When does cortical coherence carry shared content? Compatibility and self-reconstruction as testable conditions

Sebastian Bobadilla-Suarez
Independent Researcher
sebastian.bobadilla.s@gmail.com
ORCID: 0000-0002-5951-0772

20 September 2026

Abstract

Synchrony accounts treat collective phase order as the mechanism that integrates distributed cortical activity. We show formally that shared content and a self-model each need a condition beyond phase order, and state it. Three conditions are separated: coherence (collective phase order), compatibility (agreement of independently decoded local descriptions where regions overlap) and self-reconstruction (accurate recovery of the global state from an internal encoding). Compatible local descriptions determine a unique global state. Coherence bounds disagreement between overlapping descriptions only through shared encoders of bounded sensitivity, and the bound grows as the square root of the population read: on one simulated cortical sheet it constrains at most 10 sites, chaining through overlapping patches extends it to 0.312 mm, and across distant regions the link from coherence to content is an assumption. A contracting self-map fixes one state but reconstructs accurately only states within a diameter set by its error and contraction factor. Lean 4 verifies these implications; one

model satisfies all connecting hypotheses at once, showing consistency, not cortical realization. The conditions yield a disconfirmable comparison: at matched phase coherence and single-region decoding accuracy, decoders fitted independently on two regions should agree more, beyond their separate accuracies, on perceived than on unperceived stimuli classified by a no-report marker. Choosing the overlapping regions by what each decodes makes the unity correlate operational relative to a pre-registered decoder and criterion. An exploratory EEG analysis is a measurement check only. Unity and a minimal self remain interpretations that need an independent experiential endpoint.

1 Introduction

Endogenous extracellular fields can affect spike timing (Fröhlich & McCormick, 2010; Anastassiou & Koch, 2015) and may coordinate activity across cortical regions (McFadden, 2002, 2020; Pockett, 2012). Coordination raises a further question: when do those regions describe one situation, and when does the resulting state represent its own place in it?

Consider reaching for a cup. Vision and movement can describe different features while agreeing on the cup's location. Matching their timing leaves agreement about that shared quantity untested. An internal description of the perceiver's relation to the cup asks for still more. We call the three requirements *coherence* (collective phase order), *compatibility* (agreement of local descriptions on overlaps), and *self-representation* (accurate reconstruction from an internal encoding). The mathematical states are spatial measures and fields, compared with perceptual content through independently fitted decoders (Section 10.2).

Binding-by-synchrony and communication-through-coherence accounts treat phase relations as the medium of integration or of selective routing (Singer & Gray, 1995; Fries, 2015), and field theories place integration in the endogenous electromagnetic field (McFadden, 2020; Pockett, 2012). Both leave open what fact about synchronized activity makes it one content rather than coordinated activity carrying several. The contribution here is to separate the three requirements and prove which follow from which; each result names a measurement that phase data alone cannot supply.

The objection that synchrony does not fix content is older than this work. Critiques of the temporal binding hypothesis argued that synchrony, where it occurs, cannot say which features belong together or what they jointly represent: that grouping has to be computed before synchrony could signal it (Shadlen & Movshon, 1999). The formal results here are elementary: uniqueness of gluing, Banach's fixed-point theorem, the triangle inequality and a pigeonhole count. What formalization adds to the informal argument is quantitative and bookkeeping content. It gives the rate at which a coherence bound on content weakens with population size, the population at which it stops constraining anything and how far chaining through overlapping patches extends it, the diameter within which one contracting map can reconstruct, and the number of codes a separated family of states requires. It also gives a complete inventory of the premises a synchrony-to-content argument consumes, each of which names a measurement, together with one model in which all of them hold at once.

The familiar noisy mean-field oscillator supplies a stationary coherence threshold (Sakaguchi, 1988; Strogatz & Mirollo, 1991); we add a minimum cost for changing coherence quickly. A Lean 4 development checks the implications, and a joint witness shows that the eight connecting hypotheses of the formal composition are jointly satisfiable.

The empirical question is which measurements could support or disconfirm the proposed connections. A phase resultant can miss a locally ordered winding field; stationary phase distributions can leave the coupling mechanism unidentified; and coherent adaptation can lose alignment with a designated content template. We therefore distinguish stationary branches from finite observed onset, and compare prespecified field and connectivity onset models on held-out emergence recordings. The comparison tests a specified spatial prediction, not the presence of experience. Section 10.2 states the comparison that bears on conscious content directly.

The article makes six claims:

1. Compatible local descriptions determine a unique global state, while equal phases coexist with incompatible descriptions (Section 4).

2. Coherence bounds disagreement between overlapping descriptions through shared encoders of bounded sensitivity, by $L\sqrt{2N(1-r^2)}$ on the population read; on a simulated cortical sheet this constrains at most 10 sites, and chaining through overlapping patches extends it to 0.312 mm (Section 4.2).

3. A contracting self-map fixes one state and reconstructs within error ε only a family of diameter at most $2\varepsilon/(1-\Lambda)$; reconstructing M states pairwise more than 2ε apart takes M distinguishable codes (Section 5).

4. Changing coherence in time τ has a minimum current cost that grows as $1/\tau$ (Section 3.2).

5. Stationary phase statistics leave the coupling mechanism unidentified, finite ramps shift the observed onset, and a winding field keeps local order at global $r = 0$ (Section 8).

6. The link to content predicts that, at matched phase coherence and single-region accuracy, decoders fitted independently on two regions agree more, beyond their separate accuracies, on perceived than on unperceived stimuli; equal excess counts against the link to content (Section 10.2).

Sections 2–5 develop the three requirements; Section 6 applies their measurement obligations to a digital candidate. Section 7 states what the formal composition establishes, and Sections 8–9 examine observables and the cortical test, and Section 11 states the scope of every result in one place. Table 1 lists the connecting hypotheses the article argues from; Supplemental Material gives the full chain, derivations and controls. For a theoretical article the Methods are the definitions, proofs and simulations. Each result is stated in the section that uses it, Supplemental Material Table S1 maps it to its Lean declaration, and every simulation's model, parameters and controls are given

in Supplemental Material. The Results are the theorems of Sections 3–5 and the numerical controls of Section 8.

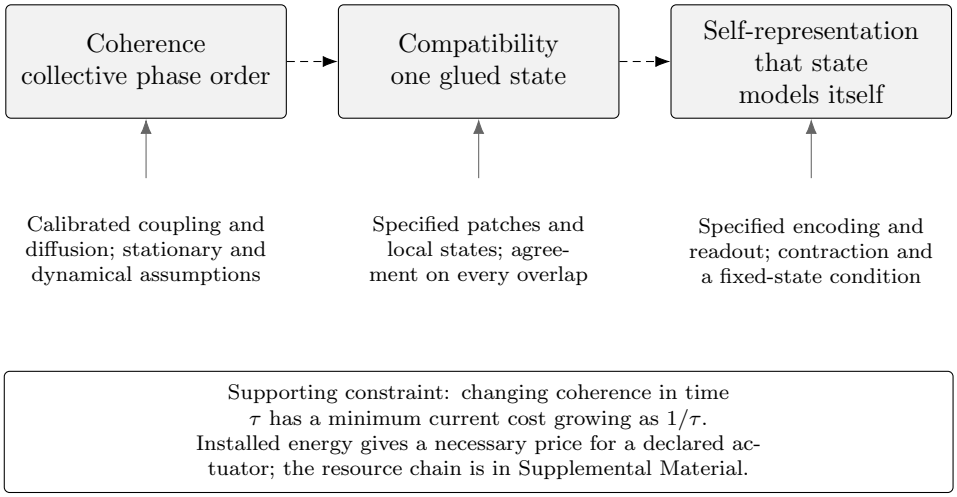


Figure 1: Three requirements of the proposed account. Each requires the inputs shown below it. Dashed arrows mark proposed connections, each of which needs the inputs of the condition it points to in addition to phase coherence. Thermodynamics constrains a candidate implementation through separate assumptions. This is the explanatory structure of the article; Table 1 records the connecting hypotheses of the formal composition that the article argues from.

2 Limits of phase order

Several limits hold without the cortical hypothesis. Equal phases can coexist with incompatible local descriptions; phase order constrains shared content only through encoders related on overlaps. A winding field can have zero global resultant while remaining locally ordered. A constant readout can have a unique fixed point without reconstructing its input. A region encodes only

the states it distinguishes, and by a declared deadline a site can report only changes inside its causal past. These limits reflect familiar phase oscillator and information constraints (Acebrón et al., 2005; Ott & Antonsen, 2008; Dörfler & Bullo, 2014; Breakspear, 2017). The sections below state each one beside the condition or measurement it motivates.

3 Coherence: physical coupling and its cost

3.1 Dynamics and the stationary threshold

The finite stochastic model is

$$d\theta_i = \left[\omega_i + \sum_j K_{ij}(t)\sin(\theta_j - \theta_i)\right] dt + \sqrt{2D}\, dW_i, \qquad r_N = \left|\frac{1}{N}\sum_i e^{i\theta_i}\right|. \tag{1}$$

Here θ_i is the phase of oscillator i, ω_i its natural frequency, $D > 0$ is phase diffusion, W_i are independent Wiener processes, and K_{ij} has inverse-time units, with angles dimensionless. Uniform mean-field coupling is $K_{ij} = K/N$; frequency heterogeneity and diffusion are distinct parameters. The stationary calculation below uses identical frequencies and this scalar coupling. Applying it to a spatial kernel requires an aggregation rule and a check of the resulting approximation.

For identical frequencies in their rotating frame, the positive normalized

stationary densities of the one-body Fokker–Planck equation are exactly

$$\rho(\theta) = \frac{e^{a\cos(\theta-\psi)}}{2\pi I_0(a)}, \qquad a = \frac{Kr}{D}, \qquad r = R(a) := \frac{I_1(a)}{I_0(a)}, \tag{2}$$

where I_n are modified Bessel functions and ψ is the collective phase. The formal derivation uses the absence of stationary current for a periodic gradient drift to obtain this density, then verifies normalization and self-consistency. At $K \leq K_c = 2D$, only $r = 0$ is self-consistent (Sakaguchi, 1988; Strogatz & Mirollo, 1991). Above threshold, exactly one positive solution coexists with zero and increases strictly with K. Its onset is

$$\frac{r^2}{K - K_c} \longrightarrow \frac{1}{D}, \qquad r \sim \sqrt{\frac{K - K_c}{D}}. \tag{3}$$

Figure 2 shows this stationary branch.

