Amplitwist Cascades, Biological Priors, and the Geometry of Cognitive Transformation

Flyxion

Independent Researcher

May 2026

Abstract

This essay develops a geometric account of cognitive transformation grounded in the amplitwist operator formalism. The central thesis is that cognition operates within a structured transformation algebra whose expressive capacity exceeds its biological use, and that this gap—far from being incidental—encodes the prior that shapes how biological systems actually compute. We formalize the amplitwist cascade as a universal transformation language, embed it within RSVP-style field dynamics, and derive efficiency bounds from the spectral geometry of the representational manifold. We then introduce a biological prior modelled on hierarchical central-pattern-generator (CPG)-like dynamics, which constrains the system to low-dimensional, rhythmically stable trajectories. Propositional representation is reconceived as an equivalence class of perturbation-stable trajectories rather than as static encoding. The two-layer structure—expressive envelope versus biological prior—generates falsifiable predictions about phase alignment, vorticity convergence, and cascade ordering in both human behaviour and neural population dynamics. An experimental design using compositional and non-compositional morphological transformation tasks is proposed to test these predictions, and an alignment loss metric is defined to compare operator geometry between humans and language models.

Contents

1 Introduction: From Expressivity to Constraint **2**

2 The Amplitwist Operator as a Universal Transformation Basis **2**
  2.1 Operator definition 2
  2.2 Cascade composition 3
  2.3 Two regimes of operation 3

3 Field-Theoretic Embedding and Phase Alignment **3**
  3.1 RSVP fields as state space 3
  3.2 Phase alignment and the amplitwist angle 4
  3.3 Entropy weighting as biological prior 4

4 Hybrid Dynamics and Convergence **4**

5 Spectral Geometry and Efficiency Constraints **5**
  5.1 Eigenmode structure 5
  5.2 Efficiency bound 5

6 The Biological Prior: Dynamical Constraint via CPG-Like Structures **6**
  6.1 The capacity–implementation gap 6
  6.2 CPG-like dynamics as a trajectory prior 6
  6.3 Errors under the prior 6

7 Representation as Trajectory Equivalence **7**
  7.1 Against static encoding 7
  7.2 Trajectory equivalence 8
  7.3 CPG dynamics as a trajectory basis 8

8 Predictions About Cognitive Dynamics **8**

9 Experimental Design: Compositional vs Non-Compositional Tasks **9**
  9.1 Controlled transformation domain 9
  9.2 Compositional task family 9
  9.3 Non-compositional task family 10
  9.4 Trajectory measurement 10

10 Human vs Model Comparison **10**
  10.1 Alignment loss 10
  10.2 Predicted asymmetries 11

11 Interpretation: Geometry as the Invariant Substrate **11**

12 Falsifiability and Outcomes **12**

13 Conclusion 12

A Amplitwist Operator Algebra 14
A.1 Local Operator Definition 14
A.2 Composition Law 14
A.3 Local Linearization 14
A.4 Density in Function Space 14
A.5 Cascades as Operator Sequences 15
A.6 Separation of Capacity and Implementation 15

B Phase Alignment and Vorticity 15
B.1 Field Structure 15
B.2 Phase Alignment 15
B.3 Alignment Functional 15
B.4 Vorticity 16
B.5 Phase–Vorticity Relationship 16
B.6 Discrete Approximation 16
B.7 Coherence Criterion 16

C Entropy Weighting as a Biological Prior 16
C.1 Entropy Field and Weighting 16
C.2 Variational Formulation 17
C.3 Admissibility and Stability 17
C.4 Entropy–Vorticity Coupling 17
C.5 Regularised Energy Functional 17

D Hybrid Dynamics and Trotter–Kato Convergence 17
D.1 Semigroup Evolution 17
D.2 Discrete Operator Updates 18
D.3 Linear Trotter–Kato Framework 18
D.4 Nonlinear Extension (Lie–Trotter Splitting) 18
D.5 Convergence Result 18

E Spectral Geometry and Efficiency Bounds 19
E.1 Laplacian and Eigenstructure 19
E.2 Efficiency Metric and Bound 19
E.3 Geometric Interpretation 19
E.4 Spectral Filtering and the Biological Subspace 20

F Dynamical Constraint Model (CPG Prior) 20
F.1 Trajectory Spaces and Biological Realisation 20
F.2 Constraint Functional 20
F.3 Low-Dimensional Manifold Constraint 20

G Admissibility Constraints and the Revised Categorical Principle **21**
G.1 Motivation 21
G.2 Admissible Transformation Set 22
G.3 Constraint Functional 22
G.4 The Revised Categorical Principle 22
G.5 Compositional Closure 22
G.6 Relation to Biological Prior and Spectral Structure 23
G.7 Alignment Without Identical Outputs 23

H Representation as Trajectory Equivalence **23**
H.1 Equivalence Relation 23
H.2 Quotient Space and Propositions 24
H.3 Stability Under Composition 24
H.4 Metric Structure 24

I Empirical Observables and Measurement Protocol **24**
I.1 Fixed Representation Manifold 24
I.2 Discrete Trajectory and Phase 24
I.3 Entropy Proxies 24
I.4 Vorticity, Aggregates, and Layer Segmentation 24
I.5 Human–Model Alignment Loss 25

J Non-Compositional Transformations and Failure Modes **25**
J.1 Compositional and Non-Compositional Transformations 25
J.2 Failure Signatures 25
J.3 Three-Regime Prediction 26
J.4 Output vs Path Dissociation 26

K Gibbs Unification: Measure, Energy, and Admissibility **26**
K.1 Motivation 26
K.2 Total Constraint Energy 26
K.3 Admissibility as a Typical Set 27
K.4 Pushforward to Function Space 27

L Variational Alignment and the Learning Objective **28**
L.1 Human and Model Gibbs Measures 28
L.2 KL Divergence as the Canonical Alignment Quantity 28
L.3 Alignment Loss as a KL Approximation 28
L.4 Variational Alignment Objective 29
L.5 Gradient of the Variational Objective 29
L.6 Pathwise Alignment Criterion 29

1. Introduction: From Expressivity to Constraint

The central thesis of this essay is not merely that cognition is nonlinear or compositional, but that it operates within a structured transformation algebra whose expressive capacity strictly exceeds its biological use. The distinction that matters is between what a system *could* compute and what it *actually* computes, between the envelope of possible transformations and the narrow, prior-shaped corridor through which biological computation travels.

Standard accounts of neural computation tend to focus on one side of this distinction or the other. Representational accounts emphasise the richness of what the brain can encode. Dynamical systems accounts emphasise the constrained, low-dimensional structure of what it actually does. The position developed here holds that both sides are necessary, and that the relationship between them—the gap between expressive capacity and biological implementation—is itself the primary theoretical object.

This gap is not a residual to be explained away. It encodes the prior that shapes cognition: the preference for rhythmically stable, compositionally ordered, dynamically supportable trajectories over the vast space of transformations that the underlying operator algebra could in principle compute. Formalising this gap, and making it empirically testable, is the aim of the essay.

The argument proceeds in four stages. Sections 2–5 develop the mathematical core: the amplitwist operator as a universal transformation basis, its embedding in field dynamics, and the spectral constraints governing efficiency. Sections 6–7 introduce the biological prior and reformulate propositional representation as trajectory equivalence. Sections 8–10 derive empirical predictions and propose an experimental design. Sections 11–13 interpret the framework and assess its falsifiability.

2. The Amplitwist Operator as a Universal Transformation Basis

2.1. Operator definition

The starting point is the classical result from complex analysis. For a holomorphic function $f\colon \mathbb{C} \to \mathbb{C}$, the derivative at a point z_0 takes the form $f'(z_0) = \rho\, e^{i\theta}$, encoding a local similarity transformation that simultaneously scales by $\rho > 0$ and rotates by $\theta \in [0, 2\pi)$ [1, 2]. Tristan Needham's *Visual Complex Analysis* gives this decomposition the name *amplitwist*, capturing the dual action of amplitude and twist on infinitesimal neighbourhoods.

Generalising to a smooth n-dimensional representational manifold M, we define the *amplitwist operator* at a point $x \in M$ as the triplet

$$\mathcal{A}(x) = \rho(x)\, e^{i\theta(x)}\, \pi(x), \tag{1}$$

where $\rho(x) \in \mathbb{R}_{>0}$ is an amplitude modulation (gain or precision), $\theta(x) \in \mathbb{R}$ is a phase encoding the direction of transformation in the tangent space T_xM, and $\pi(x)$ is a nonlinear projection or contraction that introduces the nonlinearity required for function approximation beyond the affine regime.

2.2. Cascade composition

Under sequential application, amplitwist operators compose into cascades. Given a sequence of operators $\mathcal{A}_1, \ldots, \mathcal{A}_N$ acting at successive stages, the composed transformation is

$$\mathcal{A}^{(N)} = \mathcal{A}_N \circ \cdots \circ \mathcal{A}_1. \tag{2}$$

With sufficient depth and nonlinearity, this compositional system defines a dense function class over smooth maps $M \to \mathbb{R}^n$. The amplitwist cascade is therefore analogous in expressive power to neural network universal approximation theorems [3, 8], but grounded in geometrically interpretable local operators rather than in abstract activation functions. This is the *capacity layer*: the formal claim that the operator algebra, given unconstrained composition, can approximate any sufficiently regular transformation.

2.3. Two regimes of operation

The capacity claim must be distinguished from any claim about how biological systems actually use this capacity. Formally, two regimes are identified. In *Regime A* (universal approximation), the system exploits the full depth of the cascade to approximate arbitrary continuous functions; intermediate states are determined solely by the approximation target. In *Regime B* (constrained control), the system operates through a small number of structured generators—rhythmic, hierarchical, embodied—and computation is compositional in the sense that each stage stabilises before the next is engaged. The central claim of this essay is that biological cognition predominantly occupies Regime B, and that the entropy weighting introduced in Section 3 implements the corresponding prior.

