator is not an instantiated operator. An instantiated operator is not a causal sequence. A taxonomy is not yet a theory. None of these distinctions needs to be permanent, but each requires a specific argument before it may be removed.

11. Conclusion: Preserve the Distinctions

The amplitwist program began with a relation designed to mark a distinction. Two quantities could remain unequal at every finite stage while becoming ultimately equal under a specified limit,

$$A \asymp B \iff \lim_{\varepsilon \rightarrow 0} \frac{A}{B} = 1. \quad (233)$$

The notation is modest, but the methodological principle behind it is consequential. Distinctions should not be erased merely because two objects resemble one another. They should be collapsed only when an explicit operation establishes the conditions under which the distinction ceases to matter.

The same principle provides an appropriate way to evaluate the larger amplitwist program.

At its mathematical core, the amplitwist is exact. The operator

$$A = sR(\theta) \quad (234)$$

combines uniform scaling with rotation, and its composition law follows directly:

$$s_1 R(\theta_1) s_2 R(\theta_2) = s_1 s_2 R(\theta_1 + \theta_2). \quad (235)$$

Amplitude multiplies and twist adds. The resulting transformation is conformal, invertible for $s > 0$, and closed under composition. None of these properties depends on a cross-domain interpretation.

Circuit theory supplies an equally secure physical realization because magnitude and phase already belong to the standard formalism. Complex impedance scales and rotates phasors. Resonance can produce a zero-phase condition. Rectification supplies an elementary example of non-invertibility through non-injectivity. Capacitor state demonstrates

how a present variable can compress multiple possible histories. Boolean gates make information loss countable, while Landauer’s principle and Johnson–Nyquist noise connect information processing, dissipation, and fluctuation to established thermodynamic results.

These results form the derived core of the program.

Beyond that core, the status changes.

The real flag-space formulation of quantum mechanics contains exact elementary mathematics. The matrix

$$J_F = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \quad (236)$$

satisfies

$$J_F^2 = -I, \quad (237)$$

and therefore realizes complex multiplication explicitly in a real space. The composite-system construction also identifies a genuine representational problem: naive local realification can retain distinctions that the corresponding global complex state does not possess. An admissibility quotient is a mathematically meaningful response. What has not been established here is the stronger claim that the resulting construction provides a complete reconstruction of arbitrary composite quantum mechanics. Its status is therefore not negative but unverified.

The cortical extension differs because a specific type condition has already been located. The proposed local operator

$$J = \mathcal{A}R \quad (238)$$

is more general than the strict amplitwist

$$J = sR. \quad (239)$$

The discrepancy is measured by

$$\mathcal{A}_0 = \mathcal{A} - \frac{\text{tr } \mathcal{A}}{2} I. \quad (240)$$

Only when

$$\boxed{\mathcal{A}_0 = 0} \quad (241)$$

does the general symmetric amplification collapse to isotropic scaling and recover the strict amplitwist. The cortical claim therefore has a precise outstanding obligation rather than an indefinite request for “more evidence.”

The consciousness proposal inherits that obligation but is not refuted by it. Even if cortical isotropy were established, a further result would still be required to connect local operator coherence to the proposed critical transition and that transition to consciousness. The

theory therefore remains programmatic downstream of two distinguishable obligations.

The population-vector prediction occupies a different category. It can be exposed to observation. Task switching is predicted to generate rotation-like trajectories in neural population space, and the strict amplitwist version strengthens this to an approximate conformality condition,

$$\|\mathcal{A}_{0,t}\| \approx 0. \quad (242)$$

Persistent anisotropic deformation, representation-dependent apparent rotation, or systematic predictive superiority of a more general transformation class would count against that hypothesis. The prediction is untested, but it is not epistemically empty.

The cultural and artificial-intelligence extensions remain programs. Phonetic drift, grammaticalization, and semantic bleaching can be represented as successive transformations. The history of *terrible* provides an intuitive example of changing semantic magnitude. Cultural curvature, cognitive temperature, and an amplitwist alignment loss suggest possible quantitative developments. None has yet received the examination required to establish that strict amplitwist geometry contributes a distinctive constraint rather than a common vocabulary for generic transformation.

The Lamphron–Amplitwist–Sheaf Morphism material sharpens the same lesson.

Closer inspection upgrades Lamphron from terminological to constraining but unverified because its nonlinear entropy-gradient coupling contributes structure beyond generic diffusion. Sheaf Morphism receives the same upgrade because the relevant construction contains genuine local-to-global machinery, including restriction maps and a cohomological obstruction represented through H^1 . In both cases, mathematical content is found where an initial terminological reading would have been too severe.