Linearization about the uniform density gives rate $K/2 - D$ for the first harmonic and $-Dn^2$ for harmonics $n \geq 2$. A growing mode therefore exists exactly above threshold. About the coherent density, the quadratic form is strictly dissipative transverse to collective rotation, whose rate is zero. These linear estimates concern the limiting equation. Relating them to finite stochastic trajectories requires a dynamical mean-field limit and branch selection; exchange symmetry alone does not establish the needed independence (Sznitman, 1991).

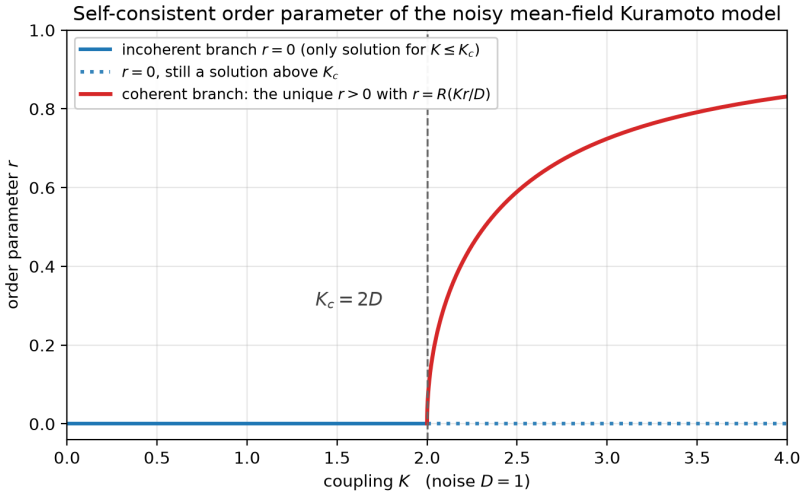


Figure 2: Stationary self-consistency for identical frequencies and phase diffusion $D = 1$. The positive coherent branch appears continuously above $K_c = 2D$ and increases with coupling. Zero remains a solution above threshold. The curve is computed from the integral expressions used in the formalization.

3.2 A speed limit on changes of coherence

One thermodynamic statement bears on coherence directly. For a smooth positive periodic phase density with probability current J and diffusion D, the current functional $\sigma_J = \int J^2/(D\rho)\, d\theta$ satisfies

$$\int_0^\tau \sigma_J \, dt \geq \frac{[\arcsin r(\tau) - \arcsin r(0)]^2}{D\tau} \tag{4}$$

while $0 < r < 1$. A change of coherence completed in time τ therefore has a minimum current cost that grows as $1/\tau$, whereas holding a stationary order has no positive floor. Read the other way, a current budget C sets a minimum

duration $\tau \geq [\Delta \arcsin r]^2/(DC)$ for a given change. This bears on ignition-like transitions, in which large-scale coherence rises abruptly (Dehaene et al., 2014), and on recovery latency after anaesthesia: each has a floor set by the current available to drive it. The bound is a heat or power statement only when J is identified with a physical flux, a reservoir at a stated temperature is specified and local detailed balance holds for the steps concerned; only then does σ_J convert to entropy production (Landauer, 1961; Norton, 2005; Shenker, 2000). Installed energy supplies a separate, necessary price: for a declared actuator whose response modes carry installation prices, mean coupling is at most κU_{inst} for a mode-data factor κ, so a coherent stationary branch requires installed energy $U_{\text{inst}} > 2D/\kappa$. Supplemental Material gives both derivations, the register and feedback budgets, and the sensitivity of the metabolic comparison to its assumptions.

4 Compatibility: from coordinated activity to a global description

4.1 Why common timing is insufficient

The formal setting is a sheaf: each patch of the substrate carries local descriptions, and restriction maps cut a description on a patch down to any smaller patch. A *section* is a description on one patch, and a *global section* is a description on the whole substrate. For overlapping patches U and V, com-

patibility means that restricting either local description to $U \cap V$ gives the same object. Descriptions can agree on shared quantities while contributing different information elsewhere. Compatible sections of the specified sheaf have a unique global extension. On a finite cover, local finite spatial measures agreeing as measures on every overlap also have a unique finite-measure extension whose sheaf image is the glued section. No mass-one normalization is imposed.

In the cup example, visual and reaching descriptions must express a shared quantity, such as cup location, in common coordinates. Restriction retains that quantity while dropping details particular to each description. Agreement permits surface appearance and hand configuration to contribute to one situation. The mathematical counterexample assigns the same phase to every patch while two local mass profiles disagree on a shared site. Those profiles have no common extension. Because they are chosen independently of phase, synchrony places no constraint on their disagreement.

This distinction limits the biological interpretation. Receptive territories and content domains must be related to spatial patches, and the descriptions must support the required extension. Matching overlap marginals of decoded probability laws need not determine a unique global joint distribution (Abramsky & Brandenburger, 2011).

The cover is itself a choice, and the construction constrains it in one way only: synchronized patches agree where they overlap. Where they do not overlap there is nothing to compare, so a cover by pairwise disjoint patches

satisfies that obligation for *every* assignment of local descriptions and phases and the gluing theorem still returns a unique global section, which is the tuple it was handed. Membership of the class then reports nothing about the system beyond the patches covering it. What a gluing claim asks of a cover is therefore a non-degeneracy the theorem does not assume: that distinct patches meet, and that what they share is a quantity fixed independently of the agreement it is being used to establish.

The cup example exposes a second gap, between two kinds of overlap. The sheaf is built over a spatial substrate, so its patches overlap where they share sites. Visual and motor territories that agree on cup location overlap in what they describe, not in where they lie, and the construction supplies no relation between the two. One route recasts the construction over content variables: each region becomes a partial decoder whose domain is the set of variables it estimates, and two regions overlap exactly where they estimate a common variable. Gluing and approximate selection hold on such a cover, and there compatibility is a consequence: regions decoding a variable within a tolerance agree on it within twice that tolerance, so only agreement beyond fidelity is evidence. The phase-to-content bounds concern spatial patches and do not transfer. One principle selects such a cover without presupposing the agreement it is used to test. Assign each region the variables that can be decoded from its activity above a declared criterion, on data held out from the agreement test, and let two regions overlap on exactly the variables decodable from both. Decodability is fixed by what each region carries before

any comparison is made, so the cover cannot be chosen to make agreement come out. It is non-degenerate exactly when its nerve is connected, meaning that every region is linked to every other through a chain of shared decodable variables, and recordings can fail that condition. The principle has three inputs, the candidate variables, the criterion and the decoder class, and each has a rule that does not consult agreement.

The candidate variables are those the experiment manipulates and the task makes behaviourally relevant, fixed before recording. Data-driven alternatives, such as directions of maximal shared variance between two regions, select for the agreement under test and are excluded. The criterion becomes a statistical one: a region encodes a variable when its cross-validated decoding exceeds a permutation null at a pre-registered level, and results are reported across a declared range of levels. The level need not be exact, because two accurate decoders agree automatically (Eq. (13)); the statistic of Section 10.2 is agreement in excess of what each region's accuracy implies, which remains valid at any admitted level. The decoder class is bounded on both sides. Unrestricted nonlinear decoders recover a variable from any region carrying information about it, so nearly every pair of regions overlaps and the cover again constrains nothing. The positive bound of Section 4.2 requires encoders with a declared Lipschitz constant L, which enters δ_r directly, so the class should consist of L-Lipschitz maps; linear readout, the usual operationalization of information explicitly available to downstream neurons (DiCarlo & Cox, 2007; Kriegeskorte & Diedrichsen, 2019), is one such class. Within the

admitted family, the decoder and the distance in which content discrepancy is measured are chosen by held-out decoding accuracy and never by agreement. Neither is neutral: which similarity measure best describes neural activity differs across tasks (Bobadilla-Suarez et al., 2020). What remains is that the cover is the experimenter's. Whether cortex realizes unity over the chosen variables is the empirical claim under test, and no choice of decoder settles it. The correlate is therefore operational relative to these pre-registered inputs, and carrying the formal results to such a cover is the recasting sketched above.

4.2 When coherence does constrain content

A declared relation between phase and content gives a positive result. Let contents take values in a real normed vector space $\mathcal{C}$. Patch i reads an assigned oscillator through $e_{i,x}$ at site x, evaluated on the circle point $(\cos\theta, \sin\theta)$. Assume a uniform Lipschitz constant L in chord distance, meaning that an encoder's output moves by at most L times the chord between two input phases, and require $e_{i,x} = e_{j,x}$ wherever patches i and j overlap. Encoders may vary across sites; their equality on overlaps is a supplied relation between descriptions of the shared quantity. Write $z_k = e^{i\theta_k}$ and $m = N^{-1}\sum_k z_k$, so $|m| = r$. The phasors' squared deviations from their mean sum to $\sum_k |z_k - m|^2 = N(1 - r^2)$, and $|z_i - z_j| \le |z_i - m| + |z_j - m|$ gives $|z_i - z_j|^2 \le 2N(1 - r^2)$ for each pair. Consequently every overlap

discrepancy is bounded by

$$\delta_r = L\sqrt{2N(1-r^2)}. \tag{5}$$

Here discrepancy is the norm of the difference, and r and N are read on whichever population the bound is applied to. That choice is what decides whether the estimate says anything. A patch carries an order parameter of its own, and reading the bound there carries that patch's resultant and that patch's site count, so a large population observed through small patches is constrained by the small count. Read instead on the whole array, the factor $\sqrt{N}$ makes the estimate weak, because it draws a pointwise guarantee from an average over N sites. The order cannot be improved in general: one oscillator in antiphase among N leaves $N(1-r^2)$ near 4. At $r = 1$ the observations agree exactly at either scale. A shared scalar encoder is a special case. Dropping the overlap relation permits disagreement even at perfect locking. A common nonexpansive linear projection preserves agreement.

