TUC-N

A Graph-Dynamical Multiscale Model of Candidate Neural States of Consciousness

Mathematical Foundations, Computational Architecture, and Falsifiable Neuroscientific Predictions

Alessio De Angelis

28 June 2026

Version 1.0 — Foundational Specification

Abstract

The neuroscience of consciousness requires formal models that can state, rather than merely describe, which properties of neural activity are hypothesized to be relevant to conscious access and reported content. TUC-N introduces a graph-dynamical framework in which neural activity is represented as a latent state evolving on a weighted neural graph. The model is organized around four elements: a reference neural state, denoted by $\mathbf{\Phi}_0$; the deviation from that reference, $\delta\mathbf{\Phi}(t)$; a normalized latent field, $\mathbf{c}(t)$; and a multiscale architecture combining global state descriptors with temporally coherent local coalitions. The present article specifies the canonical mathematical core of TUC-N. It defines the neural state space, the graph representation, the normalized graph Laplacian, the operational role of PhiZero, and the decomposition of the field into uniform and structurally differentiated components. These definitions support a stable linear graph-dynamical model, global measures of distance, spatial differentiation, temporal activity, and structural persistence, and local measures of coalition coherence and stability. The model is designed to generate falsifiable predictions through explicit observation models and out-of-sample comparisons against prespecified baseline models. The central methodological aim is to provide a coherent state-space language in which candidate neural conditions can be defined, estimated, perturbed, and tested across modalities. This article establishes the formal specification required for mathematical analysis, simulation-based calibration, and empirical evaluation in subsequent work.

Keywords: consciousness; neural state-space modelling; graph dynamics; network neuroscience; latent neural states; structural persistence; functional coalitions; identifiability.

Version note. This document is designated **Version 1.0**. It supersedes earlier internal drafts by incorporating the canonical energy normalization, threshold semantics, sampled-data observability criterion, state-reduction conditions, and computational-verification requirements stated here.

Contents

1 Introduction

A central challenge in consciousness research is the construction of models that are conceptually explicit, mathematically tractable, empirically testable, and compatible with heterogeneous neural measurements. Relevant data are multiscale by construction: they may include activity recorded from neurons, cortical sources, intracranial electrodes, electroencephalography, magnetoencephalography, functional magnetic resonance imaging, or derived regional time series. These measurements differ in temporal resolution, spatial resolution, signal-to-noise ratio, and relation to underlying neural processes. A useful formal framework must therefore distinguish the latent neural state from the observation process through which that state is measured.

Current theories of consciousness have generated influential proposals concerning widespread access, recurrent interactions, information integration, higher-order representation, predictive processing, and dynamical organization. Their empirical evaluation has increasingly emphasized preregistered predictions, explicit contrasts among theories, and the distinction between neural activity associated with perceptual content, behavioural report, task demands, and post-perceptual processing [1, 2]. This methodological development creates a need for models that state clearly what is being computed from neural data, which parameters are fixed before confirmation, and which predictions would count against the model.

A recurrent difficulty is that neural descriptions are often either predominantly global or predominantly local. Global descriptions can summarize broad properties of the neural state, such as overall magnitude, variance, spatial differentiation, temporal variability, or network-wide integration. They are useful because conscious access is frequently studied in relation to distributed neural processes and broad changes in brain-state organization. Yet global summaries can discard the local structure through which distinct functional configurations are implemented. Conversely, local descriptions can identify subnetworks, communities, assemblies, or spatially restricted patterns, but may lack a principled account of the larger neural state in which those local configurations are embedded.

The TUC-N framework addresses this problem through a graph-dynamical representation of neural activity. It does not assume that a single scalar quantity is sufficient to characterize a candidate neural condition. Instead, it represents neural activity as a time-dependent state vector on a weighted graph and distinguishes several mathematically separate properties of that state. These include deviation from a reference configuration, structural differentiation across graph edges, temporal variation, persistence of structural organization, and the presence of local coalitions with coherent temporal dynamics. The model thereby treats global organization and local functional structure as complementary rather than competing levels of description.