3. Field-Theoretic Embedding and Phase Alignment

3.1. RSVP fields as state space

The amplitwist operator acts on a structured state space defined by three coupled fields over M. The *scalar field* $\Phi\colon M \to \mathbb{R}$ represents semantic salience or potential. The *vector field* $\mathbf{v}\colon M \to TM$ represents conceptual velocity, encoding the direction and speed of representational change. The *entropy field* $S\colon M \to \mathbb{R}_{>0}$ represents cognitive uncertainty or the admissibility of a transformation under the biological prior. This triplet is the state space of the Relativistic Scalar–Vector Plenum (RSVP) framework adapted to the cognitive setting.

3.2. Phase alignment and the amplitwist angle

The central geometric quantity is the angle between the vector field $\mathbf{v}$ and the gradient of the scalar field Φ. This angle measures the degree to which the direction of representational motion aligns with the direction of increasing salience. Formally,

$$\theta(x) \;=\; \arccos\left(\frac{\mathbf{v}(x)\cdot\nabla\Phi(x)}{\|\mathbf{v}(x)\|\,\|\nabla\Phi(x)\|+\epsilon}\right), \tag{3}$$

where $\epsilon > 0$ is a small numerical stabiliser. When $\theta \approx 0$, the system moves directly along the salience gradient: the transformation is maximally coherent. When $\theta \approx \pi/2$, motion is orthogonal to the gradient, generating rotational structure without progress. The quantity $\xi = |\nabla \times \hat{\mathbf{v}}|$, where $\hat{\mathbf{v}} = (\cos\theta, \sin\theta)$, is the *vorticity* of the phase-weighted field. High vorticity signals inefficiency or instability in the transformation cascade.

3.3. Entropy weighting as biological prior

The layer-k amplitwist, weighted by entropy, is

$$\mathcal{A}^{(k)}(x) \;=\; w_k(x)\cdot\mathcal{A}(\mathfrak{R}_k(x)), \qquad w_k(x) \;=\; \exp(-\lambda S(x)), \tag{4}$$

where $\mathfrak{R}_k$ is the semantic deformation applied at layer k. The weight w_k suppresses transformations in regions of high entropy. Interpreted as a regulariser, this is a standard smoothing device. Interpreted as a biological prior, it is a hypothesis about what the nervous system computes: it continuously evaluates which transformations can be supported by available dynamical primitives, and down-weights trajectories that exceed that support. This interpretation is the key conceptual move of Section 6.

4. Hybrid Dynamics and Convergence

Amplitwist cascades do not operate in isolation from temporal dynamics. Between discrete operator applications, the state evolves continuously according to the RSVP field equations. The combined system is a *hybrid dynamical system*: continuous semigroup evolution punctuated by discrete transformation events.

Let $T(t)$ denote the strongly continuous semigroup of field evolution [4, 5], and let P_k denote the operator applied at event time t_k. The hybrid trajectory takes the form

$$T(t_N - t_{N-1})\,P_{N-1}\,\cdots\,T(t_1 - t_0)\,P_0\,T(t_0). \tag{5}$$

A Trotter–Kato-style approximation theorem ensures that, as event spacing shrinks to zero and the discrete grid refines, the hybrid trajectory converges:

$$\lim_{h\to 0,\,\delta t\to 0}\big(S_h(\delta t)\big)^{\lfloor t/\delta t\rfloor} \;=\; T(t)\circ P_t \tag{6}$$

in the operator norm on the relevant Sobolev space, directly analogous to the

Trotter product formula [6, 7]. This convergence result grounds the discrete cascade model in functional analysis, ensuring that operator composition inherits stability guarantees from the continuous semigroup. The cascade is not an algorithmic approximation whose relationship to the underlying dynamics is merely hoped for: in the continuum limit it converges to a continuous operator flow that preserves the same invariant structure as the underlying field dynamics.

5. Spectral Geometry and Efficiency Constraints

5.1. Eigenmode structure

The efficiency of transformation propagation across M is constrained by its spectral geometry. The Laplace–Beltrami operator Δ_M admits a spectral decomposition [9, 10] whose eigenmodes ψ_n and eigenvalues $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \cdots$ organise the possible patterns of activation on M. Low-frequency modes (small λ_n) are broad and spatially integrative; high-frequency modes are local and transient. A transformation event excites these modes with amplitude proportional to the overlap of the event kernel with each eigenfunction. Sequential cascades correspond to traversals in eigenmode space, with each layer engaging progressively deeper or higher-frequency structure.

5.2. Efficiency bound

The epistemic efficiency ratio $\eta^{(N)}$, defined as the ratio of useful phase alignment to total cascade cost, satisfies

$$\eta^{(N)} \geq \frac{\lambda_1(M)}{N \cdot \max_j \|\epsilon_j T_j\|_\infty}, \tag{7}$$

where $\lambda_1(M)$ is the first nonzero Laplacian eigenvalue of M, N is the cascade depth, ϵ_j are the deformation intensities, and $T_j \in \mathfrak{so}(n)$ are the Lie-algebra generators of the semantic rotations at each layer. The bound has a transparent interpretation: the higher the spectral gap of the manifold, the more easily transformations propagate between regions and the more efficient the cascade. Conversely, a manifold with a small spectral gap offers geometric resistance to propagation, requiring deeper cascades or larger deformations to achieve the same alignment. This connects local operator structure to global manifold geometry via eigenmodes and traveling waves, in direct correspondence with empirical observations of cortical dynamics in ECoG and MEG.

6. The Biological Prior: Dynamical Constraint via CPG-Like Structures

6.1. The capacity–implementation gap

The formal capacity established in Section 2 is unlimited: given sufficient depth, amplitwist cascades can approximate any sufficiently regular function. But biological systems do not search this space uniformly. The central empirical fact is that neural computation is strikingly low-dimensional relative to the dimensionality of the spaces it nominally acts upon [15, 16]. Population dynamics occupy a small number of modes; motor trajectories are organised around rhythmic generators; perceptual representations cluster into stable attractors. The capacity–implementation gap is not a deficiency but a structural signature: it reveals the shape of the prior that biological systems bring to computation.

6.2. CPG-like dynamics as a trajectory prior

The natural model for this prior is a hierarchy of central pattern generators (CPGs) or CPG-like oscillatory structures [14, 13]. CPGs are neural circuits capable of producing coordinated, rhythmic output without requiring continuous sensory input. Hierarchically composed CPG chains generate low-dimensional, stable, perturbation-resistant trajectories through the high-dimensional neural state space. CPGs are invoked here not as a literal substrate for all cognition, but as the canonical example of low-dimensional, phase-stable dynamical generators; the claim is that cortical computation inherits this structural constraint, not that it reduces to motor circuits. The amplitwist cascade, from the perspective of biological implementation, is not searching the full function class: it preferentially occupies trajectories that can be parameterised and stabilised by such generators.

This is the *biological prior layer*. Formally, the prior assigns high probability to trajectories that satisfy three conditions: they are low-dimensional relative to the ambient space of M; they are rhythmically structured, admitting a generator decomposition; and they are stable under perturbation, meaning small displacements return to the trajectory rather than diverging. The entropy weighting $w_k = \exp(-\lambda S)$ implements this prior in the operator algebra: regions of high entropy correspond to transformations that cannot be stabilised by available dynamical primitives and are accordingly suppressed, consistent with variational frameworks that penalise uncertainty and surprise [17].

6.3. Errors under the prior

A key prediction follows from the direction of failure. When a task requires a transformation that cannot be decomposed into CPG-supportable stages—a genuinely non-compositional jump through semantic space—biological systems should not fail randomly. They should fail in the direction of the nearest decomposable transformation, because entropy weighting continuously pulls the trajectory toward the highest-prior path even when that path is incorrect. This directional failure is qualitatively distinct from the behaviour of a pure function approximator, and it constitutes a falsifiable signature of the prior.

The deepest consequence of the two-layer picture is that biological cognition is not best understood as function approximation, but as constrained trajectory selection in a transformation space. Neural networks approximate functions; brains select trajectories. The following proposition captures this quantitatively.

Proposition 6.1 (Prior-Induced Measure Concentration). *Let $\mathcal{F}_{\text{all}}$ denote the full amplitwist function class, and let μ_{bio} be the probability measure induced by entropy-weighted dynamics. Then for any $\epsilon > 0$, there exists a subset $\mathcal{F}_\epsilon \subset \mathcal{F}_{\text{all}}$ such that*

$$\mu_{\text{bio}}(\mathcal{F}_\epsilon) \geq 1 - \epsilon, \qquad \dim(\mathcal{F}_\epsilon) \ll \dim(\mathcal{F}_{\text{all}}).$$

In words: the entropy-weighted prior concentrates biological computation on a low-dimensional subset of the full function space, even though the full space remains accessible in principle. The gap between $\mathcal{F}_\epsilon$ and $\mathcal{F}_{\text{all}}$ is the mathematical form of the capacity–implementation distinction that is the central theoretical claim of this paper.

Principle 6.2 (Transformation Selection Principle). Cognitive computation selects trajectories that minimise the constraint energy $\mathcal{C}$ (defined in Appendix G) within the expressive envelope $\mathcal{F}_{\text{all}}$.

This principle ties together entropy weighting (which penalises high-S trajectories), the CPG prior (which restricts to low-dimensional generators), the admissibility constraints (which impose bounds on θ, S, and ξ), and the spectral efficiency bound (which limits how rapidly transformations can propagate). Each of these is a different facet of the same structural claim: cognition operates not by searching the full function class but by selecting among trajectories that can be coherently maintained under biological constraint.

7. Representation as Trajectory Equivalence

7.1. Against static encoding

The standard view of propositional representation treats semantic content as a property of a state: a pattern of activation, a point in a representational space, a configuration of features. This view is ill-suited to a transformation-first framework, because it makes the content of a representation independent of how the representation behaves under perturbation. Two states that activate the same features but respond differently to the same sequence of transformations would count as representing the same proposition under the standard view. That is the wrong result.

7.2. Trajectory equivalence

The alternative proposed here grounds propositional identity in transformation behaviour.

Definition 7.1 (Propositional equivalence). *Let $\Pi = (P_1, \dots, P_k)$ be a sequence of operator perturbations. Two states $x, x' \in M$ are* Π-equivalent, *written $x \sim_\Pi x'$, if*

$$\mathcal{A}^{(j)}(x) \;=\; \mathcal{A}^{(j)}(x') \quad \textit{for all } j \leq |\Pi|.$$

A proposition *is an equivalence class* $[x]_\Pi = \{x' \in M : x' \sim_\Pi x\}$.