What Eq. (5) bounds is L times a chord between phasors, and a chord never exceeds 2, so the estimate constrains anything only where $\sqrt{2N(1-r^2)}$ falls below 2, whatever the value of L. On the simulated sheet of Section 8, at its patch-local order for a coupling decay length of 0.1 mm, the bound is informative for at most 10 sites; this count and the chain's reach below are values for that sheet and move with the decay length and patch order. Three conditional results decline the step that puts $\sqrt{N}$ there. A pair count bounds

the fraction of disagreeing pairs at any population size but names no pair. A chain across the sites that overlapping patches share compares distant territories through the patches between them, with error growing linearly in hops rather than with population: on the same sheet it stays informative for 19 hops, 0.312 mm. A spectral gap removes the population count at a frequency-locked configuration, given a declared algebraic connectivity of the coupling graph. Territories whose agreement unity concerns lie centimetres apart, beyond the chain's reach, and across them the coherence-to-content link is the bridge assumption E78 (Table 1). Supplemental Material gives the three derivations and the sheet's numbers.

Mix retained fields with observations at gain $\eta \in [0, 1]$: $s_i \mapsto (1 - \eta)s_i + \eta o_i$. Content and observation discrepancies bounded by ε and δ give a next discrepancy at most $(1 - \eta)\varepsilon + \eta\delta$. Repeated observations bounded by δ therefore give

$$\varepsilon_n \leq (1 - \eta)^n \varepsilon_0 + \delta. \tag{6}$$

Thus agreeing observations preserve exact compatibility and, for $\eta > 0$, reduce an initial mismatch. Approximate observations leave a nondecaying term in the error bound. Taking a uniform bound on Eq. (5) connects this update to the phase process.

4.3 What approximate agreement determines

Exact gluing requires exact restrictions. For nonnegative profiles on a finite substrate, suppose every overlap discrepancy is at most ε in uniform distance. Weights that vanish outside their patches and sum to one at each site select a global profile by weighted averaging. The selection lies within ε of each patch there, and any two such selections differ by at most ε, independently of cover multiplicity. An L-Lipschitz readout consequently varies by at most $L\varepsilon$ across selections.

At zero tolerance, the selection is the unique exact extension. At every positive tolerance there are families with no exact extension; selection then gives an approximation whose dependence on weights is bounded. Supplemental Material gives a three-site witness attaining the bound. The mass-profile construction uses the scalar case of the content dynamics and additionally requires nonnegativity of the initial profiles and observations.

These constructions concern declared objects: a specified cover, contents valued in a normed space, and encoders related on overlaps. Section 8 sets out what a measurement of their neural counterparts would have to establish.

5 Self-representation: a state that models itself

5.1 The reflexive construction

Compatibility gives a global state. Reflexivity relates it to an encoding within that state. Equip global sections with a complete metric. *Restriction resonance* is the additional condition that the encoding is the state's own restriction to a region $\mathcal{A}$, so that the region encodes the state by being part of it. Under that condition the reconstruction map is

$$F(s) = \text{readout}(s|_{\mathcal{A}}). \tag{7}$$

The region and readout are fixed across the state space. If F contracts and fixes the section s_* formed from the local descriptions, Banach's theorem makes s_* its unique fixed point (Banach, 1922). The formal bridge E89 (Table 1 gives the full composition and its edges) supplies the contraction and $F(s_*) = s_*$; restriction resonance is separate. Supercritical coupling makes the stipulated contraction factor $\Lambda = e^{-(K-2D)\tau/2}$ less than one for a relaxation time $\tau > 0$; Section 11 states the standing of that factor. A fixed contraction can reconstruct accurately only a restricted family of states: if its error is at most ε throughout that family, any two members are at most $2\varepsilon/(1-\Lambda)$ apart.

In the cup example, the global state includes the represented scene and the

perceiver's relation to it. A local encoding must retain enough information for the readout to recover that modelled state; cup position alone would be insufficient if hand configuration and viewpoint could not be reconstructed.

For the same map with Lipschitz factor $\Lambda < 1$, approximate reconstruction gives

$$d(s, F(s)) \leq \varepsilon \quad \Longrightarrow \quad d(s, s_*) \leq \frac{\varepsilon}{1 - \Lambda}. \tag{8}$$

At the stipulated rate this radius grows as $2\varepsilon/((K-2D)\tau)$ near threshold, for fixed positive ε and τ. Exact uniqueness survives while the error guarantee weakens.

5.2 Reconstruction across a family of states

Sensitivity can be required quantitatively. Declare relevant macrostates, an encoder Φ and readout R, and require $d(s, R(\Phi(s))) \leq \varepsilon$ throughout that family. If two states have the same code, the triangle inequality gives

$$\Phi(s) = \Phi(t) \quad \Longrightarrow \quad \max\{d(s, R(\Phi(s))), d(t, R(\Phi(t)))\} \geq \tfrac{1}{2} d(s, t). \tag{9}$$

Thus M relevant states pairwise more than 2ε apart require at least M distinguishable codes; an alphabet of at most 2^b codes requires $M \leq 2^b$. Applied to the reflexive boundary, E is its local encoding and $R \circ E$ is the existing self-prediction map.

This is an additional accuracy criterion on a family, with no requirement for multiple exact fixed points. For the same fixed contracting map used by

E89, however, the diameter bound above is an additional restriction: at zero error the family contains at most one metric-distinct state. A broad family needs an input-conditioned map, one contraction F_u per input u. If F is L-Lipschitz in u, members with inputs u, v lie within $(2\varepsilon + L\,d(u,v))/(1-\Lambda)$ and the fixed points within $L\,d(u,v)/(1-\Lambda)$, so a self-state can follow the scene. The code criterion then binds within each input: a readout storing one state per input reconstructs without encoding anything. Under the fixed map of E89 the minimal-self correlate is a condition on one self-state. A blind encoder or constant readout fails the criterion for separated states despite having a unique fixed point.

Where the encoding is read as the state's own restriction to a region, the criterion becomes a condition on the region. Restriction resonance requires the encoder and that restriction to identify the same pairs of states, so a boundary reconstructing two relevant states more than 2ε apart while its region reads them identically is not resonant. Read forwards, a resonant boundary that reconstructs a pairwise separated family restricts injectively on it. Which states a region reads identically is fixed by where the region sits, before an encoding is chosen, so the resolution the family demands is the region's and no readout supplies it. The implication runs one way: a region that does resolve the family is not thereby shown to carry a resonant encoding.

Resonance and resolution are jointly satisfiable without a local fold. Take the region to be the whole substrate and use the identity readout: restric-

tion is the identity, so this boundary is resonant and reconstructs every declared family exactly. Its map has Lipschitz factor one and does not satisfy E89's strict contraction premise on a nontrivial metric space. A properness condition on the region is therefore a separate requirement, and the fixed-contraction diameter bound remains in force when E89 is also required. With an L-Lipschitz readout, a family reconstructed within ε whose members lie pairwise more than $d_{\min}$ apart has δ-separated codes whenever $\delta L \leq d_{\min} - 2\varepsilon$, so the family is no larger than the δ-packing number of what the region carries. At the whole substrate it says nothing, because the region's range is everything. The gain L is a declared input, and the bound weakens without limit as the gain grows.

5.3 Agreement after a change takes communication

Snapshot compatibility does not specify how quickly descriptions respond when a shared quantity changes. In a synchronous network whose updates read only the site's own state and declared neighbours' messages, two initial configurations agreeing on a site's causal past through T rounds produce the same state there at round T. The deadline consequence is the one Section 2 states, and it follows from the reconstruction bound in Section 5.2: a report from that site is a readout of a code that the two interventions leave identical.

A three-site delay line witnesses failure after one round and exact reporting after two; cutting the path prevents delivery at every round. A report may still happen to be correct about one fixed world. The bound concerns

responsiveness across interventions.

Throughout this section *prediction* names a reconstruction map, and the relevant family of states, the metric, the tolerance, the communication graph and the deadline are all declared inputs, as are the represented variables and the fidelity counted as accurate. Reading the map as temporal neural prediction would require dynamics; reading the resulting code count as physical memory capacity would require a hardware model, which a state-dependent readout would have to extend.

6 What the conditions ask of a digital candidate

A claim that a digital system, such as a large language model (LLM) running on graphics processing units (GPUs), has unified content or represents itself commits its maker to the measurements this account asks of cortex. The maker must identify the shared quantities the system's parts describe, show that decoders fitted independently on those parts agree about them, and show that a declared local code separates the relevant family of states and survives intervention. The indicator-properties approach assesses artificial systems against computational properties drawn from theories of consciousness (Butlin et al., 2023); the conditions here are fewer and each is a measurement rather than an architectural property. A digital system is also where they can be measured today, with full access to its state and

exact interventions on it: fitting decoders for one shared quantity on two separate layers, heads or token positions, and testing their agreement and whether reconstruction survives intervention, would carry out the proposed measurement directly.

One constraint follows from the architecture alone. Within one forward pass of a causally masked model, position t reads only its causal past, so an intervention outside that past cannot be reconstructed at t by that step. This is the deadline result of Section 5 applied to the mask, and it bounds which variables a position can encode about the rest of the sequence within one step. A second constraint needs a premise the architecture does not supply: a linear encoder and readout factoring through d coordinates pay reconstruction error at least the discarded spectral tail of the coupling operator, but query/key dimension does not supply such a bottleneck for softmax attention. These results rank no architecture family, imply no general inferiority of GPU hardware and settle no question of machine experience. A companion paper (Bobadilla-Suarez, 2026) adds a condition on how agreement is produced: it proposes that unity requires agreement enforced by a coupling through which the graded physical state of each region pulls the regions it overlaps toward agreement, and proves that content read through noise margins, as content read from bits is, never meets that condition. Under that premise the measurements proposed here are necessary for unity in any substrate, and a digital system that passes them computes its agreement rather than enforcing it. Supplemental Material gives the bottleneck deriva-

tion, a synthetic two-view task with its controls, the information ceiling of token outputs and the hardware scope.