The framework rests on five methodological commitments. First, it is formulated in terms of a latent neural state rather than raw sensor measurements alone. This is necessary because observed neural signals are generated through modality-specific measurement processes. A sensor-level EEG voltage, an fMRI blood-oxygen-level-dependent signal, and an intracranial local field potential cannot be assumed to be interchangeable coordinates of the same neural state. State-space modelling provides a language for separating latent dynamics from observation processes and for estimating temporal structure under uncertainty [3, 4].

Second, neural relations are represented by a weighted graph. Graph-based models describe distributed neural systems while preserving distinctions among nodes, edges, coupling structure, and multiscale organization. The meaning of an edge must be specified carefully because it may represent structural connectivity, a symmetric functional association, anatomical proximity, or an independently defined similarity rule. Different edge definitions can lead to different network inferences; graph construction is therefore part of the model specification rather than a neutral preprocessing detail [5].

Third, the primary dynamical model is deliberately minimal. The canonical TUC-N dynamics use a stable linear graph-diffusive system with an external input term. The linear model is not intended to capture every possible neural nonlinearity. Its role is to provide a baseline that is analytically transparent, identifiable under explicit conditions, and suitable for controlled simulation. More complex nonlinear or stochastic variants become scientifically warranted only if they provide demonstrable advantages in prespecified model comparison, parameter recovery, predictive performance, or explanatory scope.

Fourth, TUC-N separates structural repertoire from momentary activation. A neural system may possess a stable family of candidate coalitions without every such coalition being active at a given time. The framework therefore distinguishes a structural repertoire, $\Sigma_\Phi^{\mathrm{str}}$, from the time-dependent family of active coalitions, $\Sigma_\Phi^{\mathrm{act}}(t)$. This distinction prevents the disappearance of an active coalition from being conflated with the disappearance of the underlying graph-local organization from which that coalition can re-emerge.

Fifth, the model is constructed to be falsifiable. Its global and local features are not intended to be self-validating descriptors. They must improve prespecified predictions of independent endpoints, such as conscious access and reported content, relative to appropriate baseline models. The relevant empirical question is not whether a feature can be computed, but whether it adds predictive information beyond simpler alternatives under held-out evaluation.

The purpose of the present article is to define the canonical mathematical core of TUC-N. Sections 2 and 3 introduce the reference state $\mathbf{\Phi}_0$, the deviation field $\delta\mathbf{\Phi}(t)$, the normalized state $\mathbf{c}(t)$, the graph $\mathcal{G} = (V, W)$, and the normalized Laplacian $\widehat{L}$. Sections 4 to 6 specify the primary dynamics and global descriptors. Sections 7 and 8 define the local coalition architecture and the multiscale operational condition. The final sections state the observation model, reduced-state representation, identifiability conditions, synthetic verification protocol, and empirical prediction framework.

The contribution is methodological and foundational. The present work does not claim completed empirical validation. Instead, it specifies a formal object that can be analysed mathematically, evaluated in synthetic data, reconstructed from neural measurements under explicit observability conditions, and tested against preregistered behavioural or phenomenological endpoints.

2 Foundations of TUC-N

2.1 Latent neural state

Let

$$\mathbf{\Phi}(t) = \begin{pmatrix} \Phi_1(t) \\ \vdots \\ \Phi_N(t) \end{pmatrix} \in \mathbb{R}^N \tag{1}$$

denote the latent neural state at time t. The index $i \in \{1, \ldots, N\}$ refers to a node in a predefined neural representation. Depending on the experimental modality and modelling choices, a node may correspond to a cortical parcel, a reconstructed source, an intracranial contact, an EEG or MEG channel, a voxel, a neuronal population, or another predefined unit of analysis.

The vector $\mathbf{\Phi}(t)$ is not assumed to be a direct copy of the observed measurement vector. Rather, it represents the state in which the model is formulated. The relation between $\mathbf{\Phi}(t)$ and empirical measurements is specified later through an explicit observation model.

The temporal domain is denoted by

$$t \in \mathcal{T}, \qquad \mathcal{T} \subseteq \mathbb{R}_{\geq 0}.$$

Unless otherwise stated, all state trajectories are assumed sufficiently regular for the operations used in the corresponding section.