This is a behavioural criterion for semantic content, not a representational one. Meaning is constituted by transformation invariance rather than by any intrinsic property of the state. The connection to inferential role semantics—specifically, Brandom's account of meaning as constituted by normative inferential commitment [20]—is immediate, but the present formulation provides a geometric implementation that the purely linguistic account lacks: Π-equivalence is in principle measurable, whereas inferential role has historically resisted quantification.

7.3. CPG dynamics as a trajectory basis

The trajectory-equivalence account avoids the overclaim that CPG chains directly encode arbitrary propositions. Instead, CPG-like dynamics provide the primitive trajectory basis from which propositional invariants can be constructed. A proposition is not mapped to a single motor pattern: it is represented by the equivalence class of all states that respond identically to the relevant sequence of cascade perturbations, where that sequence is grounded in, though not reducible to, the CPG-generated trajectory basis.

8. Predictions About Cognitive Dynamics

The two-layer model—expressive capacity plus biological prior—generates a set of strong, specific predictions.

Within a single transformation layer, as a participant learns the relevant rule, phase alignment θ_t should improve monotonically over trials: the participant's representational trajectory should come to point increasingly in the direction of the task gradient. Vorticity $\xi^{(N)}$ should decrease and stabilise within the layer, converging as the cascade becomes ordered. These within-layer predictions follow from the attractor stability theorem, which ensures vorticity convergence given small enough deformation intensities.

Across transformation layers, the cascade should compose in a lawful order. Participants should stabilise layer $\mathfrak{R}_1$ before $\mathfrak{R}_2$ becomes coherent, and $\mathfrak{R}_2$ before $\mathfrak{R}_3$. This ordering is a direct prediction of the biological prior: CPG-like dynamics engage one level of the hierarchy at a time, and entropy weighting suppresses premature engagement of higher layers before lower ones are stable.

The failure modes of the framework are as diagnostic as its successes. A participant who shortcuts—who moves directly from layer $\mathfrak{R}_1$ to the final output without passing through $\mathfrak{R}_2$—should show a phase discontinuity at layer 2: a sudden high-amplitude jump with poor alignment between $\mathbf{v}$ and $\nabla\Phi$ at that boundary. A participant confronted with a genuinely non-compositional transformation should either impose a spurious intermediate stage or show high phase error and boundary-localised vorticity. These signatures distinguish structured cascades from generic nonlinear learning, because a generic nonlinear model predicts smooth but arbitrarily shaped trajectories rather than vorticity specifically localised at decomposition boundaries.

9. Experimental Design: Compositional vs Non-Compositional Tasks

9.1. Controlled transformation domain

The experimental domain is designed to be deliberately artificial, low-dimensional, and over-instrumented. The goal at this stage is not ecological validity but falsifiability: the domain must be rich enough that participants can discover structure rather than memorise labels, while simple enough that the true transformation geometry can be computed in advance and used as ground truth. Linguistic morphology with nonce words satisfies both conditions.

A fixed feature manifold is defined prior to data collection, using phonological feature vectors from a standard inventory, morphological role vectors drawn from Universal Dependencies, and semantic-pragmatic features established through a norming study on the nonce items. The ground-truth Φ and $\nabla\Phi$ are computed from this manifold and held fixed throughout. This non-circularity is methodologically essential: if the embedding space were reconstructed from participant responses, the test of phase alignment would be vacuous.

9.2. Compositional task family

The first family presents transformations smoothly decomposable across three layers. Layer $\mathfrak{R}_1$ involves *phonological change* (feature scaling): a nonce stem undergoes a regular phonological modification, describable as a small-magnitude movement in phonological feature space. Layer $\mathfrak{R}_2$ involves *morphological role change* (rotation): the item is recategorised—from noun to verb, agent to instrument, or singular to plural—requiring a directional reorientation in role-feature space without necessarily changing magnitude. Layer $\mathfrak{R}_3$ involves *semantic-pragmatic shift* (entropy-weighted transformation): the item acquires a contextually modulated meaning such as a diminutive, augmentative, or metaphorical extension, where the correct transformation depends on contextual entropy and is accordingly suppressed by high uncertainty.

A representative cascade might run: *blick* → *blicken* → *blicker* → *tool-for-blicking*. The lexical items are irrelevant; what matters is that each layer has a known target direction in the pre-defined feature space and that participants' responses can be embedded and measured against that direction.

9.3. Non-compositional task family

The second family requires discontinuous jumps through semantic space. These tasks have correct answers that cannot be reached by any smooth decomposition into the three layers above. The correct transformation is opaque from the standpoint of the cascade structure: it requires moving to a point in the manifold with no stable intermediate under available dynamical primitives. The biological prior predicts that humans will either fail (showing high vorticity and poor alignment) or succeed by imposing a spurious intermediate stage not present in the task structure. A pure function approximator makes neither prediction: it should succeed or fail based solely on output accuracy, with no characteristic geometric signature.

9.4. Trajectory measurement

From participant responses across trials, the following quantities are computed at each time step t:

$$\mathbf{v}_t \;=\; x_{t+1} - x_t, \qquad \theta_t \;=\; \angle(\mathbf{v}_t,\, \nabla\Phi), \qquad \rho_t \;=\; \|\mathbf{v}_t\|, \tag{8}$$

where x_t is the embedding of the participant's response at trial t and $\nabla\Phi$ is the pre-defined task gradient. The triplet $(\mathbf{v}_t, \theta_t, \rho_t)$ constitutes the empirical estimate of the amplitwist operator at time t. Vorticity is estimated from the curl of the phase-weighted field across the trial sequence.

10. Human vs Model Comparison

10.1. Alignment loss

To compare operator geometry between human participants and language models, we define the *geometric alignment loss*. Let $\mathcal{A}^{(k)} = (\rho_k, \theta_k, S_k)$ denote the operator triplet at layer k. Then

$$\mathcal{L}_A \;=\; \sum_{k=1}^{N} \Big(\alpha_\rho \, (\rho_k^{(h)} - \rho_k^{(m)})^2 + \alpha_\theta \, d_{S^1}(\theta_k^{(h)},\, \theta_k^{(m)})^2 + \alpha_S \, (S_k^{(h)} - S_k^{(m)})^2 \Big), \tag{9}$$

where $d_{S^1}(\theta, \theta') = \min(|\theta - \theta'|, 2\pi - |\theta - \theta'|)$ is the angular geodesic distance on the circle, and $\alpha_\rho, \alpha_\theta, \alpha_S > 0$ are weighting coefficients. Using d_{S^1} rather than the naive Euclidean difference ensures that the metric is correct for the periodic phase variable and that the loss is genuinely a geodesic distance on the operator space $\mathbb{R}_{>0} \times S^1 \times \mathbb{R}_{\geq 0}$, not an ad hoc sum of squared differences. The critical comparison is not whether the model and the human produce the same final answer, but whether they traverse the transformation space through the same local rotation-scaling path. Endpoint equivalence with path divergence is the key empirical signature: if a model consistently matches human outputs while exhibiting disordered intermediate operator geometry, it has learned to produce the same results without learning the underlying transformation structure.

10.2. Predicted asymmetries

The framework predicts three distinct patterns. Humans on compositional tasks should show low vorticity, improving phase alignment, and ordered cascade composition, with each layer stabilising before the next is engaged. Humans on non-compositional tasks should show high vorticity or the characteristic signature of imposed spurious structure: a phase discontinuity followed by an intermediate state that reduces vorticity before the discontinuous jump is completed. Language models, given either task family, should show correct or near-correct outputs but weaker geometric coherence: higher vorticity at layer boundaries, less stable intermediate operator structure, and entropy insensitivity (because models lack anything analogous to the biological prior implemented by entropy weighting).

The asymmetry between human failure modes—directional, prior-shaped, localised at decomposition boundaries—and model failure modes—random, entropy-insensitive, uniformly distributed across layers—is the central empirical contrast. If this asymmetry is not observed, the biological prior interpretation is falsified independently of the operator formalism.

11. Interpretation: Geometry as the Invariant Substrate

If the predicted invariants are confirmed, the framework supports the following interpretation. Cognition, perception, and semantic processing are governed by shared transformation operators whose local structure is captured by the amplitwist formalism. Biological systems restrict the expressive envelope of these operators via dynamical priors that reflect the rhythmic, hierarchical, embodied structure of the nervous system. Meaning is not stored in states but emerges from the invariant structure of transformations: two states mean the same thing when they behave the same way under the same cascade of perturbations.

The same formalism describes local computation (individual operator applications), global coherence (cascade composition and eigenmode propagation), and cross-domain transfer (the alignment loss between human and model transformation geometry). This unified description is not merely analogical: it is grounded in the functional-analytic convergence of hybrid dynamics (Section 4) and the spectral geometry of the manifold (Section 5).

The framework also clarifies a connection to the geometry of conscious states without presupposing any particular answer to the hard problem. If conscious states are characterised by particularly stable, low-entropy, high-phase-alignment trajectories through the representational manifold—a claim consistent with integrated information theory and the entropic brain hypothesis [18, 19]—then the amplitwist formalism provides a language for describing the geometric structure that distinguishes such states from non-conscious processing.

12. Falsifiability and Outcomes

The theory is falsified by any of the following outcomes. If human trajectories on compositional tasks do not exhibit improving phase alignment and converging vorticity, the within-layer prediction fails. If compositional tasks do not produce ordered cascade composition—if participants show no tendency to stabilise lower layers before engaging higher ones—the across-layer prediction fails. If non-compositional tasks do not produce the characteristic signature of high vorticity or imposed spurious structure, the biological prior prediction fails. If the alignment loss metric does not discriminate human from model operator geometry despite output-level accuracy differences, the transformation-coherence hypothesis is not supported.

Partial failures are informative rather than merely negative. If within-layer alignment improves but across-layer ordering does not appear, it suggests the biological prior operates within layers but not across them—a finding that would call for revision of the CPG-hierarchy account. If the alignment loss discriminates humans from models on compositional tasks but not on non-compositional ones, it suggests the prior is doing discriminative work only where it has traction, which would refine rather than refute the framework.