7 The formal composition

The Lean development composes the three requirements into one conditional theorem. Its passive and feedback branches each accept eight connecting hypotheses, running from finite capacity through a register, dissipation and a continuum coupling to the field, coherence, the cover and the self-prediction fixed point. Each edge is a hypothesis of the composition, and its inferential content is its last step, where contraction and a fixed glued section give uniqueness. Table 1 lists the four edges the article argues from; Supplemental Material gives the full chain, its thermodynamic edges and the register ledger of the feedback branch. Both branches also take a cover carrying exact overlap compatibility and a reflexive boundary specifying encoding and readout, and conclude that the glued section is the unique fixed point of self-prediction. Composing the edges establishes that the premises are consistent, named in one place and machine-checked together, and joint witnesses satisfy every edge of each branch at once. The development declares no additional axioms beyond its libraries' foundational axioms.

Table 1: The four connecting hypotheses the article argues from, shared by both branches of the formal composition. Supplemental Material gives the full eight-edge chain and, in Table S1, maps the node-level results.

Connection	Status	Required assumption
E56 *field identification*	Physical commitment	A specified physical field realizes the limiting coupling and satisfies the kernel conditions.
E67 *above threshold*	Physical commitment	That field operates above the synchronization threshold.
E78 *coherence to cover*	Bridge assumption	Coherence relaxes the supplied cover at coupling at least the named mean-field value. Shared phase encoders bound vector-content mismatch; the exact cover and timely responsiveness remain separate requirements. Patches and oscillators carry one index, so one oscillator carries exactly one content patch: a modelling assumption rather than a consequence, and not an innocent one, since field-mediated coupling is spatial while two territories describing one quantity need not be neighbours. Where the cover's nerve is connected the same edge carries through the sites its patches share, at a residual linear in that nerve's diameter; where the configuration is frequency-locked it carries through the coupling graph's algebraic connectivity λ_2 instead of through any patch count, which asks the cover for nothing and is declared hardware data in the standing of κ.
E89 *contraction to fixed point*	Modelling assumption	A map mediated by local self-encoding fixes the glued section and has the stipulated Lipschitz factor $\Lambda < 1$. Reconstruction over a declared family imposes an additional code-separation criterion.

8 What the observations establish

The composition holds under its premises. Whether observation can establish those premises is a separate question, and the three requirements it connects concern different objects: phase dynamics, local descriptions and internal encodings. Their observables must retain these distinctions. The following controls show which mechanisms remain indistinguishable, which temporal approximations fail, which states one observable cannot separate, and what independent content measurements must establish. Each is a finite numerical regime on a declared model, with full parameters and controls in Supplemental Material, and none of them calibrates a cortical quantity.

8.1 Stationary distributions leave the mechanism undetermined

Equation (2) gives a concentration–coherence relation with no fitted shape parameter. Estimate r from circular resultant length and a from density shape: inverting R from the measured r would make agreement automatic. Even independently estimated agreement identifies a statistical family shared by different mechanisms.

Let $q(t)$ cross 2, let $r(q)$ be its stationary branch, and fix $D_0 > 0$. Increasing coupling $K(t) = D_0 q(t)$ at diffusion D_0 and decreasing diffusion $D(t) = K_0/q(t)$ at coupling K_0 both give concentration $a = q(t)r(q(t))$. Uncoupled oscillators with prescribed drift $-h(t)\sin(\theta - \psi)$, where $h(t) = D_0 q(t) r(q(t))$,

and diffusion D_0 have that same density. Their stationary distributions coincide window by window despite different causes.

Resolved dynamics supply additional information. In the ideal model with continuous latent paths, quadratic variation accumulates at rate $2D$ and conditional increment means constrain drift; a filtered observation of the phase needs its own observation model before either is read off. Separating endogenous coupling from a matched common drive additionally requires a measured drive or controlled perturbation; identifying electromagnetic coupling requires evidence specific to that interaction. Supplemental Material gives the equivalence construction and its numerical check. Phase statistics pooled over time and over bipolar pairs cannot supply these estimates: on public sedation EEG (Bajwa et al., 2025), the pooled pipeline places a perfectly phase-locked signal at $r = 0.0001$, so a spatial test needs an estimand that removes the common rotation and a montage whose effect on spatial phase is modelled; Supplemental Material reports this estimator exercise and the dependence-aware design.

8.2 Finite ramps change the observable onset

The square-root law in Eq. (3) describes stationary solutions, and a changing coupling needs time for the phase distribution to approach them. In a deterministic mean-field control crossing the threshold, none of the scanned rates meets the declared tracking tolerance at onset. In a finite-population ramp study ($N = 2000$, $D = 1$, 32 replicas, 8 speeds), faster ramps delay escape.

The fitted delay exponent 0.444 has a replica-bootstrap interval $[0.398, 0.496]$ that excludes the ideal $1/2$, a deterministic reference shows that the escape criterion changes the fitted exponent, and no common asymptotic exponent is established. Delayed passage through a noisy bifurcation (Berglund & Gentz, 2002) makes the apparent exponent depend on ramp speed and fitting window. Supplemental Material gives the sweep and its figure, the fluctuating-coupling and population-size controls, and the signed-network crossing bracket.

8.3 A winding state leaves the order parameter ambiguous

The global order parameter r discards spatial arrangement. On a periodic sheet, a phase field can wind once around the sheet and have $r = 0$ while neighbouring sites remain locked. At identical frequencies, every such winding is stationary for an isotropic symmetric kernel: contributions at opposite separations cancel. It violates the quarter-turn initial-spread premise of the convergence result, which is why that result can still select the uniform state.

In the spatial control, the winding persists below the interpolated decay length 0.1871 mm and unwinds above it (Fig. 3). At 0.1 mm, uniform and winding starts under the same coupling, diffusion and kernel end at $r = 0.9657$ and $r = 0.0083$; patch-local order is 0.9015. Under quenched frequency spread 1.5 $\mathrm{rad\,s^{-1}}$, the measured retention boundary is 0.1421 mm,

and the corresponding order comparison is 0.9418 against 0.0348 over 3 frozen realizations. These are finite controls, not a cortical wave identification or a derived transition.

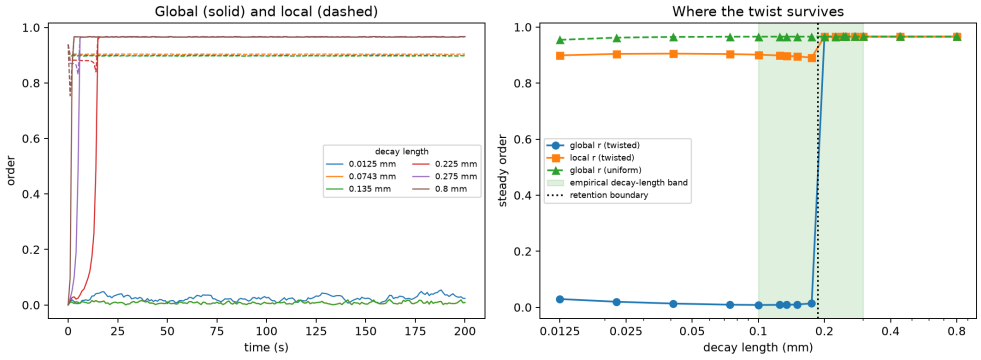


Figure 3: A winding phase field on the spatial sheet, at identical frequencies. Left: global order (solid) against patch-local order (dashed) across a subset of coupling decay lengths. Right: the twisted and uniform initial conditions compared across lengths, with the shaded 0.1–0.3 mm band and the interpolated retention boundary. Below that boundary the two initial conditions reach different stationary values of r at the same coupling, diffusion and kernel, so the resultant does not report the coupling. The sweep is sampled at the lengths marked.

Linearization about a winding gives explicit mode rates, which concern linear stability rather than the measured retention boundary. An amplitude extension adds an existence condition under which a broad kernel can remove a winding that the phase-only equations retain as a stationary state.

Neither model turns these states into measured cortical spirals. The periodic sheet cannot carry a lone phase defect, the relaxational dynamics has no propagation speed, and the amplitude variable has no calibrated relation to a recorded envelope. Supplemental Material gives the spectral, amplitude

and defect results. Locally, Eq. (5) can still bound content disagreement on a winding patch when shared encoders are supplied, although its global $r = 0$ form is uninformative. A nonzero ring winding number q also obstructs a global real-valued phase lift; it has no identified content meaning here. Thus an onset criterion based only on r can register a change of basin as well as a change of coupling. The protocol of Section 9.3 needs a spatial phase observable.

8.4 Coherent adaptation can lose environmental alignment

An adaptive coupling rule might supply representational structure, but its objective must reward the relevant agreement. For symmetric coupling, pairwise sine terms cancel in the population mean drift. Minimizing squared drift at that fixed mean therefore rewards equalizing drift across nodes. In the frequency-cluster control, this favors coupling between unlike-frequency clusters, whereas the diagnostic environmental template rewards coupling within clusters. Supplemental Material proves descent of this objective for a specified gradient flow at fixed phases and distinguishes it from entropy production.

The finite control co-evolves noisy phases and symmetric, nonnegative couplings at fixed total resource. Its template is a diagnostic covariance of latent signals, never a training input. Across three illustrative seeds,

second-half order remains between 0.9167 and 0.9222, while the within-cluster to between-cluster coupling ratio falls from 0.9824–0.9948 to 0.3507–0.6135 (Fig. 4). Random-update, permuted-gradient and frozen-kernel controls are reported in Supplemental Material.

The immediate post-update objective reduction disappears when averaged over complete update intervals. A prespecified grid of 6 rate–cadence combinations contains no case meeting the 5% sustained-reduction criterion while retaining coherence and losing alignment. The finding is coherent adaptation away from the specified structure, with no demonstrated sustained cost reduction.

Kernel alignment and overlap compatibility are different observables. This control addresses acquisition of its designated template under its particular objective. Learning models can explicitly include data in an objective (Corberi et al., 2026); an account of agreement about the cup must specify the content its learning rule is meant to preserve.