2.2 PhiZero as neural reference state

The first canonical object in TUC-N is the neural reference state,

$$\mathbf{\Phi}_0 = \begin{pmatrix} \Phi_{0,1} \\ \vdots \\ \Phi_{0,N} \end{pmatrix} \in \mathbb{R}^N. \tag{2}$$

Within a given analysis, $\mathbf{\Phi}_0$ is fixed. It defines the reference configuration relative to which deviations, spatial differentiation, temporal evolution, and multiscale organization are evaluated. The reference state must therefore be specified or estimated independently of the data used for confirmatory evaluation.

Let $\mathcal{T}_{\mathrm{cal}}$, $\mathcal{T}_{\mathrm{train}}$, and $\mathcal{T}_{\mathrm{test}}$ denote respectively calibration, model-development, and confirmatory evaluation sets. The minimum separation condition is

$$\mathcal{T}_{\mathrm{cal}} \cap \mathcal{T}_{\mathrm{test}} = \varnothing. \tag{3}$$

Whenever model parameters or thresholds are estimated on training data, the additional separation conditions are

$$\mathcal{T}_{\mathrm{train}} \cap \mathcal{T}_{\mathrm{test}} = \varnothing, \qquad \mathcal{T}_{\mathrm{cal}} \cap \mathcal{T}_{\mathrm{train}} = \varnothing. \tag{4}$$

One possible estimator of the reference state is the calibration average,

$$\widehat{\mathbf{\Phi}}_0 = \frac{1}{K_{\mathrm{cal}}} \sum_{k=1}^{K_{\mathrm{cal}}} \widehat{\mathbf{\Phi}}(t_k), \qquad t_k \in \mathcal{T}_{\mathrm{cal}}. \tag{5}$$

Equation (5) is an example rather than a universal requirement. Alternative estimators may be used when justified by the experimental design, provided that the estimation rule is fixed before confirmatory analysis.

2.3 DeltaPhi and the normalized neural field

The fundamental deviation field is

$$\delta\mathbf{\Phi}(t) = \mathbf{\Phi}(t) - \mathbf{\Phi}_0. \tag{6}$$

At the nodal level,

$$\delta\Phi_i(t) = \Phi_i(t) - \Phi_{0,i}, \qquad i = 1, \ldots, N. \tag{7}$$

The inverse relation is

$$\mathbf{\Phi}(t) = \mathbf{\Phi}_0 + \delta\mathbf{\Phi}(t). \tag{8}$$

The condition

$$\delta\mathbf{\Phi}(t) = \mathbf{0} \tag{9}$$

is equivalent to

$$\mathbf{\Phi}(t) = \mathbf{\Phi}_0. \tag{10}$$

This equality identifies coincidence with the reference state. It does not replace the multiscale operational condition introduced later.

To permit dimensionless comparisons within a fixed analysis, TUC-N introduces a positive calibration scale,

$$\phi_\star > 0. \tag{11}$$

The normalized neural state field is

$$\mathbf{c}(t) = \frac{\delta\mathbf{\Phi}(t)}{\phi_\star} = \frac{\mathbf{\Phi}(t) - \mathbf{\Phi}_0}{\phi_\star}. \tag{12}$$

Equivalently,

$$\delta\mathbf{\Phi}(t) = \phi_\star \mathbf{c}(t), \qquad \mathbf{\Phi}(t) = \mathbf{\Phi}_0 + \phi_\star \mathbf{c}(t). \tag{13}$$

A possible root-mean-square calibration is

$$\phi_\star^2 = \frac{1}{NK_{\text{cal}}} \sum_{k=1}^{K_{\text{cal}}} \left\| \widehat{\mathbf{\Phi}}(t_k) - \widehat{\mathbf{\Phi}}_0 \right\|_2^2. \tag{14}$$

The estimator of $\phi_\star$ must be fixed before confirmatory analyses and retained throughout all feature calculations, threshold definitions, and model comparisons.

2.4 Reference-relative interpretation

The triplet

$$(\mathbf{\Phi}_0, \delta\mathbf{\Phi}(t), \mathbf{c}(t)) \tag{15}$$

provides the minimal state representation required by the model. Its components denote respectively the reference neural configuration, the neural deviation relative to the reference, and the normalized state used in the canonical dynamics. This reference-relative formulation permits global displacement, structural differentiation, temporal variation, and local functional organization to be analysed as distinct quantities.