The decisive null test is the following.

Null Hypothesis. If phase alignment θ_t and vorticity ξ_t show no systematic difference between (i) compositional and non-compositional tasks, and (ii) human and model trajectories, then the amplitwist framework does not capture a distinct computational invariant and the framework is falsified. This null can be tested without any commitment to the full RSVP field theory: it requires only that the discrete observables (θ_t, ξ_t) be computed from a pre-specified embedding and compared across conditions.

If the predictions are confirmed, the consequences are precise. Transformation coherence is a real, measurable dimension of cognitive processing, distinct from output accuracy. Biological systems operate as constrained operator cascades with a specific geometric signature. Current language models capture output statistics without capturing the operator geometry that generates human outputs. The amplitwist cascade is not an elegant reinterpretation of known ideas under a new vocabulary: it identifies structure in data that output-level models do not.

13. Conclusion

The amplitwist cascade is simultaneously a universal transformation language and a description of biological constraint. As a universal language, it encompasses the full expressive envelope of local rotation-scaling-projection operators under entropic regularisation. As a description of biological constraint, it formalises the prior that shapes cognition: the preference for CPG-supportable, rhythmically stable, compositionally ordered trajectories over the vast space of transformations the algebra could in principle compute.

The critical insight is the separation: capacity defines possibility; constraint defines reality. The gap between them is not an embarrassment to be minimised but the primary theoretical object—the structure that encodes the biological prior, generates the falsifiable predictions, and grounds the alignment loss metric that distinguishes human from model operator geometry.

The next step is empirical. The toy-domain experiment proposed in Section 9 is designed to be small enough that the true transformation geometry can be computed and used as ground truth, but rich enough that participants must discover structure rather than memorise labels. Running the experiment, computing phase alignment and vorticity trajectories, and comparing human and model operator geometry will determine whether transformation coherence is a real cognitive variable or an elegant fiction. The decisive outcome is not whether the framework is elegant, but whether it identifies structure in data that other models miss.

Acknowledgements. This essay was developed as part of the RSVP/Amplitwistor research programme.

Appendix A: Amplitwist Operator Algebra

A.1 Local Operator Definition

Let M be a smooth n-dimensional manifold and $x \in M$. Let T_xM denote the tangent space at x.

Definition A.1 (Amplitwist Operator). *An* amplitwist operator *at x is a triple*

$$\mathcal{A}_x = (\rho_x,\, R_x,\, \pi_x),$$

where $\rho_x \in \mathbb{R}_{>0}$ is a scalar gain (amplitude), $R_x \in SO(n)$ is a rotation acting on T_xM, and $\pi_x \colon T_xM \to T_xM$ is a (possibly nonlinear) projection map.

The action of $\mathcal{A}_x$ on a vector $u \in T_xM$ is

$$\mathcal{A}_x(u) = \pi_x(\rho_x\, R_x u).$$

A.2 Composition Law

Let $\mathcal{A}_1 = (\rho_1, R_1, \pi_1)$ and $\mathcal{A}_2 = (\rho_2, R_2, \pi_2)$. Define composition by

$$(\mathcal{A}_2 \circ \mathcal{A}_1)(u) = \pi_2\Big(\rho_2 R_2\, \pi_1(\rho_1 R_1 u)\Big).$$

Proposition A.2. *The set of amplitwist operators equipped with composition forms a nonlinear semigroup $(\mathcal{A}, \circ)$.*

Proof. Closure follows from closure of scalar multiplication, rotations, and function composition. Associativity follows from associativity of function composition. Nonlinearity arises from the presence of π_x. □

Remark A.3. In general $\mathcal{A}_2 \circ \mathcal{A}_1 \neq \mathcal{A}_1 \circ \mathcal{A}_2$, due to non-commutativity of both the rotation group ($R_2R_1 \neq R_1R_2$ for $n \geq 3$) and the projection maps ($\pi_2 \circ \pi_1 \neq \pi_1 \circ \pi_2$ in general).

A.3 Local Linearization

If π_x is differentiable, the Jacobian of $\mathcal{A}_x$ at u is

$$D\mathcal{A}_x(u) = D\pi_x(\rho_x R_x u) \cdot \rho_x R_x.$$

In the linear case $\pi_x = \mathrm{Id}$, this reduces to $D\mathcal{A}_x = \rho_x R_x$, a similarity transformation.

A.4 Density in Function Space

Let $K \subset \mathbb{R}^n$ be compact and let $C(K)$ denote the space of continuous functions equipped with the uniform norm.

Theorem A.4 (Expressive Density). *Suppose π_x contains a non-polynomial nonlinearity and amplitwist operators are composed in finite depth. Then the set of functions* *representable as* $f(x) = (\mathcal{A}_N \circ \cdots \circ \mathcal{A}_1)(x)$ *is dense in* $C(K)$*.*

Sketch. Each amplitwist layer can be written locally as $u \mapsto \pi(Wu + b)$ with $W = \rho R$. Since compositions of affine maps and non-polynomial nonlinearities are dense in $C(K)$, the result follows from standard universal approximation arguments. ☐

A.5 Cascades as Operator Sequences

Define an *amplitwist cascade* of depth N as $\mathcal{A}^{(N)} = \mathcal{A}_N \circ \cdots \circ \mathcal{A}_1$. Given an initial state u_0, the induced trajectory is $u_{k+1} = \mathcal{A}_{k+1}(u_k)$, defining a discrete dynamical system on T_xM.

A.6 Separation of Capacity and Implementation

Definition A.5 (Expressive Capacity)**.** $\mathcal{F}_{\text{all}} = \overline{\{\mathcal{A}_N \circ \cdots \circ \mathcal{A}_1\}}$*, the closure in* $C(K)$*.*

Definition A.6 (Biological Subspace)**.** $\mathcal{F}_{\text{bio}} \subset \mathcal{F}_{\text{all}}$ *denotes functions realisable under biological constraints (formalised in Appendix F).*

Proposition A.7. *In general,* $\mathcal{F}_{\text{bio}} \subsetneq \mathcal{F}_{\text{all}}$.

This strict inclusion formalises the separation between capacity (what amplitwist cascades can represent) and implementation (what biological systems actually compute).

Appendix B: Phase Alignment and Vorticity

B.1 Field Structure

Let M be a smooth manifold. Define a scalar field $\Phi\colon M \to \mathbb{R}$ (salience), a vector field $v\colon M \to TM$ (trajectory velocity), and an entropy field $S\colon M \to \mathbb{R}_{\geq 0}$ (uncertainty). Assume $v(x) \neq 0$ and $\nabla\Phi(x) \neq 0$ almost everywhere.

B.2 Phase Alignment

Definition B.1 (Phase Angle)**.** *The local phase alignment is*

$$\theta(x) = \arccos\left(\frac{v(x) \cdot \nabla\Phi(x)}{\|v(x)\|\,\|\nabla\Phi(x)\| + \epsilon}\right), \quad \epsilon > 0.$$

Proposition B.2. $\theta(x) = 0 \iff v(x) \parallel \nabla\Phi(x)$*, and* $\theta(x) = \pi \iff v(x) \parallel -\nabla\Phi(x)$*. Thus* $\theta(x)$ *measures directional coherence between the trajectory and the salience gradient.*

B.3 Alignment Functional

Define the global alignment energy $\mathcal{E}_{\text{align}} = \int_M (1 - \cos\theta(x))\, d\mu(x)$.

Theorem B.3. $\mathcal{E}_{\text{align}} = 0 \iff v(x) \parallel \nabla\Phi(x)$ *for all* $x \in M$*.*

Proof. Follows immediately from $1 - \cos\theta = 0 \iff \theta = 0$. ☐

B.4 Vorticity

Definition B.4 (Vorticity)**.** *Let* $\hat{v}(x) = v(x)/\|v(x)\|$ *be the normalised direction field. The vorticity is*

$$\xi(x) = \|\nabla \times \hat{v}(x)\|.$$

The integrated vorticity is $\Xi = \int_M \xi(x)\, d\mu(x)$, which equals zero if and only if $\hat{v}$ is irrotational.

B.5 Phase–Vorticity Relationship

Proposition B.5. *Large spatial gradients of* θ *imply increased vorticity: if* $\|\nabla\theta(x)\| \gg 0$ *then* $\xi(x)$ *is large.*

Sketch. Since $\hat{v}$ depends on the angular variable, spatial variation in phase induces curl in the direction field. ☐

B.6 Discrete Approximation

Given a discrete response trajectory $\{x_t\}$, define $v_t = x_{t+1} - x_t$ and $\hat{v}_t = v_t/(\|v_t\| + \epsilon)$. The discrete phase estimate is

$$\theta_t = \arccos\left(\frac{v_t \cdot \nabla\Phi(x_t)}{\|v_t\|\,\|\nabla\Phi(x_t)\| + \epsilon}\right),$$

and the discrete vorticity approximation is $\xi_t \approx \|\hat{v}_{t+1} - \hat{v}_t\|$.

B.7 Coherence Criterion

A trajectory is *coherent* if $\sup_t \theta_t \leq \theta_{\text{crit}}$ and $\sum_t \xi_t \leq \Xi_{\text{crit}}$. Low θ signals aligned transformation; low ξ signals stable cascade dynamics. High θ signals misalignment; high ξ signals rotational instability. These quantities define measurable invariants of transformation coherence that can be estimated from behavioural response sequences.

Appendix C: Entropy Weighting as a Biological Prior

C.1 Entropy Field and Weighting

Let $S\colon M \to \mathbb{R}_{\geq 0}$ denote an entropy (uncertainty) field.

Definition C.1 (Entropy Weight and Weighted Operator)**.** *For layer k, define* $w_k(x) = e^{-\lambda S(x)}$ *with* $\lambda > 0$. *The entropy-weighted amplitwist is*

$$\widetilde{\mathcal{A}}_x^{(k)}(u) = w_k(x)\, \pi_x\big(\rho_x R_x u\big).$$

High entropy suppresses operator influence; low entropy amplifies stable transformations.