8.5 Content agreement needs independent decoders

A neural test needs independently identified domains, local descriptions and restriction maps. Two decoders should estimate a shared content on overlapping domains before their agreement is examined. Discrepancy requires a null under decoding noise and controls using deliberately compatible and incompatible examples. An implementation on synthetic data separates these cases at AUC $= 0.986$, on a 95% interval of $[0.980, 0.992]$ resampled by trial;

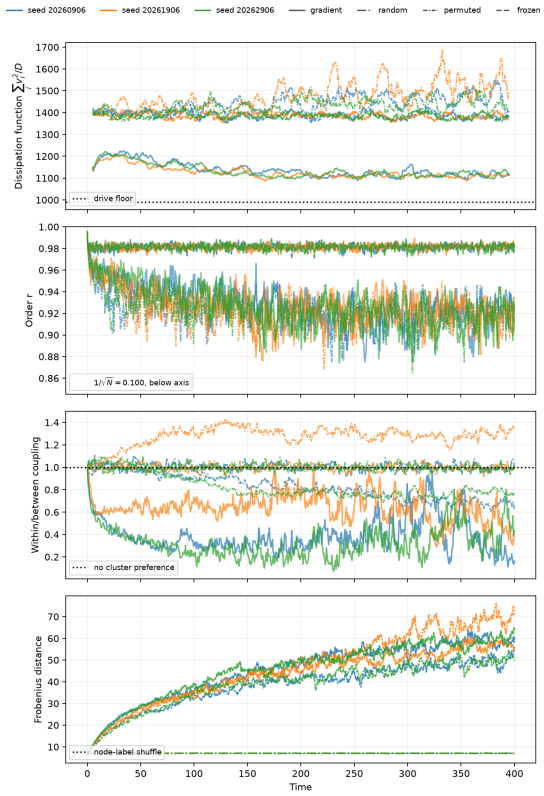


Figure 4: Coherence persists while coupling loses alignment with the frequency-cluster template. Gradient updates (solid), random updates (dash-dotted), permuted gradient updates (dash-dot-dotted), and a frozen kernel (dashed) share one phase-noise stream per seed. Panels show the immediate post-update squared-drift objective, phase order, within-cluster to between-cluster coupling, and template distances. The gradient arm preserves order while its coupling ratio falls; growth in both distances mainly reflects kernel norm. The objective panel samples the instant after updates; complete-interval measurements find no sustained reduction in the bounded study.

the compatible construction scores 0.476 on an interval containing one half. Shared-prior shrinkage can nevertheless manufacture agreement, and regional bias can manufacture disagreement. Per-patch decoding diagnostics detect these failures; the compatibility statistic alone does not.

On constructed content, phase carries a share 1.000 of the content entropy when content is built from it and 0.000 when it is not. A phase-derived statistic reports compatibility on phase-locked patches whose local descriptions disagree. The implemented content statistic detects their incompatibility. These controls validate a measurement on known constructions; a neural content model and decoder remain to be established.

Choosing the cover also chooses a candidate subject: trivial descriptions can agree while carrying little content, and a partition can conceal disagreement. Section 4 gives the decodability rule that fixes the cover without consulting agreement.

9 The cortical hypothesis and an awakening test

The cortical proposal identifies a physical source of coupling and asks whether its changes help explain recovery. The tests in Section 8 determine what observations must distinguish before a phase trace can support this identification.

9.1 Physical identification and calibration

The cortical hypothesis identifies part of effective coupling with endogenous extracellular electromagnetic interactions. Fields can modulate spike timing (Fröhlich & McCormick, 2010; Anastassiou & Koch, 2015), and field models propose coordination of distributed neural information (Pinotsis & Miller, 2022; McFadden, 2020). Phase organization under propofol motivates studying this regime (Bastos et al., 2021).

The formal field predicate requires a jointly continuous kernel, a normalized substrate, coupling equal to the kernel average, and no single site providing more than half that strength. Toy constructions satisfy it. A cortical realization must specify the scale at which the kernel and its reduction apply: calcium imaging reports different state-dependent modularity patterns at single-cell and mesoscale resolution (Kiyooka et al., 2026).

Cortical columns are a serious candidate for the oscillator index. They are delineated by response properties and anatomy rather than by the agreement they would be used to establish (Mountcastle, 1997), and a phase can be estimated from a column-sized territory at the electrode spacings used for cortical wave measurements (Townsend et al., 2015). They are a poor content cover: non-overlapping columns form a disjoint cover (Section 4), much of association cortex has no accepted columnar parcellation, and the column's standing as a canonical unit is contested (Horton & Adams, 2005; Rakic, 2008). The content cover is the decodability cover of Section 4, and the bridge assumption E78 records its identification with the oscillator index.

The spatial reduction is also testable. A row-sum aggregation at zero diagonal tracks the uniform-coupling control over the tested spatial decay range, while frequency heterogeneity displaces the operational threshold. The fit window biases both estimates relative to $2D$; agreement with the control is therefore the supported conclusion. Supplemental Material reports the reduction, its numerical error and the distinct assumptions needed to identify the kernel average with the scalar coupling of the installed-energy condition.

Converting a spike-time shift into coupling needs a rate: the dimensionless product of population size N, timing sensitivity $\mathcal{T}$, field amplitude E and frequency f becomes a coupling through a conversion rate γ,

$$K_{\mathrm{eff}} = \gamma N \mathcal{T} E f. \tag{10}$$

This rate describes how induced timing shifts accumulate into drift.

9.2 Extracellular geometry and recovery

The proposed control variable is extracellular geometry. Extracellular volume fraction is approximately $\alpha = 0.2$ at tortuosity $\lambda \approx 1.6$ (Syková & Nicholson, 2008); a sleep study reports a rise from roughly 0.14 to 0.23 (Xie et al., 2013). Below 10 kHz, cortical tissue conducts approximately resistively at 0.3–0.6 $\mathrm{S\,m^{-1}}$ (Barbour, 2017). At fixed transmembrane current density, an effective-medium description predicts that a larger extracellular fraction raises conductivity and reduces field amplitude.

The coupling-gated hypothesis consequently predicts that extracellular contraction on awakening raises E and, holding the other factors in Eq. (10) fixed, raises effective coupling. This argument supplies a direction; the geometry–coupling relation, its magnitude and its role in human recovery require measurement.

For a smooth crossing at nonzero speed, evaluating the stationary branch along the changing control gives a deterministic tracking approximation that fails at onset in the tested crossing control:

$$r_{\mathrm{ss}}(K(t)) \sim \sqrt{\frac{\dot{K}(t_0)}{D}}\,(t-t_0)^{1/2}, \qquad K(t_0)=K_c, \quad \dot{K}(t_0)>0. \tag{11}$$

This is a deterministic tracking approximation from Eq. (3) at fixed D; the crossing controls in Section 8.2 show that it fails as an observed-onset approximation under the tested initialization. A separation of geometry and phase time scales can support tracking away from threshold. The awakening hypothesis needs a calibrated dynamical prediction for an operational onset above the finite-size fluctuation floor. Crossing K_c alone does not supply a recovery latency.

9.3 A protocol for the awakening experiment

Neural inertia describes a behavioural barrier and hysteresis between anaesthetic induction and emergence (Friedman et al., 2010); recovery can pass through metastable activity states (Hudson et al., 2014; Proekt & Hudson,

2018). The $v^{1/2}$ exponent of an ideal delayed bifurcation is generic (Baer et al., 1989; Berglund & Gentz, 2002) and cannot alone discriminate a field mechanism from a synaptic one. A possible discriminator is spatial onset relative to independently estimated local field coupling and network connectivity. The field-based onset map below is a separate model from the mean-field one, specified before the observations are inspected.

For that comparison, partition the recording into prespecified regions i and measure pre-onset order $r_i(t_0)$ and diffusion D_i. Define a field score $h_i = N_i \mathcal{T}_i E_i f_i$ from independent pre-onset field and timing measurements, and a connectivity score c_i from a separately estimated structural network. For the same measured emergence drive $g(t)$, compare $K_i^F(t) = a_F g(t) h_i$ with $K_i^C(t) = a_C g(t) c_i$. Each model has one fitted scale. The local first-harmonic approximation predicts

$$\widehat{r}_i^M(t) = r_i(t_0) \exp\left(\int_{t_0}^{t} [K_i^M(s)/2 - D_i]\, ds \right), \qquad M \in \{F, C\}. \tag{12}$$

Predicted onset is the first crossing of a region-specific threshold fixed from pre-onset fluctuations. Fit the two scales on training sessions, then compare held-out onset times and site order under the same observation model. Larger field scores predict earlier onset only when baseline order and diffusion do not reverse the ranking. This declared approximation omits interregional coupling and saturation; a successful ranking would support its field covariate, not prove ephaptic causation. Supplemental Material specifies the

controls and failure conditions.

A within-subject protocol can manipulate emergence rate, with spatially resolved recordings to compare the proposed field and connectivity models:

1. **Manipulate emergence rate.** Vary the propofol infusion taper across sessions within the same subject, recording drug concentrations or a validated pharmacokinetic proxy to quantify the manipulation.

2. **Measure behavioural recovery.** Record time-to-response to a standard stimulus as the independent behavioural endpoint, rather than estimating ΔK directly.

3. **Resolve spatial onset.** Identify anatomical locations where coherence first crosses the baseline-scaled threshold and score the two maps of Eq. (12) on held-out sessions. The two covariate maps are fixed before the outcome sessions are inspected.

4. **Match physical controls.** Confirm that the observed coherence-to-amplitude relation holds under the specific local diffusion D and observation noise. The model's quadratic-variation rate $2D$ applies to the latent phase; a band-limited analytic phase is smooth below its correlation time, and on a constructed field its squared-increment rate changes 37-fold between 0.2-ms and 20-ms steps. Estimating D therefore requires an explicit observation model for filtering and sampling.

Predeclare the recording window, regional fluctuation thresholds, and a shared timing-error model. Fit the two scales and that observation model

on training sessions only; on held-out sessions, score the observed site-onset times and sites that never cross by the same regional predictive rule. Report site order as a secondary outcome and behavioural time-to-response separately. Sessions without a crossing are censored observations, not zero delays or exclusions. Compare the two scores only where independently estimated field and connectivity maps predict distinguishable onset patterns at the recording resolution. A held-out field-model advantage would support that specified spatial predictor; it would not identify an ephaptic cause, shared content or experience. Supplemental Material gives the controls and pilot quantities needed to judge whether the comparison is informative.