3 Neural State Space, Graph Representation, and Discrete Geometry

3.1 Weighted neural graph

TUC-N represents the neural system by a weighted undirected graph,

$$\mathcal{G} = (V, W), \qquad V = \{1, \dots, N\}. \tag{16}$$

The adjacency matrix $W = (W_{ij})_{i,j=1}^N$ satisfies

$$W = W^\top, \qquad W_{ij} \geq 0, \qquad W_{ii} = 0. \tag{17}$$

The primary TUC-N specification assumes

$$N \geq 2, \qquad \mathcal{G} \text{ is connected.} \tag{18}$$

The construction rule for W, including its biological or statistical interpretation, must be specified before confirmatory evaluation.

3.2 Graph Laplacian and canonical normalization

The weighted degree of node i is

$$d_i = \sum_{j=1}^N W_{ij}, \qquad D_W = \operatorname{diag}(d_1, \dots, d_N). \tag{19}$$

The combinatorial graph Laplacian is

$$L = D_W - W. \tag{20}$$

For every $\mathbf{z} \in \mathbb{R}^N$,

$$\mathbf{z}^\top L \mathbf{z} = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N W_{ij} (z_i - z_j)^2 \geq 0. \tag{21}$$

Thus $L \succeq 0$, and connectedness implies

$$\ker(L) = \operatorname{span}\{\mathbf{1}_N\}. \tag{22}$$

The canonical normalized Laplacian is

$$\widehat{L} = \frac{L}{\lambda_{\max}(L)}. \tag{23}$$

Since the graph is nontrivial and connected, $\lambda_{\max}(L) > 0$. Therefore,

$$\widehat{L} \succeq 0, \qquad \lambda_{\max}(\widehat{L}) = 1. \tag{24}$$

3.3 Incidence representation and discrete differentiation

Let E denote the edge set of $\mathcal{G}$, choose an arbitrary orientation of the edges, and let $B_{\mathcal{G}} \in \mathbb{R}^{N\times|E|}$ be the oriented incidence matrix. If $W_E = \operatorname{diag}(w_e : e \in E)$, then

$$L = B_{\mathcal{G}} W_E B_{\mathcal{G}}^{\top}. \tag{25}$$

The graph gradient of $\mathbf{z} \in \mathbb{R}^N$ is

$$\nabla_{\mathcal{G}}\mathbf{z} = B_{\mathcal{G}}^{\top}\mathbf{z}, \qquad [\nabla_{\mathcal{G}}\mathbf{z}]_{(i\to j)} = z_i - z_j. \tag{26}$$

The weighted graph gradient of the neural deviation field is

$$\nabla_{\mathcal{G},W}\delta\boldsymbol{\Phi}(t) = W_E^{1/2} B_{\mathcal{G}}^{\top}\delta\boldsymbol{\Phi}(t), \tag{27}$$

and it satisfies

$$\|\nabla_{\mathcal{G},W}\delta\boldsymbol{\Phi}(t)\|_2^2 = \delta\boldsymbol{\Phi}(t)^{\top} L\delta\boldsymbol{\Phi}(t). \tag{28}$$

The normalized diffusive operator is

$$\Delta_{\mathcal{G}}\mathbf{z} = -\widehat{L}\mathbf{z}. \tag{29}$$

3.4 Uniform and structurally differentiated components

Define the orthogonal projector

$$P_{\perp} = I_N - \frac{1}{N}\mathbf{1}_N\mathbf{1}_N^{\top}. \tag{30}$$

The graph-wide mean is

$$\bar{c}(t) = \frac{1}{N}\mathbf{1}_N^{\top}\mathbf{c}(t), \tag{31}$$

the uniform component is $\mathbf{c}_{\parallel}(t) = \bar{c}(t)\mathbf{1}_N$, and the structurally differentiated component is

$$\widetilde{\mathbf{c}}(t) = P_{\perp}\mathbf{c}(t). \tag{32}$$

Thus

$$\mathbf{c}(t) = \mathbf{c}_{\parallel}(t) + \widetilde{\mathbf{c}}(t), \qquad \mathbf{1}_N^{\top}\widetilde{\mathbf{c}}(t) = 0, \tag{33}$$

and

$$\|\mathbf{c}(t)\|_2^2 = N\bar{c}(t)^2 + \|\widetilde{\mathbf{c}}(t)\|_2^2. \tag{34}$$