The ordered composition does not receive the same upgrade.

A toy realization shows that the three operators can commute, so the order

$$L \rightarrow A \rightarrow S \quad (243)$$

does not follow from their simplest algebra. No domain or precondition implication forces Lamphron to precede Amplitwist or Amplitwist to precede the Sheaf stage. The sequence therefore functions as a taxonomy of relaxation, transformation, and coherence rather than as a demonstrated causal theory.

That verdict is itself provisional in the proper mathematical sense. It identifies what would change it. A future result could establish a gating condition such as

$$\text{Im}(L) \subseteq \text{Dom}(A), \quad (244)$$

followed by

$$\text{Im}(A \circ L) \subseteq \text{Dom}(S), \quad (245)$$

while showing that competing orders fail the required preconditions. Such a result

would supply content to the arrows that the present taxonomy lacks.

The resulting status ledger is therefore intentionally uneven:

Claim	Current status
Ultimately-equal calculus	Derived
Strict amplitwist algebra	Derived
Circuit magnitude–phase realization	Derived / established physics
Quantum flag-space reconstruction	Unverified extension
Cortical strict-amplitwist identification	Type-mismatch result at the present formulation
Consciousness as coherence tile	Programmatic, downstream of unmet conditions
Population-vector rotation	Falsifiable but untested
Cultural amplitwist cascade	Programmatic extension
Amplitwist alignment objective	Programmatic extension
Lamphron	Constraining but unverified
Sheaf Morphism	Constraining but unverified
$L \rightarrow A \rightarrow S$ as necessary causal order	Taxonomy, not yet theory

The purpose of this ledger is not to rank the ideas from respectable to disposable. It records what kind of statement each one presently is. A derived result, an unverified extension, a type-mismatch result, an empirical prediction, and a programmatic conjecture require different next steps. Treating all of them as equally established would obscure those next steps. Treating all of them as equally speculative would do the same.

The tempered account therefore leaves the amplitwist program neither universally confirmed nor generally repudiated. Its mathematical nucleus survives intact. Its strongest physical application survives intact. Some extensions become explicit research programs. One prominent cross-domain identification acquires a precise correction criterion. A speculative neuroscientific interpretation yields a testable prediction. Two later operators prove more substantive than an initial terminological reading suggested, while their proposed causal ordering remains unproved.

This is a more useful form of unity than the claim that one operator has already been discovered everywhere.

A unifying theory earns its unity by surviving distinctions, not by suppressing them. If the same structure genuinely recurs across quantum representation, circuits, cortical dynamics, cultural evolution, and machine alignment, independent analysis should progressively force those descriptions toward the same restricted mathematics. Where it does not, the differences should remain visible.

The original relation provides an appropriate final notation for that standard. Resemblance alone does not justify equality. Even ultimate equality requires a limit:

$$A \asymp B \iff \lim_{\varepsilon \rightarrow 0} \frac{A}{B} = 1. \quad (246)$$

The burden of a cross-domain theory is analogous. It must specify the operation under which apparently different structures become equivalent, and then demonstrate that the required limit, quotient, constraint, or invariant actually exists. Until then, the distinction remains.

References

- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., and Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487(7405):51–56.
- Flyxion (2026a). Amplitwist cascades: Recursive epistemic geometry in cultural-semantic evolution. Unpublished manuscript.
- Flyxion (2026b). Amplitwist operators in cortical columns. Unpublished manuscript.
- Flyxion (2026c). The Lamphrodynamic self-modulation protocol. Unpublished manuscript.
- Flyxion (2026d). Rotation before number: Admissibility and the real structure of quantum mechanics. Unpublished manuscript.
- Flyxion (2026e). The TARTAN update as a sheaf morphism: Local-to-global coherence in the RSVP framework. Unpublished manuscript.
- Hartshorne, R. (1977). *Algebraic Geometry*. Springer-Verlag, New York.
- Johnson, J. B. (1928). Thermal agitation of electricity in conductors. *Physical Review*, 32(1):97–109.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191.
- Needham, T. (1997). *Visual Complex Analysis*. Oxford University Press, Oxford.
- Needham, T. (2021). *Visual Differential Geometry and Forms: A Mathematical Drama in Five Acts*. Princeton University Press, Princeton, NJ.
- Nyquist, H. (1928). Thermal agitation of electric charge in conductors. *Physical Review*, 32(1):110–113.
- Sedra, A. S. and Smith, K. C. (2015). *Microelectronic Circuits*, 7th edition. Oxford University Press, New York.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423.