C.2 Variational Formulation

Consider $f\colon M \to \mathbb{R}^m$. Define the weighted Dirichlet energy

$$\mathcal{L}[f] = \int_M e^{-\lambda S(x)} \, \|\nabla f(x)\|^2 \, d\mu(x).$$

Proposition C.2. *Critical points of $\mathcal{L}$ satisfy the weighted Euler–Lagrange equation $\nabla \cdot (e^{-\lambda S} \nabla f) = 0$.*

Proof. Computing the first variation and integrating by parts: $\delta\mathcal{L} = -2 \int_M \delta f \, \nabla \cdot (e^{-\lambda S} \nabla f) \, d\mu = 0$ for all admissible δf. $\square$

Regions with large $S(x)$ are penalised (exploration suppressed); regions with small $S(x)$ are favoured (stable exploitation). Thus w_k encodes a biological prior toward low-uncertainty trajectories.

C.3 Admissibility and Stability

Definition C.3 (Admissible State and Trajectory). *A point $x \in M$ is admissible if $S(x) \leq S_{\max}$. A trajectory $\gamma\colon [0, T] \to M$ is admissible if* $\sup_t S(\gamma(t)) \leq S_{\max}$.

Proposition C.4. *If $S(x) \leq S_{\max}$ and $\lambda > 0$, then $e^{-\lambda S_{\max}} \leq w_k(x) \leq 1$.*

C.4 Entropy–Vorticity Coupling

Under the assumption that high-entropy regions induce higher phase variability ($\|\nabla\theta(x)\| \propto S(x)$), Appendix B implies that $S(x) \uparrow$ entails $\xi(x) \uparrow$, since vorticity scales with the spatial gradient of direction.

C.5 Regularised Energy Functional

The combined functional

$$\mathcal{J}[f] = \int_M \left[e^{-\lambda S(x)} \|\nabla f(x)\|^2 + \alpha S(x)^2 \right] d\mu(x)$$

enforces coherent transformations through its first term and penalises entropy growth through its second. Admissible dynamics minimise $\mathcal{J}$ subject to task constraints. The limiting cases are informative: as $\lambda \to 0$, the entropy weight vanishes and the full expressive capacity of $\mathcal{F}_{\text{all}}$ is recovered; as $\lambda \to \infty$, only minimal-entropy trajectories survive, corresponding to the most strongly constrained biological regime.

Appendix D: Hybrid Dynamics and Trotter–Kato Convergence

D.1 Semigroup Evolution

Let X be a Banach space (e.g. $L^2(M)$ or $H^k(M)$) and let $F\colon \mathcal{D}(F) \subset X \to X$ generate a well-posed flow.

Definition D.1 (Strongly Continuous Semigroup). $\{T(t)\}_{t \geq 0}$ *satisfies* $T(0) = I$,

$T(t+s) = T(t)T(s)$, and $\lim_{t\downarrow 0} \|T(t)u - u\| = 0$ for all $u \in X$. For linear $F = A$ (densely defined and closed), $T(t) = e^{tA}$.

D.2 Discrete Operator Updates

Let $\{P_k\}$ be a sequence of (possibly nonlinear) amplitwist operators acting on X. Given a time grid $0 = t_0 < t_1 < \cdots < t_N = t$ with $\Delta t_k = t_{k+1} - t_k$, the hybrid evolution is

$$u_{k+1} = P_k \, T(\Delta t_k) \, u_k, \qquad u(t) = \Big(\prod_{k=0}^{N-1} P_k T(\Delta t_k) \Big) u_0,$$

with time-ordered product.

D.3 Linear Trotter–Kato Framework

Theorem D.2 (Trotter Product Formula)**.** *If A and B generate strongly continuous semigroups $T_A(t)$ and $T_B(t)$ on X, then under standard domain conditions,*

$$\lim_{n\to\infty} \Big(T_A(t/n) \, T_B(t/n) \Big)^n = T_{A+B}(t)$$

strongly on X.

Identifying $T_A(\Delta t) \equiv T(\Delta t)$ and $T_B(\Delta t) \equiv P(\Delta t)$, the hybrid scheme approximates $\exp(t(A+B))$ in the linearised regime.

D.4 Nonlinear Extension (Lie–Trotter Splitting)

For nonlinear evolution $u' = F(u) + G(u)$, define flows $\Phi_F(t)$ and $\Phi_G(t)$.

Theorem D.3 (Lie–Trotter Splitting)**.** *Under Lipschitz conditions on F and G, the scheme $u^{(n)}(t) = (\Phi_G(t/n) \circ \Phi_F(t/n))^n u_0$ converges to the solution of $u' = F(u) + G(u)$ as $n \to \infty$.*

The identification is $\Phi_F \leftrightarrow T(t)$ and $\Phi_G \leftrightarrow P(t)$.

D.5 Convergence Result

Theorem D.4 (Hybrid Convergence)**.** *Assume $T(t)$ is strongly continuous, $P_k \to P_{t_k}$ uniformly, the scheme is stable (in the sense that $\sup_n \| \prod_{k=0}^{n-1} P_k T(\Delta t_k) \| \le C$), and the local truncation error is $O(\Delta t^2)$. Then*

$$\lim_{\max \Delta t_k \to 0} \Big(\prod_{k=0}^{N-1} P_k \, T(\Delta t_k) \Big) = T(t) \circ P_t$$

strongly on X.

The consequence for cascades is direct: discrete amplitwist cascades are not arbitrary updates but approximate a continuous operator flow, and in the limit they inherit the same invariant structure as the underlying continuous field dynamics.

Appendix E: Spectral Geometry and Efficiency Bounds

E.1 Laplacian and Eigenstructure

Let M be a compact Riemannian manifold with metric g. The Laplace–Beltrami operator $\Delta_M f = \nabla \cdot \nabla f$ has eigenpairs (λ_k, ψ_k) satisfying $\Delta_M \psi_k = -\lambda_k \psi_k$ with $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \cdots$. Any $f \in L^2(M)$ admits the spectral expansion $f(x) = \sum_k a_k \psi_k(x)$, and semigroup evolution satisfies $T(t)f = \sum_k a_k e^{-\lambda_k t} \psi_k$.

E.2 Efficiency Metric and Bound

Definition E.1 (Cascade Efficiency). $\eta^{(N)} = \|f^{(N)} - f^{(0)}\|_{L^2} / \sum_{j=1}^{N} \|\epsilon_j T_j\|_{L^2 \to L^2}$.

Theorem E.2 (Spectral Efficiency Bound). *Let M be a compact Riemannian manifold and let $\{\psi_k\}$ be the orthonormal eigenbasis of Δ_M. Let the cascade act on $f^{(0)} \in L^2(M)$ with $f^{(0)} \perp \psi_0$, and assume each layer admits the first-order expansion*

$$\mathcal{A}_j = I + \epsilon_j T_j + O(\epsilon_j^2),$$

where T_j is a bounded operator with $\|T_j\|_{L^2 \to L^2} \leq C$. Then

$$\|f^{(N)} - f^{(0)}\|_{L^2} \;\geq\; \lambda_1(M) \left\| \sum_{j=1}^{N} \epsilon_j T_j f^{(0)} \right\|_{L^2} - O\Big(\sum_j \epsilon_j^2\Big),$$

and consequently

$$\eta^{(N)} \;\gtrsim\; \frac{\lambda_1(M)}{N \cdot \max_j \|\epsilon_j T_j\|_{L^2 \to L^2}}.$$

Sketch. Since $f^{(0)} \perp \psi_0$, the Rayleigh quotient gives $\|f^{(0)}\|_{L^2}^2 \leq \lambda_1(M)^{-1} \|\nabla f^{(0)}\|_{L^2}^2$. Expanding the cascade to first order in ϵ_j and applying the Rayleigh quotient to the accumulated deformation $\sum_j \epsilon_j T_j f^{(0)}$ yields the lower bound on $\|f^{(N)} - f^{(0)}\|_{L^2}$. The second-order remainder is $O(\sum_j \epsilon_j^2)$ [9, 10]. □

The bound arises because any nontrivial deformation must excite modes orthogonal to the constant eigenfunction, and the Rayleigh quotient enforces a minimum energy cost proportional to $\lambda_1(M)$. The spectral gap therefore sets a geometric lower limit on how efficiently transformations propagate across M [11].

E.3 Geometric Interpretation

Large λ_1 indicates strong connectivity and rapid propagation; small λ_1 indicates geometric resistance. The Cheeger constant $h(M)$ satisfies $h(M)^2/4 \leq \lambda_1(M) \leq 2h(M)$, so geometric bottlenecks directly limit cascade efficiency. Low-frequency modes (small λ_k) correspond to long-range integration and slow dynamics; high-frequency modes correspond to local refinement and transient responses.

E.4 Spectral Filtering and the Biological Subspace

Entropy-weighted amplitwists preferentially damp high-frequency modes: the map $a_k \mapsto a_k e^{-\lambda S}$ suppresses higher-λ_k coefficients more strongly, favouring smooth, low-energy representations. The biologically realisable function space can accordingly be characterised spectrally as

$$\mathcal{F}_{\text{bio}} = \Big\{ f \in \mathcal{F}_{\text{all}} \;\Big|\; \sum_k \lambda_k a_k^2 \le C \Big\}$$

for some bound C, defining a spectrally constrained subset of the full function space in which biological systems operate within low-energy spectral bands.

Appendix F: Dynamical Constraint Model (CPG Prior)

F.1 Trajectory Spaces and Biological Realisation

Let $\mathcal{T}_{\text{all}} = \{\gamma \colon [0,T] \to M \mid \gamma \text{ absolutely continuous}\}$.

Definition F.1 (Biological Trajectory)**.** *A trajectory $\gamma \in \mathcal{T}_{\text{bio}} \subset \mathcal{T}_{\text{all}}$ if it is generated by a finite composition of low-dimensional dynamical primitives $\gamma(t) = \Phi_k \circ \cdots \circ \Phi_1(t)$, where each Φ_i is generated by a smooth, bounded vector field $\dot{x} = f_i(x)$ with $\|\nabla f_i\| \le C$. A primitive is* CPG-like *if $f_i(x) = A_i x + g_i(x)$, where A_i has imaginary eigenvalues (oscillatory structure) and g_i is a small nonlinear perturbation.*

F.2 Constraint Functional

Define the trajectory energy

$$\mathcal{J}[\gamma] = \int_0^T \Big(\|\ddot{\gamma}(t)\|^2 + \alpha\, \xi(\gamma(t)) + \beta\, S(\gamma(t)) \Big)\, dt.$$

This functional simultaneously penalises rapid acceleration (lack of smoothness), vorticity (rotational instability), and entropy (uncertainty).