Since $\widehat{L}\mathbf{1}_N = \mathbf{0}$,

$$\mathbf{c}(t)^{\top}\widehat{L}\mathbf{c}(t) = \widetilde{\mathbf{c}}(t)^{\top}\widehat{L}\widetilde{\mathbf{c}}(t). \tag{35}$$

3.5 Spectral representation

Because $\widehat{L}$ is real and symmetric, it admits an orthonormal eigendecomposition,

$$\widehat{L} = U\Lambda U^\top, \qquad U = (\mathbf{v}_1, \dots, \mathbf{v}_N), \qquad \Lambda = \mathrm{diag}(\lambda_1, \dots, \lambda_N). \tag{36}$$

For a connected graph,

$$0 = \lambda_1 < \lambda_2 \le \cdots \le \lambda_N = 1, \qquad \mathbf{v}_1 = \frac{1}{\sqrt{N}}\mathbf{1}_N. \tag{37}$$

The state expansion is

$$\mathbf{c}(t) = \sum_{m=1}^{N} a_m(t)\mathbf{v}_m, \qquad a_m(t) = \mathbf{v}_m^\top \mathbf{c}(t), \tag{38}$$

and

$$\mathbf{c}(t)^\top \widehat{L}\mathbf{c}(t) = \sum_{m=1}^{N} \lambda_m a_m(t)^2. \tag{39}$$

4 Primary Graph Dynamics and Stability

4.1 Canonical dynamical specification

The canonical TUC-N dynamics are

$$\dot{\mathbf{c}}(t) = -\kappa_G \widehat{L}\mathbf{c}(t) - \mu\mathbf{c}(t) + B\mathbf{u}(t), \tag{40}$$

where

$$\kappa_G > 0, \qquad \mu > 0, \qquad B \in \mathbb{R}^{N\times q}, \qquad \mathbf{u}(t) \in \mathbb{R}^q. \tag{41}$$

Define

$$A = \kappa_G \widehat{L} + \mu I_N. \tag{42}$$

Then

$$\dot{\mathbf{c}}(t) = -A\mathbf{c}(t) + B\mathbf{u}(t). \tag{43}$$

For every nonzero $\mathbf{z} \in \mathbb{R}^N$,

$$\mathbf{z}^\top A\mathbf{z} = \kappa_G \mathbf{z}^\top \widehat{L}\mathbf{z} + \mu\|\mathbf{z}\|_2^2 \ge \mu\|\mathbf{z}\|_2^2 > 0. \tag{44}$$

Hence $A \succ 0$. In DeltaPhi coordinates,

$$\delta\dot{\mathbf{\Phi}}(t) = -A\delta\mathbf{\Phi}(t) + \phi_\star B\mathbf{u}(t), \tag{45}$$

and in full-state coordinates,

$$\dot{\mathbf{\Phi}}(t) = -A[\mathbf{\Phi}(t) - \mathbf{\Phi}_0] + \phi_\star B\mathbf{u}(t). \tag{46}$$

4.2 Existence, uniqueness, and exact solution

Assume $\mathbf{u} \in L^1_{\mathrm{loc}}([0,\infty); \mathbb{R}^q)$ and $\mathbf{c}(0) = \mathbf{c}_0 \in \mathbb{R}^N$.

Proposition 4.1 (Existence and uniqueness)**.** *For every initial state $\mathbf{c}_0 \in \mathbb{R}^N$ and every locally integrable input $\mathbf{u}$, (43) has a unique absolutely continuous solution on $[0,\infty)$.*

The exact solution is

$$\mathbf{c}(t) = e^{-At}\mathbf{c}_0 + \int_0^t e^{-A(t-s)}B\mathbf{u}(s)\,ds. \tag{47}$$

Consequently,

$$\delta\mathbf{\Phi}(t) = e^{-At}\delta\mathbf{\Phi}(0) + \phi_\star \int_0^t e^{-A(t-s)}B\mathbf{u}(s)\,ds. \tag{48}$$