10 Discussion

10.1 Unity and a minimal self as interpretive motivation

We propose the state glued from compatible local descriptions as a formal correlate of conscious unity, and its accurate internal reconstruction as a correlate of a minimal self. In the cup example, this would mean one description that preserves the cup's relation to the perceiver as well as the scene. Compatibility permits distinct contents and uncertainty; the shared quantities must be chosen from independent evidence about what each region encodes. The reconstruction need not be a verbal or autobiographical account.

The minimal self, in Gallagher's sense, is the pre-reflective sense of being the subject of an experience, present before any narrative or reflective self-description (Gallagher, 2000). Two of its features motivate the proposed correlate. It is not a separate object of experience but a perspective built into the experienced scene; correspondingly, the reconstruction is not a second representation beside the glued state but a map whose fixed point is that state, read through a part of it. It is immediate rather than inferred from a higher-order state; correspondingly, the fixed point needs no further state that represents the first. The correspondence is structural. Self-model theory locates the phenomenal self in a transparent model a system has of itself (Metzinger, 2003), and interoceptive-inference accounts ground it in prediction of the body's own states (Seth, 2013). The present condition is narrower than either: it asks only that a declared local encoding recover the global state, including perceiver-relative variables, at a declared fidelity, and it says nothing about transparency or about which bodily variables must be modelled. Either account could supply the family of states and the fidelity that the condition leaves as inputs. Under E89 that family is one state at zero error and a family of diameter at most $2\varepsilon/(1-\Lambda)$ within error ε (Section 5.2), so under E89 the condition constrains a single self-state; a self that persists through change needs the input-conditioned map of Section 5.2. Its proposed relation to experienced perspective remains an interpretation. Neither identification is among the connecting hypotheses of the composition (Table 1), and satisfying them would not measure experience. The unity

proposal is also only as definite as its cover. The gluing theorem holds for whatever cover it is given, a disjoint cover satisfies it trivially, and the cover of interest overlaps in content rather than in space. Section 4 proposes fixing that cover by decodability on held-out data, which excludes covers chosen for agreement and makes non-degeneracy a testable property of the recordings. Its variables, criterion and decoder class are pre-registered by rules that never consult agreement.

10.2 A comparison that could disconfirm the link to content

The framework predicts that conscious content tracks compatibility, not coherence. The comparison must also separate compatibility from how well each region encodes the stimulus, because two accurate decoders agree automatically. If decoders d_A and d_B estimate a shared stimulus quantity x with errors $e_A = d_A - x$ and $e_B = d_B - x$, then

$$|d_A - d_B| \le |e_A| + |e_B|, \qquad |d_A - d_B|^2 = |e_A|^2 + |e_B|^2 - 2\langle e_A, e_B\rangle. \quad (13)$$

Consciously perceived stimuli are usually decoded better, so higher agreement on those trials would follow from single-region fidelity alone. At matched single-region error, agreement differs only through the covariance of the two errors. The comparison adds one premise to link compatibility to that term: two regions describing one content err together, so that compatibility appears

as error covariance beyond what single-region fidelity implies.

Present matched stimuli that are perceived on some trials and not on others, for example at masking or rivalry threshold, and fit decoders for one shared stimulus quantity independently on two regions that both encode it, selected by the decodability criterion of Section 4 on separate data. Classify trials by perceptual state with a no-report measure, such as reflexive eye movements or pupil responses that track the dominant percept in rivalry (Tsuchiya et al., 2015), so that the classification does not also select for the neural correlates of reporting; report-based classification is a secondary analysis. Match the two classes on phase coherence between the regions and on each region's own decoding error (or, where the stimulus value is uncertain, on each region's signal strength). The statistic is the excess agreement over a fidelity-matched null: pair region A's decode on one trial with region B's decode on another trial of the same stimulus and error stratum, which preserves both marginal accuracies and removes any coupling of the errors.

The proposed link predicts a larger excess on perceived trials at matched coherence and matched single-region accuracy (Fig. 5). Equal excess across classes, or excess that follows coherence and not perceptual state, counts against it. A positive result is not yet specific: noise shared upstream of both regions also correlates their errors, so the excess should be compared against an upstream region's decoding error, matched in the same way. A self-model comparison has the same form: at matched accuracy for the scene variable, the reconstruction of perceiver-relative variables, such as hand position rel-

ative to the target, from a declared local encoding should be more accurate on perceived trials. These are correlational tests of proposed correlates, and the no-report classification is itself an indirect marker of experience.

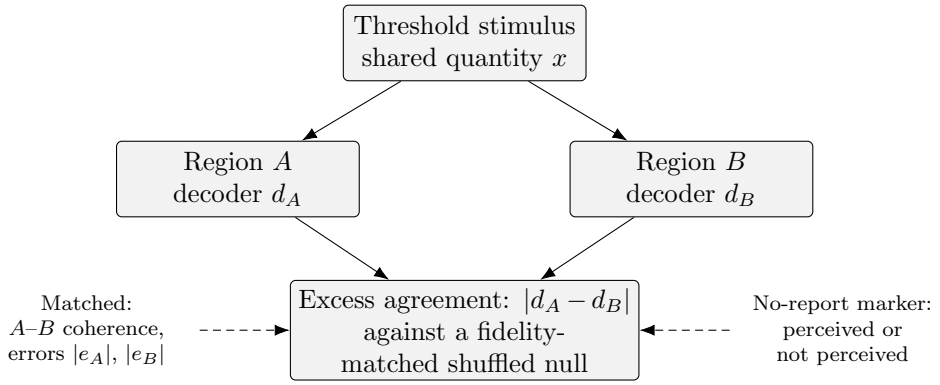


Figure 5: Design of the disconfirming comparison. Decoders for one shared stimulus quantity are fitted independently on two regions. Trials are classed by a no-report marker of perceptual state and matched on inter-regional phase coherence and on each region's own decoding error, so that agreement cannot follow from single-region fidelity (Eq. (13)). The statistic is agreement in excess of a null that pairs decodes across trials within the same stimulus and error stratum.

10.3 What a test of the full proposal must establish

A test needs separate estimates of effective coupling and diffusion, shared content on overlapping regions, and reconstruction from a specified encoding at a declared fidelity. The cover, relevant state family and decoder must be fixed without selecting them for agreement. An independent experiential endpoint is needed to evaluate the interpretation; absence of a report alone

does not define absence of experience. A unified episode with incompatible decoded descriptions would challenge the unity identification. A unified episode lacking the proposed self-encoding would challenge the minimal-self claim. Conversely, satisfying the mathematical conditions without such an episode would challenge their sufficiency as correlates.

For the physical route, cortical actuator modes and prices, the geometry–coupling relation and a dynamical onset model remain to be measured. The field/connectivity comparison in Section 9.3 evaluates one spatial onset prediction. Controlled changes in coordination and independently decoded content agreement would address the representational connection. The synthetic decoder exercise calibrates part of these measurements, and the EEG exercise identifies the estimand a spatial test needs.

Integrated-information, global-workspace and predictive-processing accounts organize unity, access and modelling differently (Tononi et al., 2016; Baars, 1997; Dehaene et al., 2014; Friston, 2010). This account specifies conditional relations among coupling, local content and reconstruction, each tied to a measurement.

11 Limitations

Experience and interpretation. Synchrony is not itself an experiential measure, and nothing here establishes a consciousness criterion: the conditions are proposed correlates to be tested against report or another independent

endpoint. The mathematical states are spatial measures and fields, and identifying them with perceptual content requires independent decoding evidence. The constructions of Section 4 select neither the neural cover nor the decoder, and identifying their contents with perceptual content, or establishing that the identification bears on experience, is a further modelling task. The correspondence between the reconstruction's fixed point and the minimal self is structural and does not show that a fixed point is experienced as a perspective. The remaining explanatory question is whether these relations occur in neural activity and matter for experience; the framework does not explain why satisfying them would be accompanied by experience.

Formal scope. The formal results are conditional on declared encoders, covers and maps whose cortical counterparts are unmeasured. Each edge of the composition is supplied rather than derived from the one above it, so the composition does not establish that self-representation follows from coherence. The joint witnesses satisfy every edge in a specified model: they prove neither that cortex satisfies those hypotheses nor that the glued state is experienced, and establish no common cortical mechanism. The contraction factor Λ is stipulated, not derived. It borrows the rate $(K - 2D)/2$ at which the first harmonic of the uniform density grows (Section 3), a rate about a different state from the one F acts on, and no result derives a Lipschitz law for an encoding or readout from the phase dynamics; constructing a map with that factor remains a modelling obligation. No physical input-conditioned map F_u is constructed, the gain L of Section 5.2 is derived here for no physical

readout, and E89 supplies no resolution of a declared family by the encoding region. Agreement of a common projection does not establish agreement of the full vector, and the general vector result of Section 4 supplies no identification with the spatial-measure sheaf. The cover on which compatibility is assessed is fixed only relative to pre-registered variables, decoder class and criterion, and the phase-to-content bounds are stated for spatial covers rather than for the content cover that criterion yields. No result connects exact or approximate compatibility to correlated decoding errors; the comparison of Section 10.2 takes that link as a premise.

Physical identification. The formal field condition requires a continuous kernel, a normalized substrate, a specified average coupling and noise, and a site-contribution bound. Those conditions do not identify an electromagnetic substrate; the extracellular interpretation needs the separate calibration of Section 9, and the field commitment requires physical evidence beyond a choice of mathematical representation. Phase organization under propofol does not measure coupling, and the support comparisons in Supplemental Material concern specified support constraints and supply neither that evidence nor energy prices. Neither the conversion rate γ nor the installation factor κ follows from a heat bound, and no cortical value of κ is known. Testing the physical threshold requires independent estimates of the interaction kernel, its aggregation and diffusion in the same cortical regime. No measured connectivity yet supplies the spectral route of Section 4.2. The mean-field model has no spatially varying κ and derives no nucleation or-

der, which is why the onset comparison of Section 9.3 rests on a separately specified model.