Theorem F.2. *Biological trajectories approximately satisfy $\gamma \in \arg\min_{\tilde{\gamma} \in \mathcal{T}_{\text{all}}} \mathcal{J}[\tilde{\gamma}]$, with Euler–Lagrange equation $\frac{d^2}{dt^2}\ddot{\gamma}(t) - \alpha \nabla \xi(\gamma(t)) - \beta \nabla S(\gamma(t)) = 0$.*

F.3 Low-Dimensional Manifold Constraint

Proposition F.3. *There exists a manifold $\mathcal{M}_{\text{bio}} \subset \mathcal{T}_{\text{all}}$ of dimension $d \ll \dim(\mathcal{T}_{\text{all}})$ such that $\mathcal{T}_{\text{bio}} \subset \mathcal{M}_{\text{bio}}$.*

Sketch. Follows from restriction to a finite set of dynamical primitives and smooth compositions. □

Although amplitwist cascades are dense in $C(K)$ by Theorem A.4, biological trajectories realise only the restricted subset $\mathcal{F}_{\text{bio}} = \{f \in \mathcal{F}_{\text{all}} \mid f \text{ induced by } \gamma \in \mathcal{T}_{\text{bio}}\}$. A transformation f is *non-dynamical* if no $\gamma \in \mathcal{T}_{\text{bio}}$ generates it with bounded $\mathcal{J}$; such transformations produce high vorticity, high entropy, and discontinuities in θ.

The measure-theoretic form of this restriction is given by the following proposition, which is the formal statement of the paper's central thesis.

Definition F.4 (Entropy-Induced Path Measure)**.** *Let $\mathcal{T}_{\text{all}}$ be the space of admissible trajectories $\gamma\colon [0, T] \to M$. Define the Gibbs measure over paths*

$$d\mu_{\text{bio}}(\gamma) = \frac{1}{Z} \exp\Big(-\lambda \int_0^T S(\gamma(t))\, dt\Big)\, d\mu_0(\gamma),$$

where μ_0 is a reference measure on trajectory space and Z is the normalising partition function. The induced measure on function space is $\Phi_* \mu_{\text{bio}}$, where $\Phi\colon \mathcal{T}_{\text{all}} \to \mathcal{F}_{\text{all}}$ maps trajectories to their induced transformations.

Proposition F.5 (Entropy-Induced Concentration)**.** *Assume the entropy functional S is coercive in the sense that trajectories leaving the finite-dimensional submanifold $\mathcal{M}_{\text{bio}} \subset \mathcal{T}_{\text{all}}$ incur entropy growth $S(\gamma(t)) \geq c \operatorname{dist}(\gamma(t), \mathcal{M}_{\text{bio}})^2$ for some $c > 0$. Then for any $\epsilon > 0$, there exists a tubular neighbourhood $\mathcal{U}_\epsilon$ of $\mathcal{M}_{\text{bio}}$ such that*

$$\mu_{\text{bio}}(\mathcal{U}_\epsilon) \geq 1 - \epsilon.$$

Consequently, the induced function class $\mathcal{F}_\epsilon = \Phi(\mathcal{U}_\epsilon)$ has effective dimension bounded by $\dim(\mathcal{M}_{\text{bio}})$.

Sketch. Under the coercivity assumption, trajectories at distance r from $\mathcal{M}_{\text{bio}}$ carry entropy weight at most $e^{-\lambda c r^2 T}$. Standard large-deviation estimates then show that $\mu_{\text{bio}}(\mathcal{T}_{\text{all}} \setminus \mathcal{U}_\epsilon) \leq e^{-\lambda c \operatorname{dist}(\mathcal{U}_\epsilon, \mathcal{M}_{\text{bio}})^2 T} / Z$, which vanishes as $\epsilon \to 0$ since $Z > 0$ and the reference measure is finite. This is a standard Gibbs concentration result: the measure concentrates exponentially around the low-entropy manifold. □

The result follows the standard large-deviation phenomenon: the entropy weighting induces a Gibbs measure that concentrates exponentially around low-entropy trajectories. Since these trajectories lie on a finite-dimensional manifold generated by the dynamical primitives, biological computation occupies a sharply restricted subset of the full function space despite its theoretical expressivity.

Appendix G: Admissibility Constraints and the Revised Categorical Principle

G.1 Motivation

The biological prior of Appendix F constrains which trajectories occur. A complementary question concerns which *transformations* are structurally valid: not what a system does in practice, but what it is permissible for it to do while remaining composable with other systems operating under the same constraints. The answer takes the form of a bounded region in the transformation space defined by the operator geometry already in place, rather than a prescription over outputs. A transformation is admissible if and only if it remains within the region where phase alignment, entropy, and vorticity are all bounded—precisely the conditions under which operators can be coherently composed. This is the sense in which the framework yields something analogous to a categorical constraint, but one that is grounded in geometry rather than derived from abstract ethical reasoning.

G.2 Admissible Transformation Set

Definition G.1 (Admissible Operator and Admissible Set). *An operator* $\mathcal{A}_x$ *is* admissible *at* $x \in M$ *if*

$$\theta(x) \leq \theta_{\max}, \quad S(x) \leq S_{\max}, \quad \xi(x) \leq \xi_{\max}.$$

The admissible set is $\mathcal{A}_{\text{adm}} = \{\mathcal{A}_x \in \mathcal{A} \mid \theta(x) \leq \theta_{\max},\ S(x) \leq S_{\max},\ \xi(x) \leq \xi_{\max}\}$.

G.3 Constraint Functional

Define the global constraint energy

$$\mathcal{C}[\gamma] = \int_M \left[\alpha\, \theta(x)^2 + \beta\, S(x)^2 + \gamma\, \xi(x)^2\right] d\mu(x).$$

A trajectory γ is admissible if $\mathcal{C}[\gamma] \leq C_{\max}$.

G.4 The Revised Categorical Principle

Principle G.2 (Admissibility Constraint)**.** A transformation is admissible if and only if it remains within bounded regions of phase alignment, entropy, and vorticity such that it preserves composability with other admissible transformations.

This principle differs from classical categorical formulations in two ways. First, it operates at the level of transformations, not outputs: it does not prescribe what function to compute, only which transformations are structurally valid. Second, it implies both a minimum and a maximum. The lower bound—$\theta(x) \geq 0$, ensuring non-degenerate transformation—prevents degenerate (null) operators that contribute no information to the cascade. The upper bounds—on θ, S, and ξ—ensure stability and composability. An operator that violates the upper bounds produces non-composable transformations, cascade breakdown, and unstable trajectories.

G.5 Compositional Closure

Theorem G.3. *If* $\mathcal{A}_1, \mathcal{A}_2 \in \mathcal{A}_{\text{adm}}$ *and the constraint functional* $\mathcal{C}$ *is subadditive over compositions, then* $\mathcal{A}_2 \circ \mathcal{A}_1 \in \mathcal{A}_{\text{adm}}$.

Sketch. Composition preserves the admissibility bounds when $\mathcal{C}$ satisfies the subadditivity condition $\mathcal{C}[\mathcal{A}_2 \circ \mathcal{A}_1] \leq \mathcal{C}[\mathcal{A}_1] + \mathcal{C}[\mathcal{A}_2]$. □

G.6 Relation to Biological Prior and Spectral Structure

The admissibility set $\mathcal{T}_{\text{bio}}$ from Appendix F can be recovered as $\{\gamma \mid \mathcal{C}[\gamma] \leq C_{\max}\}$, confirming that biological systems operate within admissible bounds rather than maximising expressivity. The admissibility constraint also implies a spectral bound: $\sum_k \lambda_k a_k^2 \leq E_{\max}$, consistent with the spectral characterisation of $\mathcal{F}_{\text{bio}}$ in Appendix E.

G.7 Alignment Without Identical Outputs

The admissibility framing resolves a persistent ambiguity in alignment discussions. Two agents are aligned not because they produce identical outputs, but because their transformations lie within the same admissible region and can be composed without breaking coherence. Alignment is a property of transformation geometry, not output agreement. This follows directly from the trajectory-equivalence account of propositional identity (Appendix H): two agents that produce different surface responses may nonetheless be Π-equivalent if their operator geometry is compatible under the same perturbation sequence.

The Transformation Selection Principle from Section 6 can be restated in this appendix's language.

Principle G.4 (Transformation Selection—Formal Statement)**.** Cognitive computation selects trajectories that minimise the constraint energy

$$\mathcal{C}[\gamma] = \int_M \left[\alpha\, \theta(x)^2 + \beta\, S(x)^2 + \gamma\, \xi(x)^2\right] d\mu(x)$$

within the expressive envelope $\mathcal{F}_{\text{all}}$, subject to task constraints.

This formulation is not a prescription over outputs. It is a constraint over the geometry of admissible computation: cognitive systems do not compute any function they could in principle compute, but the specific sub-class of functions reachable by trajectories that remain within $\mathcal{A}_{\text{adm}}$.

Appendix H: Representation as Trajectory Equivalence

H.1 Equivalence Relation

Let $\mathcal{X} = \{x \in M\}$ and let $\mathcal{P} = \{P_k\}$ be a sequence of perturbations. Define the perturbed cascade $x_{k+1} = \mathcal{A}_k(P_k(x_k))$.

Definition H.1 (Trajectory Equivalence)**.** *Two states $x, y \in M$ are equivalent, written $x \sim y$, if for all admissible perturbation sequences $\mathcal{P}$,*

$$\mathcal{A}_{\mathcal{P}}^{(N)}(x) = \mathcal{A}_{\mathcal{P}}^{(N)}(y) \quad \textit{for all } N.$$

H.2 Quotient Space and Propositions

The quotient $\mathcal{R} = \mathcal{X}/\sim$ represents propositions: an element $[x] \in \mathcal{R}$ is a proposition. A property $F\colon \mathcal{X} \to \mathbb{R}$ is representationally invariant if and only if $x \sim y$ implies $F(x) = F(y)$.