4.3 Closed-system relaxation and Lyapunov analysis

For the closed system $\mathbf{u}(t) \equiv \mathbf{0}$,

$$\mathbf{c}(t) = e^{-At}\mathbf{c}(0). \tag{49}$$

Since $A \succeq \mu I_N$,

$$\|e^{-At}\|_2 \le e^{-\mu t}, \qquad \|\mathbf{c}(t)\|_2 \le e^{-\mu t}\|\mathbf{c}(0)\|_2. \tag{50}$$

Proposition 4.2 (Global exponential stability)**.** *Under the closed-system dynamics,* $\mathbf{c} = \mathbf{0}$ *is globally exponentially stable. Consequently,*

$$\lim_{t\to\infty} \mathbf{c}(t) = \mathbf{0}, \qquad \lim_{t\to\infty} \delta\mathbf{\Phi}(t) = \mathbf{0}, \qquad \lim_{t\to\infty} \mathbf{\Phi}(t) = \mathbf{\Phi}_0. \tag{51}$$

A Lyapunov function is

$$\mathcal{V}_0(t) = \frac{1}{2}\|\mathbf{c}(t)\|_2^2 = \frac{1}{2\phi_\star^2}\|\delta\mathbf{\Phi}(t)\|_2^2. \tag{52}$$

Along closed-system trajectories,

$$\frac{d\mathcal{V}_0}{dt} = -\kappa_G \mathbf{c}(t)^\top \widehat{L}\mathbf{c}(t) - \mu\|\mathbf{c}(t)\|_2^2 \le -2\mu\mathcal{V}_0(t). \tag{53}$$

4.4 Bounded, constant, and periodic inputs

If $\|\mathbf{u}(t)\|_2 \le U_{\max}$ for $t \ge 0$, then

$$\|\mathbf{c}(t)\|_2 \le e^{-\mu t}\|\mathbf{c}(0)\|_2 + \frac{\|B\|_2}{\mu}(1 - e^{-\mu t})U_{\max}. \tag{54}$$

Thus,

$$\limsup_{t\to\infty} \|\mathbf{c}(t)\|_2 \le \frac{\|B\|_2}{\mu}U_{\max}. \tag{55}$$

For constant input $\mathbf{u}(t) = \mathbf{u}_0$, the unique stationary state is

$$\mathbf{c}_{\mathrm{ss}} = A^{-1}B\mathbf{u}_0, \qquad \delta\mathbf{\Phi}_{\mathrm{ss}} = \phi_\star A^{-1}B\mathbf{u}_0. \tag{56}$$

For a periodic input $\mathbf{u}(t + T_u) = \mathbf{u}(t)$, stability yields a unique periodic response

$$\mathbf{c}_{\mathrm{per}}(t) = \int_{-\infty}^{t} e^{-A(t-s)}B\mathbf{u}(s)\, ds. \tag{57}$$

The primary model therefore does not posit autonomous sustained oscillations in the absence of forcing.

4.5 Spectral dynamics and exact discretization

Projection onto $\mathbf{v}_m$ gives

$$\dot{a}_m(t) = -(\mu + \kappa_G\lambda_m)a_m(t) + \mathbf{v}_m^\top B\mathbf{u}(t). \tag{58}$$

In the closed system,

$$a_m(t) = e^{-(\mu+\kappa_G\lambda_m)t}a_m(0). \tag{59}$$

At sampling times $t_k = k\Delta t$, under zero-order hold input $\mathbf{u}(t) = \mathbf{u}_k$ on $[t_k, t_{k+1})$, the exact discrete model is

$$\mathbf{c}_{k+1} = F\mathbf{c}_k + G_c\mathbf{u}_k, \qquad F = e^{-A\Delta t}, \qquad G_c = A^{-1}(I_N - F)B. \tag{60}$$

5 Global Reference-Relative Descriptors

5.1 Distance, dispersion, and homogeneity

The root-mean-square distance from the reference is

$$\rho_\Phi(t) = \left[\frac{1}{N}\|\mathbf{c}(t)\|_2^2\right]^{1/2} = \frac{1}{\phi_\star}\left[\frac{1}{N}\|\delta\mathbf{\Phi}(t)\|_2^2\right]^{1/2}. \tag{61}$$