Numerical and empirical controls. The controls of Section 8 are finite regimes on declared models and calibrate no cortical quantity. A slow control variable supplies no general guarantee of branch tracking; the winding sweep establishes no transition and no biological identification; phase agreement cannot replace an independently decoded content measurement. The synthetic decoder exercise validates a measurement on known constructions, and the EEG exercise checks a measurement pipeline: it identifies what a pooled bipolar estimand cannot measure, and neither exercise identifies a neural content code, tests experience or validates any experiential claim.

12 Conclusions

Phase coherence, compatibility and self-reconstruction are separable conditions. Coherence constrains shared content only through encoders shared on overlaps and with a bound that grows as the square root of the population read. That bound is vacuous beyond 10 sites at the simulated sheet's patch order; chaining it through overlapping patches reaches 0.312 mm of that sheet, and a spectral route removes the population count only for a declared connectivity at a frequency-locked configuration. Across the centimetre separations of distinct cortical territories, the positive link from coherence to content is the bridge assumption E78 rather than a derived result. Com-

patible local descriptions, not synchrony, determine one global state; and a contracting self-map fixes that state while reconstructing only a family of bounded diameter. Each result replaces an assumption about synchrony by a measurement: independently decoded agreement on overlaps, a declared encoding and its fidelity, and calibrated coupling and diffusion. The comparison of Section 10.2 states how such measurements could count against the proposed link to conscious content.

Notation

Table 2: Symbols used throughout, with the section that defines each. Where one letter carries more than one meaning, the meanings are separated by the section they are read in.

Symbol	Defined	Meaning
θ_i, ω_i	Eq. (1)	Phase and natural frequency of oscillator i
N	Eq. (1)	Number of oscillators in the population the quantity is read on
D	Eq. (1), Supp.	Phase diffusion in inverse time; in the Supplemental Material's convergence proof, the phase spread $D \leq \pi/2$
K_{ij}, K	Eq. (1)	Pairwise and scalar mean-field coupling, $K_{ij} = K/N$
K_c	Eq. (2)	Synchronization threshold, $K_c = 2D$
r, r_N	Eq. (1)	Resultant length (order parameter), on the population it is read on
a, ψ	Eq. (2)	Concentration $a = Kr/D$ and collective phase of the stationary density
q	Section 8.3	Winding number of a phase field read around a ring

Table 2 (continued)

Symbol	Defined	Meaning
U_{inst}, κ	Section 3.2	Installed energy of a declared actuator and the mode-data factor bounding its mean coupling by κU_{inst}
L	Section 4.2	Lipschitz constant of an encoder or readout
ε, δ	Section 4.2	Content and observation discrepancy tolerances
M, b	Eq. (9)	Size of a declared family of states and code length in bits
Λ, τ	Eq. (7)	Contraction factor $\Lambda = e^{-(K-2D)\tau/2}$ and its relaxation time; in Eq. (4), τ is the duration of a change
J, σ_J, C	Eq. (4)	Probability current of the phase density, its current functional and a current budget
T	Section 5	Number of communication rounds to a deadline
ν	Supplemental Material	Fan-in of a communication graph: the most sites any one site hears from
γ, $\mathcal{T}$, f	Eq. (10)	Field-to-coupling conversion rate, timing sensitivity and frequency
α	Section 9	Extracellular volume fraction
λ	Section 9	Extracellular tortuosity
Ω, λ_2	Section 4.2	Common rotation rate of a frequency-locked configuration, and the algebraic connectivity of the coupling graph
β, k	Section 4.2	Chord diameter of a patch and the diameter of a cover's nerve, in hops
$\mathcal{C}$, $e_{i,x}$	Section 4.2	Content space and the phase encoder patch i reads at site x
Φ, R	Eq. (9)	Encoder and readout of the reconstruction criterion
$\mathcal{A}$, F	Eq. (7)	Encoding region and the reconstruction map $F(s) = \mathrm{readout}(s\|_{\mathcal{A}})$
E	Eq. (10)	Field amplitude, in the cortical identification
s, s_*	§4–5	A global section and the one the reconstruction fixes

Data and code availability

The Lean development, simulation scripts, tests and saved numerical summaries are maintained in the GitHub repository `bobaseb/coupling_gated_cortical_coherence`. The formalization uses Lean 4.34.0-rc2 and Mathlib commit `5ce203ecd53f8c3e47ee5530d7c86b94a55f0183`. Publication simulation values are generated from saved summaries without rerunning production sweeps. The exploratory EEG analysis uses OpenNeuro dataset ds005620, version 1.0.0 (Bajwa et al., 2024, 2025): subjects `1010`, `1016`, `1022`, `1033`, `1045`, `1054`, `1060`, `1067`, first resting block under sedation. The data, subset and preprocessing are described in Supplemental Material and implemented in `simulations/empirical_collapse.py`. The repository contains the reproducible analysis and does not redistribute the raw recordings.

Funding and competing interests

This research received no external funding. The author declares no competing interests.

Use of AI tools

AI assistance to this work was substantial and is not separable by task. Several coding assistants and chat interfaces were used in an interleaved way throughout: Google's Antigravity CLI (Gemini 3.1 Pro, Claude 4.6 Opus,

Claude 4.6 Sonnet), OpenAI's Codex CLI (GPT-6 Astra and Sol; GPT-5.6 Terra, Luna and Sol), Claude Code (Claude Opus 5 and Opus 5.5), Gemini Pro through its web interface, and Anthropic's Fable and DeepSeek v4 Flash through OpenRouter, accessed with the Hermes CLI. Between them they contributed to the Lean development, the simulation and analysis code, and the drafting and revision of both documents. Because the tools were used in combination and in alternation, no section, proof, figure or result is attributable to any one of them, and none is claimed to be. The author is responsible for all content, including everything a tool produced. Machine checking establishes the stated Lean results under their hypotheses; it does not validate biological interpretations or replace scientific review of generated code and prose.

References

Pinotsis, D. A., & Miller, E. K. (2022). Beyond dimension reduction: Stable electric fields emerge from and allow representational drift. *NeuroImage*, 253, 119058.

Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17(7), 450-461.

Singer, W., & Gray, C. M. (1995). Visual feature integration and the temporal

correlation hypothesis. *Annual Review of Neuroscience*, 18, 555-586.

Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron*, 88(1), 220-235.

Shadlen, M. N., & Movshon, J. A. (1999). Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron*, 24(1), 67-77.

Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. F. (2015). No-report paradigms: extracting the true neural correlates of consciousness. *Trends in Cognitive Sciences*, 19(12), 757-770.

Metzinger, T. (2003). *Being No One: The Self-Model Theory of Subjectivity*. Cambridge, MA: MIT Press.

Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, 17(11), 565-573.

Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., et al. (2023). Consciousness in artificial intelligence: insights from the science of consciousness. *arXiv preprint* arXiv:2308.08708.

Bobadilla-Suarez, S., Ahlheim, C., Mehrotra, A., Panos, A., & Love, B. C. (2020). Measures of neural similarity. *Computational Brain & Behavior*, 3, 369-383.

DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, 11(8), 333-341.

Kriegeskorte, N., & Diedrichsen, J. (2019). Peeling the onion of brain representations. *Annual Review of Neuroscience*, 42, 407-432.

Gallagher, S. (2000). Philosophical conceptions of the self: implications for cognitive science. *Trends in Cognitive Sciences*, 4(1), 14-21.

Baars, B. J. (1997). In the Theater of Consciousness: The Workspace of the Mind. *Oxford University Press*.

Dehaene, S., Charles, L., King, J. R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, 25, 76-84.

Friston, K. (2010). The free-energy principle: a unified brain theory?. *Nature Reviews Neuroscience*, 11(2), 127-138.

Milinkovic, B., & Aru, J. (2026). On biological and artificial consciousness: A case for biological computationalism. *Neuroscience & Biobehavioral Reviews*, 181, 106524. doi:10.1016/j.neubiorev.2025.106524.

Aru, J., Larkum, M. E., & Shine, J. M. (2023). The feasibility of artificial consciousness through the lens of neuroscience. *Trends in Neurosciences*, 46(12), 1008–1017. doi:10.1016/j.tins.2023.09.009.

Seth, A. K. (2026). Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences*, 49, e315. doi:10.1017/S0140525X25000032.

Seth, A. K. (2026). The stuff matters: Consciousness, computation, and biology. *Behavioral and Brain Sciences*, 49, e366. doi:10.1017/S0140525X26106906.

Gurnee, W., et al. (2026). Verbalizable Representations Form a Global Workspace in Language Models. *Transformer Circuits Thread*. arXiv:2607.15495.

Aguilera, M., Millidge, B., Tschantz, A., & Buckley, C. L. (2022). How particular is the physics of the free energy principle?. *Physics of Life Reviews*, 40, 24-50.

Biehl, M., Pollock, F. A., & Kanai, R. (2021). A technical critique of some parts of the free energy principle. *Entropy*, 23(3), 293.

Bruineberg, J., Dolega, K., Dewhurst, J., & Baltieri, M. (2022). The Emperor's New Markov Blankets. *Behavioral and Brain Sciences*, 45, e183.

Pinotsis, D. A., & Miller, E. K. (2023). In vivo ephaptic coupling allows memory network formation. *Cerebral Cortex*, 33(17), 9877-9895.

Pinotsis, D. A., Fridman, G., & Miller, E. K. (2023). Cytoelectric coupling: Electric fields sculpt neural activity and "tune" the brain's infrastructure. *Progress in Neurobiology*, 226, 102465.

McFadden, J. (2002). The conscious electromagnetic information (cemi) field theory: the hard problem made easy?. *Journal of Consciousness Studies*, 9(8), 45-60.

969 Pockett, S. (2012). The electromagnetic field theory of consciousness: a
970 testable hypothesis about the characteristics of conscious as opposed to
971 non-conscious fields. *Journal of Consciousness Studies*, 19(11-12), 191-223.

972 McFadden, J. (2020). Integrating information in the brain's EM field: the
973 cemi field theory of consciousness. *Neuroscience of Consciousness*, 2020(1),
974 niaa016.

975 Fröhlich, F., & McCormick, D. A. (2010). Endogenous electric fields may
976 guide neocortical network activity. *Neuron*, 67(1), 129-143.