H.3 Stability Under Composition

Theorem H.2. *If $x \sim y$, then $\mathcal{A}^{(N)}(x) \sim \mathcal{A}^{(N)}(y)$ for any cascade.*

Proof. Closure follows from associativity of operator composition. □

H.4 Metric Structure

Define the pseudometric $d(x, y) = \sup_{\mathcal{P},N} \|\mathcal{A}_{\mathcal{P}}^{(N)}(x) - \mathcal{A}_{\mathcal{P}}^{(N)}(y)\|$. Then $x \sim y \iff d(x, y) = 0$. Restricting to biologically admissible trajectories yields $\sim_{\text{bio}} \subseteq \sim$: biological systems identify propositions only up to admissible transformations.

Propositions under this account are not static encodings. They are invariants under transformation cascades, and meaning is defined by stability under perturbation rather than by any intrinsic property of the state.

Appendix I: Empirical Observables and Measurement Protocol

I.1 Fixed Representation Manifold

Let $M \subset \mathbb{R}^d$ be a pre-specified feature manifold with metric g, defined prior to data collection (non-circular design). All embeddings $x_t \in M$ are constructed from this manifold; no parameters of M are fitted to participant trajectories.

I.2 Discrete Trajectory and Phase

Given responses $\{x_t\}_{t=0}^{T} \subset M$, define $v_t = x_{t+1} - x_t$, $\hat{v}_t = v_t/(\|v_t\| + \epsilon)$, and the discrete phase and amplitude

$$\theta_t = \arccos\left(\frac{v_t \cdot \nabla\Phi(x_t)}{\|v_t\|\,\|\nabla\Phi(x_t)\| + \epsilon}\right), \qquad \rho_t = \|v_t\|,$$

where $\Phi(x) = -d_M(x, x^\star)$ for the task target $x^\star$.

I.3 Entropy Proxies

The entropy field S_t may be estimated by several proxies: the Shannon entropy of the response choice distribution ($-\sum_i p_{t,i} \log p_{t,i}$), the surprisal of the observed transition ($-\log p(x_{t+1} \mid x_t)$), or a monotone function of reported confidence. Any of these provide an operationalisation of S that can be inserted into the entropy-weighted operator.

I.4 Vorticity, Aggregates, and Layer Segmentation

Discrete vorticity is approximated as $\xi_t \approx \|\hat{v}_{t+1} - \hat{v}_t\|$. Global aggregates are $\bar{\theta} = T^{-1} \sum_t \theta_t$ and $\Xi = \sum_t \xi_t$. Partitioning indices into layers $\mathcal{I}_\ell$ (phonological, morphological, semantic-pragmatic), layer-wise quantities are $\bar{\theta}^{(\ell)} = |\mathcal{I}_\ell|^{-1} \sum_{t \in \mathcal{I}_\ell} \theta_t$ and $\Xi^{(\ell)} = \sum_{t \in \mathcal{I}_\ell} \xi_t$.

A layer ℓ is *convergent* if $\theta_{t+1} \leq \theta_t + \delta$ for most $t \in \mathcal{I}_\ell$ (small tolerance δ). At a layer boundary b, the boundary vorticity $\xi_b = \|\hat{v}_{b+1} - \hat{v}_b\|$ measures the discontinuity between cascade stages; large ξ_b indicates a phase jump rather than smooth composition.

I.5 Human–Model Alignment Loss

Let $\{x_t^{(h)}\}$ and $\{x_t^{(m)}\}$ be human and model trajectories embedded in the same M.

Definition I.1 (Geometric Alignment Loss). *Let $\mathcal{A}^{(k)} = (\rho_k, \theta_k, S_k)$. Define the layer-k operator distance*

$$d_A^{(k)} = \sqrt{\alpha_\rho (\rho_k^{(h)} - \rho_k^{(m)})^2 + \alpha_\theta \, d_{S^1}(\theta_k^{(h)}, \theta_k^{(m)})^2 + \alpha_S (S_k^{(h)} - S_k^{(m)})^2},$$

where $d_{S^1}(\theta, \theta') = \min(|\theta - \theta'|, 2\pi - |\theta - \theta'|)$ is the angular geodesic distance on S^1, and $\alpha_\rho, \alpha_\theta, \alpha_S > 0$ are weighting coefficients. The alignment loss is $\mathcal{L}_A = \sum_k d_A^{(k)}$.

Using d_{S^1} corrects the periodicity of the phase variable and establishes $\mathcal{L}_A$ as a proper geodesic distance on the product space $\mathbb{R}_{>0} \times S^1 \times \mathbb{R}_{\geq 0}$.

Since M and Φ are fixed prior to data collection, the estimates θ_t and ξ_t are independent of trajectory inference, ensuring that the test of phase alignment is non-circular. The primary experimental outputs are the phase trajectory $\{\theta_t\}$, the vorticity profile $\{\xi_t\}$, the layer-wise convergence $\bar{\theta}^{(\ell)}$, boundary discontinuities ξ_b, and the alignment loss $\mathcal{L}_A$. Together these define a fully operational pipeline mapping behavioural responses to the tuple $(v_t, \theta_t, \rho_t, S_t, \xi_t)$ that constitutes the empirical amplitwist estimate.

Appendix J: Non-Compositional Transformations and Failure Modes

J.1 Compositional and Non-Compositional Transformations

Definition J.1 (Compositional Transformation). *A transformation $f\colon M \to M$ is compositional if there exists a finite admissible cascade $f \approx \mathcal{A}_N \circ \cdots \circ \mathcal{A}_1$ with each $\mathcal{A}_k \in \mathcal{A}_{\mathrm{adm}}$.*

Definition J.2 (Non-Compositional Transformation). *f is non-compositional if no finite admissible cascade approximates it with bounded $\theta \leq \theta_{\max}$, $\xi \leq \xi_{\max}$, and $S \leq S_{\max}$.*

Proposition J.3. *If f is non-compositional, then any approximating trajectory γ must contain a time t at which $\theta_t > \theta_{\max}$ or $\xi_t > \xi_{\max}$.*

J.2 Failure Signatures

At a layer boundary b, the boundary jump is $\Delta v_b = \|\hat{v}_{b+1} - \hat{v}_b\|$. Non-compositional tasks produce $\Delta v_b \gg 0$. Biological systems minimising $\mathcal{J}$ may respond by introducing spurious decompositions: a trajectory $\gamma = \gamma_3 \circ \gamma_2 \circ \gamma_1$ where no true compositional structure exists in f, reducing ξ and θ locally at the cost of traversing an intermediate state that is not part of the task.

J.3 Three-Regime Prediction

Theorem J.4 (Behavioural Regimes). *Let tasks be partitioned into compositional and non-compositional families. Then three distinct patterns are predicted. On compositional tasks, human trajectories show $\theta_t \downarrow$ (improving phase alignment), $\xi_t \downarrow$ (stabilising vorticity), and an ordered cascade in which lower layers stabilise before higher ones engage. On non-compositional tasks, human trajectories show $\theta_t \uparrow$ or $\xi_t \uparrow$, or exhibit spurious intermediate stages. Language models, on either task family, exhibit low final output error alongside higher vorticity at layer boundaries and disordered intermediate operator structure.*

J.4 Output vs Path Dissociation

Definition J.5 (Endpoint Accuracy and Path Coherence Error). *Let $\epsilon_{\text{out}} = \|f(x) - \hat{f}(x)\|$ and $\epsilon_{\text{path}} = \sum_t (\theta_t^2 + \xi_t^2)$.*

Proposition J.6. *It is possible that $\epsilon_{\text{out}} \approx 0$ while $\epsilon_{\text{path}} \gg 0$.*

This distinguishes correct outputs from coherent transformations. Non-compositional transformations also excite high-frequency spectral modes (a_k large for high λ_k), simultaneously increasing entropy and vorticity.

A task is identified as non-compositional empirically if the observed data satisfy $\Xi > \Xi_{\max}$ and $\bar{\theta} > \theta_{\max}$ despite low output error. This constitutes the falsifiable signature of the framework: the predicted asymmetry between human failure modes (directional, prior-shaped, localised at decomposition boundaries) and model failure modes (random, entropy-insensitive, uniformly distributed across layers) is either present in the data or it is not.

Appendix K: Gibbs Unification: Measure, Energy, and Admissibility

K.1 Motivation

The constraint functional $\mathcal{C}$, the entropy-weighted measure μ_{bio}, and the admissibility region $\mathcal{A}_{\text{adm}}$ are three descriptions of the same object in different languages. This appendix makes their identification explicit, showing that admissible trajectories are precisely the typical set of a Gibbs distribution whose energy functional is $\mathcal{C}$.

K.2 Total Constraint Energy

Definition K.1 (Total Constraint Energy). *Let $\gamma \colon [0, T] \to M$. Define*

$$\mathcal{C}[\gamma] = \int_0^T \left(\alpha\, \theta(\gamma(t))^2 + \beta\, S(\gamma(t)) + \delta\, \xi(\gamma(t))^2 \right) dt.$$

Definition K.2 (Gibbs Measure on Trajectories).

$$d\mu_{\text{bio}}(\gamma) = \frac{1}{Z} \exp(-\lambda\, \mathcal{C}[\gamma])\, d\mu_0(\gamma),$$

where μ_0 *is a reference measure and* $Z = \int \exp(-\lambda\mathcal{C})\, d\mu_0$ *is the partition function.*

The entropy weighting of Appendix C and the constraint functional of Appendix G are now unified: μ_{bio} is the Gibbs measure induced by $\mathcal{C}$, and the exponential weight $w_k = e^{-\lambda S}$ is the single-time-step restriction of the full path measure.