For a prespecified $\rho_\star > 0$, define

$$d_\Phi(t) = \frac{\rho_\Phi(t)}{\rho_\Phi(t)+\rho_\star}, \qquad A_\Phi(t) = \frac{\rho_\star}{\rho_\Phi(t)+\rho_\star} = 1 - d_\Phi(t). \tag{62}$$

The global structural dispersion is

$$V_{\mathrm{glob}}(t) = \frac{1}{N}\|\widetilde{\mathbf{c}}(t)\|_2^2. \tag{63}$$

The norm decomposition implies

$$\rho_\Phi(t)^2 = \bar{c}(t)^2 + V_{\mathrm{glob}}(t). \tag{64}$$

For $V_\star > 0$, define global homogeneity as

$$A_{\mathrm{hom}}(t) = \exp\left[-\frac{V_{\mathrm{glob}}(t)}{V_\star}\right], \qquad 0 < A_{\mathrm{hom}}(t) \le 1. \tag{65}$$

5.2 Spatial, temporal, and curvature descriptors

The canonical global spatial activity is

$$E_S^{\mathrm{glob}}(t) = \frac{1}{N}\mathbf{c}(t)^\top \widehat{L}\mathbf{c}(t). \tag{66}$$

Its edgewise form is

$$E_S^{\mathrm{glob}}(t) = \frac{1}{2N\lambda_{\max}(L)}\sum_{i=1}^N\sum_{j=1}^N W_{ij}[c_i(t)-c_j(t)]^2, \tag{67}$$

and in DeltaPhi coordinates,

$$E_S^{\mathrm{glob}}(t) = \frac{1}{2N\lambda_{\max}(L)\phi_\star^2}\sum_{i=1}^N\sum_{j=1}^N W_{ij}[\delta\Phi_i(t)-\delta\Phi_j(t)]^2. \tag{68}$$

Since the spectrum of $\widehat{L}$ lies in $[0,1]$,

$$0 \le E_S^{\mathrm{glob}}(t) \le V_{\mathrm{glob}}(t). \tag{69}$$

This normalization is invariant under a global rescaling $W \mapsto \alpha W$, $\alpha > 0$, because $\widehat{L}$ is invariant under that rescaling.

For a prespecified temporal scale $\tau_G > 0$, global temporal activity is

$$E_T^{\mathrm{glob}}(t) = \frac{\tau_G^2}{N}\|\dot{\mathbf{c}}(t)\|_2^2 = \frac{\tau_G^2}{N}\|-A\mathbf{c}(t)+B\mathbf{u}(t)\|_2^2. \tag{70}$$

A graph-relative curvature-like descriptor is

$$\boldsymbol{\kappa}_\Phi(t) = \widehat{L}\mathbf{c}(t), \qquad K_\Phi^{\mathrm{glob}}(t) = \frac{1}{N}\|\widehat{L}\mathbf{c}(t)\|_2^2. \tag{71}$$

Because the spectrum of $\widehat{L}$ lies in $[0,1]$,

$$0 \le K_\Phi^{\mathrm{glob}}(t) \le V_{\mathrm{glob}}(t). \tag{72}$$

The curvature diagnostic is optional and is not part of the minimal global activation criterion.

6 Global Structural Persistence

Let $T_R > 0$ be a memory horizon, $\tau_R > 0$ a kernel time scale, and $\varepsilon_R > 0$ a regularization parameter. Unless an initial history is specified on $[-T_R, 0]$, the persistence score is defined only for

$$t \geq T_R. \tag{73}$$

For $s \in [0, T_R]$, define the regularized structural correlation

$$r_{\text{str}}^{\text{glob}}(t,s) = \frac{N^{-1}\widetilde{\mathbf{c}}(t)^\top \widetilde{\mathbf{c}}(t-s)}{\sqrt{[V_{\text{glob}}(t) + \varepsilon_R^2][V_{\text{glob}}(t-s) + \varepsilon_R^2]}}. \tag{74}$$