977 Anastassiou, C. A., & Koch, C. (2015). Ephaptic coupling to endogenous
978 electric field activity: why bother?. *Current Opinion in Neurobiology*, 31,
979 95-103.

980 Pockett, S. (2002). Difficulties with the electromagnetic field theory of con-
981 sciousness. *Journal of Consciousness Studies*, 9(4), 51-56.

982 Vöröslakos, M., et al. (2018). Direct effects of transcranial electric stimulation
983 on brain circuits in rats and humans. *Nature Communications*, 9, 483.

984 Verkhratsky, A., & Nedergaard, M. (2018). Physiology of Astroglia. *Physio-*
985 *logical Reviews*, 98(1), 239-389.

986 Mountcastle, V. B. (1997). The columnar organization of the neocortex.
987 *Brain*, 120(4), 701-722. doi:10.1093/brain/120.4.701.

Horton, J. C., & Adams, D. L. (2005). The cortical column: a structure without a function. *Philosophical Transactions of the Royal Society B*, 360(1456), 837-862. doi:10.1098/rstb.2005.1623.

Rakic, P. (2008). Confusing cortical columns. *Proceedings of the National Academy of Sciences*, 105(34), 12099-12100. doi:10.1073/pnas.0807271105.

Bastos, A. M., Donoghue, J. A., Brincat, S. L., Mahnke, M., Yanar, J., Correa, J., ... & Miller, E. K. (2021). Neural effects of propofol-induced unconsciousness and its reversal using thalamic stimulation. *eLife*, 10, e60824.

Bajwa, I. J., Nilsen, A. S., Skukies, R., Aamodt, A., Ernst, G., Storm, J. F., & Juel, B. E. (2025). A repeated awakening study exploring the capacity of complexity measures to capture dreaming during propofol sedation. *Scientific Reports*, 15, 32746. doi:10.1038/s41598-025-12695-z.

Bajwa, I. J., Nilsen, A. S., Skukies, R., Aamodt, A., Ernst, G., Storm, J. F., & Juel, B. E. (2024). A repeated awakening study exploring the capacity of complexity measures to capture dreaming during propofol sedation (Version 1.0.0) [Data set]. *OpenNeuro*. doi:10.18112/openneuro.ds005620.v1.0.0.

Xie, L., Kang, H., Xu, Q., Chen, M. J., Liao, Y., Thiyagarajan, M., ... & Nedergaard, M. (2013). Sleep drives metabolite clearance from the adult brain. *Science*, 342(6156), 373-377.

Syková, E., & Nicholson, C. (2008). Diffusion in brain extracellular space. *Physiological Reviews*, 88(4), 1277-1340.

Barbour, B. (2017). Analysis of claims that the brain extracellular impedance is high and non-resistive. *Biophysical Journal*, 113(7), 1636-1638.

Logothetis, N. K., Kayser, C., & Oeltermann, A. (2007). In vivo measurement of cortical impedance spectrum in monkeys: implications for signal propagation. *Neuron*, 55(5), 809–823. doi:10.1016/j.neuron.2007.07.027.

Gabriel, S., Lau, R. W., & Gabriel, C. (1996). The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz. *Physics in Medicine and Biology*, 41(11), 2251–2269. doi:10.1088/0031-9155/41/11/002.

Norton, J. D. (2005). Eaters of the lotus: Landauer's principle and the return of Maxwell's demon. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 36(2), 375-411.

Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183-191.

Kibble, T. W. B. (1976). Topology of cosmic domains and strings. *Journal of Physics A: Mathematical and General*, 9(8), 1387-1398.

Sakaguchi, H. (1988). Cooperative phenomena in coupled oscillator systems under external fields. *Progress of Theoretical Physics*, 79(1), 39-46.

Strogatz, S. H., & Mirollo, R. E. (1991). Stability of incoherence in a population of coupled oscillators. *Journal of Statistical Physics*, 63, 613–635.

Acebrón, J. A., Bonilla, L. L., Pérez Vicente, C. J., Ritort, F., & Spigler, R. (2005). The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77, 137–185. doi:10.1103/RevModPhys.77.137.

Ott, E., & Antonsen, T. M. (2008). Low dimensional behavior of large systems of globally coupled oscillators. *Chaos*, 18, 037113. doi:10.1063/1.2930766.

Dörfler, F., & Bullo, F. (2014). Synchronization in complex networks of phase oscillators: A survey. *Automatica*, 50, 1539–1564. doi:10.1016/j.automatica.2014.04.012.

Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature Neuroscience*, 20, 340–352. doi:10.1038/nn.4497.

Still, S., Sivak, D. A., Bell, A. J., & Crooks, G. E. (2012). Thermodynamics of prediction. *Physical Review Letters*, 109(12), 120604.

Kawai, R., Parrondo, J. M. R., & Van den Broeck, C. (2007). Dissipation: The phase-space perspective. *Physical Review Letters*, 98(8), 080602.

Horowitz, J. M., & Esposito, M. (2014). Thermodynamics with continuous information flow. *Physical Review X*, 4, 031015. doi:10.1103/PhysRevX.4.031015.

Ito, S., & Sagawa, T. (2013). Information thermodynamics on causal networks. *Physical Review Letters*, 111, 180603. doi:10.1103/PhysRevLett.111.180603.

Corberi, F., dello Russo, S., Messuti, G., Scarpetta, S., & Smaldone, L. (2026). Thermodynamic learning. *arXiv preprint*. arXiv:2609.04732.

Banach, S. (1922). Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta Mathematicae*, 3(1), 133-181.

Shenker, O. R. (2000). Logic and entropy. *PhilSci-Archive preprint*. philsci-archive.pitt.edu/115.

Rosas, F. E., Mediano, P. A. M., Luppi, A. I., Boyd, A. B., Berta, M., Jensen, H. J., Gastpar, M., & Carroll, S. M. (2026). Four faces of information in natural systems. *Nature Reviews Physics*. doi:10.1038/s42254-026-00971-4.

Sznitman, A.-S. (1991). Topics in propagation of chaos. In *École d'Été de Probabilités de Saint-Flour XIX—1989*, Lecture Notes in Mathematics, vol. 1464, pp. 165–251. Springer. doi:10.1007/BFb0085169.

Dai Pra, P., & den Hollander, F. (1996). McKean–Vlasov limit for interacting random processes in random media. *Journal of Statistical Physics*, 84, 735–772. doi:10.1007/BF02179656.

Bertini, L., Giacomin, G., & Pakdaman, K. (2010). Dynamical aspects of

mean field plane rotators and the Kuramoto model. *Journal of Statistical Physics*, 138, 270–290. doi:10.1007/s10955-009-9908-9.

Berglund, N., & Gentz, B. (2002). Pathwise description of dynamic pitchfork bifurcations with additive noise. *Probability Theory and Related Fields*, 122, 341–388. doi:10.1007/s004400100174.

Abramsky, S., & Brandenburger, A. (2011). The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13, 113036. doi:10.1088/1367-2630/13/11/113036.

Townsend, R. G., Solomon, S. S., Chen, S. C., Pietersen, A. N. J., Martin, P. R., Solomon, S. G., & Gong, P. (2015). Emergence of complex wave patterns in primate cerebral cortex. *The Journal of Neuroscience*, 35(11), 4657-4662. doi:10.1523/JNEUROSCI.4509-14.2015.

Xu, Y., Long, X., Feng, J., & Gong, P. (2023). Interacting spiral wave patterns underlie complex brain dynamics and are related to cognitive processing. *Nature Human Behaviour*, 7(7), 1196-1215. doi:10.1038/s41562-023-01626-5.

Nicoletti, G., Di Terlizzi, I., & Busiello, D. M. (2026). Balancing information and dissipation with partially observed fluctuating signals. *Physical Review Letters*, 137(12), 127101. doi:10.1103/8gx2-rbns.

Kiyooka, D., Oomoto, I., Kitazono, J., Saito, Y., Kobayashi, M., Matsubara, C., Kobayashi, K., Murayama, M., & Oizumi, M. (2026). Single-

cell resolution functional networks during unconsciousness are segregated into spatially intermixed modules. *Cell Reports*, 45(2), 116902. doi:10.1016/j.celrep.2025.116902.

Oomoto, I., Kiyooka, D., Oizumi, M., & Murayama, M. (2026). Multi-area single-cell calcium imaging dataset of the mouse cortex across wakefulness, sleep, and anesthesia. *Scientific Data*. doi:10.1038/s41597-026-08202-2.

Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism*, 21(10), 1133-1145. doi:10.1097/00004647-200110000-00001.

Harris, J. J., Jolivet, R., & Attwell, D. (2012). Synaptic energy use and supply. *Neuron*, 75(5), 762-777. doi:10.1016/j.neuron.2012.08.019.

Nicholls, D. G., & Ferguson, S. J. (2013). *Bioenergetics 4*. Academic Press.

Baer, S. M., Erneux, T., & Rinzel, J. (1989). The slow passage through a Hopf bifurcation: delay, memory effects, and resonance. *SIAM Journal on Applied Mathematics*, 49(1), 55–71. doi:10.1137/0149003.

Friedman, E. B., Sun, Y., Moore, J. T., Hung, H.-T., Meng, Q. C., Perera, P., Joiner, W. J., Thomas, S. A., Eckenhoff, R. G., Sehgal, A., & Kelz, M. B. (2010). A conserved behavioral state barrier impedes transitions between anesthetic-induced unconsciousness and wakefulness: evidence for neural inertia. *PLoS ONE*, 5(7), e11903. doi:10.1371/journal.pone.0011903.

Hudson, A. E., Calderon, D. P., Pfaff, D. W., & Proekt, A. (2014). Recovery of consciousness is mediated by a network of discrete metastable activity states. *Proceedings of the National Academy of Sciences*, 111(25), 9283–9288. doi:10.1073/pnas.1408296111.

Proekt, A., & Hudson, A. E. (2018). A stochastic basis for neural inertia in emergence from general anaesthesia. *British Journal of Anaesthesia*, 121(1), 86–94. doi:10.1016/j.bja.2018.02.035.

Bobadilla-Suarez, S. (2026). Unity requires physically enforced agreement. *PsyArXiv*. doi:10.31234/osf.io/dkqep_v1.