K.3 Admissibility as a Typical Set

Definition K.3 (Admissible Trajectory Set). *For* $\epsilon > 0$*, define*

$$\mathcal{T}^{\epsilon}_{\text{adm}} = \{\gamma \in \mathcal{T}_{\text{all}} \mid \mathcal{C}[\gamma] \leq C_{\min} + \epsilon\}, \quad C_{\min} = \inf_{\gamma} \mathcal{C}[\gamma].$$

Theorem K.4 (Typical-Set Characterisation of Admissibility). *Let* μ_{bio} *be as in Definition K.2. Then for any* $\delta > 0$ *there exists* $\epsilon > 0$ *such that*

$$\mu_{\text{bio}}(\mathcal{T}^{\epsilon}_{\text{adm}}) \geq 1 - \delta.$$

Sketch. Since $\mu_{\text{bio}}(\gamma) \propto e^{-\lambda\mathcal{C}[\gamma]}$, trajectories with energy exceeding $C_{\min} + \epsilon$ are suppressed by a factor of $e^{-\lambda\epsilon}$. By standard large-deviation estimates, the measure concentrates on a thin shell around the minimisers of $\mathcal{C}$, yielding the claim. □

Corollary K.5 (Variational Form of the Biological Prior). Biological trajectories satisfy $\gamma \in \arg\min_{\tilde{\gamma}} \mathcal{C}[\tilde{\gamma}]$ almost surely under μ_{bio}.

In this formulation, the biological prior is not an external constraint imposed on computation but the statistical consequence of a Gibbs distribution over trajectories. Admissible trajectories are precisely those lying in the typical set of this measure. Cognitive computation can therefore be understood as sampling from a low-temperature Gibbs measure over transformation trajectories, where $\mathcal{C}$ defines the geometry of admissible thought.

K.4 Pushforward to Function Space

Proposition K.6 (Concentration in Function Space). *Let* $\Phi\colon \mathcal{T}_{\text{all}} \to \mathcal{F}_{\text{all}}$ *map trajectories to induced transformations. Then the pushforward measure* $\Phi_*\mu_{\text{bio}}$ *concentrates on a subset* $\mathcal{F}_{\text{bio}} \subset \mathcal{F}_{\text{all}}$ *of effective dimension bounded by the dimension of the minimiser set of* $\mathcal{C}$.

Thus the restriction of biological computation to a low-dimensional function class is a direct consequence of measure concentration in trajectory space.

Appendix L: Variational Alignment and the Learning Objective

L.1 Human and Model Gibbs Measures

Let $\mathcal{C}_h$ and $\mathcal{C}_m$ denote the constraint functionals for human and model systems respectively.

Definition L.1 (Induced Trajectory Measures)**.**

$$d\mu_h(\gamma) = \frac{1}{Z_h} e^{-\lambda \mathcal{C}_h[\gamma]}\, d\mu_0(\gamma), \qquad d\mu_m(\gamma) = \frac{1}{Z_m} e^{-\lambda \mathcal{C}_m[\gamma]}\, d\mu_0(\gamma).$$

Humans and models are treated as different Gibbs samplers over the same trajectory space, distinguished by their effective constraint geometry.

L.2 KL Divergence as the Canonical Alignment Quantity

Definition L.2 (Geometric Alignment Divergence)**.**

$$D_{\mathrm{KL}}(\mu_h \,\|\, \mu_m) = \int \log\left(\frac{d\mu_h}{d\mu_m}\right) d\mu_h.$$

Expanding using the Gibbs forms:

$$D_{\mathrm{KL}}(\mu_h \,\|\, \mu_m) = \lambda\, \mathbb{E}_{\mu_h}\big[\mathcal{C}_m - \mathcal{C}_h\big] + \log \frac{Z_m}{Z_h}.$$

Misalignment is therefore the expected difference in constraint energy under human trajectories.

L.3 Alignment Loss as a KL Approximation

Proposition L.3 (Alignment Loss as KL Surrogate)**.** *Assume trajectories are observed at discrete steps $\{t_k\}$ and that $\mathcal{C}$ decomposes locally as a quadratic form in (ρ, θ, S). Then*

$$\begin{aligned} D_{\mathrm{KL}}(\mu_h \,\|\, \mu_m) \;\approx\; \mathbb{E}_{\mu_h}\Big[\sum_k \big(\alpha_\rho(\rho_k^{(h)} - \rho_k^{(m)})^2 + \alpha_\theta\, d_{S^1}(\theta_k^{(h)}, \theta_k^{(m)})^2 \\ + \alpha_S (S_k^{(h)} - S_k^{(m)})^2\big)\Big]. \end{aligned}$$

Sketch. Expanding $\mathcal{C}_m$ around $\mathcal{C}_h$ to second order yields a quadratic form in local deviations of (ρ, θ, S). Discretising the trajectory integral produces the stated sum, with the angular geodesic d_{S^1} arising from the periodicity of θ. □

The alignment loss $\mathcal{L}_A$ defined in Appendix I is therefore not an ad hoc metric but a Monte Carlo estimate of a KL divergence between trajectory distributions. Minimising $\mathcal{L}_A$ corresponds to matching the Gibbs measures induced by human and model constraint functionals. Alignment is achieved not when outputs coincide, but when the two systems assign similar probability mass to the same regions of trajectory space.

L.4 Variational Alignment Objective

Proposition L.4 (Variational Alignment Objective)**.** *Let* μ_h *be fixed. Minimising* $D_{\mathrm{KL}}(\mu_h \,\|\, \mu_m)$ *over model parameters is equivalent to minimising*

$$\mathcal{J}_{\mathrm{var}}(\theta) = \mathbb{E}_{\mu_h}\left[\mathcal{C}_m[\gamma]\right] + \frac{1}{\lambda}\log Z_m.$$

This is a free-energy objective: expected energy plus log-partition. Given sampled human trajectories $\{\gamma^{(i)}\}$, the empirical approximation is

$$\mathcal{J}_{\mathrm{var}} \;\approx\; \frac{1}{N}\sum_{i=1}^{N}\mathcal{C}_m[\gamma^{(i)}] + \frac{1}{\lambda}\log Z_m,$$

which, under the quadratic local approximation of Proposition L.3, reduces to the alignment loss $\mathcal{L}_A$ up to normalisation.

L.5 Gradient of the Variational Objective

Proposition L.5 (Gradient Form)**.**

$$\nabla_\theta \mathcal{J}_{\mathrm{var}} = \mathbb{E}_{\mu_h}[\nabla_\theta \mathcal{C}_m] - \mathbb{E}_{\mu_m}[\nabla_\theta \mathcal{C}_m].$$

The first term pushes the model to assign low constraint energy to human trajectories; the second pushes energy up on the model's own typical trajectories. This is structurally identical to energy-based model training.

L.6 Pathwise Alignment Criterion

Theorem L.6 (Pathwise Alignment)**.** *If* $D_{\mathrm{KL}}(\mu_h \,\|\, \mu_m) = 0$*, then for all admissible perturbation sequences* Π*, the induced operator cascades satisfy*

$$\mathcal{A}^{(N)}_{\Pi,h}(x) = \mathcal{A}^{(N)}_{\Pi,m}(x) \quad \textit{almost surely under } \mu_h.$$

Perfect alignment means identical transformation geometry, not merely identical outputs. This reframes alignment from a problem of output matching to a problem of distributional equivalence over transformation paths: two systems may agree on outputs while remaining maximally misaligned in the geometry of their underlying trajectory measures.

Standard supervised learning minimises discrepancies in outputs via pointwise losses. The present objective minimises a divergence between distributions over trajectories. Output agreement is neither necessary nor sufficient for alignment; what is required is agreement in the induced Gibbs measure over transformation paths. Cognitive alignment is therefore properly formulated as variational inference over trajectory space: the objective is to learn a constraint functional whose induced Gibbs measure matches that of human cognition. The amplitwist formalism provides the local parameterisation of transformations, the constraint functional defines admissibility, and the alignment loss serves as a tractable estimator of divergence between the resulting trajectory distributions.

References

[1] T. Needham, *Visual Complex Analysis*. Oxford University Press, 1997.

[2] L. V. Ahlfors, *Complex Analysis*. McGraw–Hill, 3rd edition, 1979.

[3] J. B. Conway, *Functions of One Complex Variable*. Springer, 2nd edition, 1978.

[4] K.-J. Engel and R. Nagel, *One-Parameter Semigroups for Linear Evolution Equations*. Springer, 2000.

[5] A. Pazy, *Semigroups of Linear Operators and Applications to Partial Differential Equations*. Springer, 1983.

[6] H. F. Trotter, On the product of semi-groups of operators. *Proceedings of the American Mathematical Society*, 10(4):545–551, 1959.

[7] T. Kato, Trotter's product formula for an arbitrary pair of self-adjoint contraction semigroups. *Topics in Functional Analysis*, pp. 185–195, 1978.

[8] L. C. Evans, *Partial Differential Equations*. American Mathematical Society, 2nd edition, 2010.

[9] I. Chavel, *Eigenvalues in Riemannian Geometry*. Academic Press, 1984.

[10] S. Rosenberg, *The Laplacian on a Riemannian Manifold*. Cambridge University Press, 1997.

[11] J. Cheeger, A lower bound for the smallest eigenvalue of the Laplacian. *Problems in Analysis*, pp. 195–199, 1970.

[12] S. H. Strogatz, *Nonlinear Dynamics and Chaos*. CRC Press, 2nd edition, 2018.

[13] E. M. Izhikevich, *Dynamical Systems in Neuroscience*. MIT Press, 2007.

[14] E. Marder and R. L. Calabrese, Principles of rhythmic motor pattern generation. *Physiological Reviews*, 76(3):687–717, 1996.

[15] M. M. Churchland et al., Neural population dynamics during reaching. *Nature*, 487:51–56, 2012.

[16] J. A. Gallego et al., Neural manifolds for the control of movement. *Neuron*, 94(5):978–984, 2017.

[17] K. Friston, The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11:127–138, 2010.

[18] G. Tononi, Consciousness as integrated information. *Biological Bulletin*, 215(3):216–242, 2008.

[19] R. Carhart-Harris et al., The entropic brain: a theory of conscious states. *Frontiers in Human Neuroscience*, 8:20, 2014.

[20] R. Brandom, *Making It Explicit*. Harvard University Press, 1994.

[21] J. Pearl, *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009.