By Cauchy–Schwarz,

$$-1 \leq r_{\text{str}}^{\text{glob}}(t,s) \leq 1. \tag{75}$$

The normalized truncated exponential kernel is

$$K_{T_R,\tau_R}(s) = \frac{e^{-s/\tau_R}}{\tau_R[1 - e^{-T_R/\tau_R}]}\mathbf{1}_{[0,T_R]}(s), \qquad \int_0^{T_R} K_{T_R,\tau_R}(s)\,ds = 1. \tag{76}$$

The global persistence score is

$$R_{\text{str}}^{\text{glob}}(t) = \frac{1}{2}\left[1 + \int_0^{T_R} K_{T_R,\tau_R}(s) r_{\text{str}}^{\text{glob}}(t,s)\,ds\right]. \tag{77}$$

Therefore $0 \leq R_{\text{str}}^{\text{glob}}(t) \leq 1$. A score of $1/2$ corresponds to average zero structural correlation. A threshold intended to require positive structural alignment must consequently satisfy

$$\frac{1}{2} < R_\star^{\text{glob}} < 1. \tag{78}$$

When the structural component is constant over the full memory window,

$$R_{\text{str}}^{\text{glob}}(t) = \frac{1}{2}\left[1 + \frac{V_{\text{glob}}(t)}{V_{\text{glob}}(t) + \varepsilon_R^2}\right]. \tag{79}$$

For $R_\star^{\text{glob}} > 1/2$, the inequality $R_{\text{str}}^{\text{glob}}(t) \geq R_\star^{\text{glob}}$ is equivalent to

$$V_{\text{glob}}(t) \geq \frac{2R_\star^{\text{glob}} - 1}{2(1 - R_\star^{\text{glob}})}\varepsilon_R^2. \tag{80}$$

7 Structural Repertoire and Local Coalitions

7.1 Candidate family and structural repertoire

Let

$$\mathcal{C} = \{S_1, \ldots, S_M\} \tag{81}$$

be a finite family of candidate coalitions. Each $S \subseteq V$ must satisfy

$$n_S = |S| \geq n_\star \geq 2, \qquad \mathcal{G}[S] \text{ is connected.} \tag{82}$$

Let $\text{Stab}_{\text{set}}(S) \in [0,1]$ be a prespecified stability score. The structural repertoire is

$$\Sigma_\Phi^{\text{str}} = \{S \in \mathcal{C} : \text{Stab}_{\text{set}}(S) \geq \pi_{\text{set}}\}, \qquad 0 < \pi_{\text{set}} \leq 1. \tag{83}$$

The candidate-generation rule, stability procedure, and threshold π_{set} must be calibrated, preregistered, or selected on training data only. One admissible implementation is

$$\text{Stab}_{\text{set}}(S) = \frac{1}{B_{\text{rep}}}\sum_{b=1}^{B_{\text{rep}}} \mathbf{1}\left[S \in \mathcal{C}_{\text{ret}}^{(b)}\right]. \tag{84}$$

7.2 Local neural state and induced graph geometry

For $S \in \Sigma_\Phi^{\mathrm{str}}$, let $\mathbf{c}_S(t)$ denote restriction of $\mathbf{c}(t)$ to the nodes in S. Thus

$$\delta\boldsymbol{\Phi}_S(t) = \boldsymbol{\Phi}_S(t) - \boldsymbol{\Phi}_{0,S}, \qquad \mathbf{c}_S(t) = \frac{\delta\boldsymbol{\Phi}_S(t)}{\phi_\star}. \tag{85}$$

The local mean and centred state are

$$\bar{c}_S(t) = \frac{1}{n_S}\mathbf{1}_{n_S}^\top \mathbf{c}_S(t), \qquad \widetilde{\mathbf{c}}_S(t) = \mathbf{c}_S(t) - \bar{c}_S(t)\mathbf{1}_{n_S}. \tag{86}$$

The local dispersion and homogeneity score are

$$V_S(t) = \frac{1}{n_S}\|\widetilde{\mathbf{c}}_S(t)\|_2^2, \qquad A_{\mathrm{hom},S}(t) = \exp\left[-\frac{V_S(t)}{V_{\star,S}}\right], \qquad V_{\star,S} > 0. \tag{87}$$

Let L_S^{ind} be the Laplacian of the induced graph $\mathcal{G}[S]$. Its normalized form is

$$\widehat{L}_S = \frac{L_S^{\mathrm{ind}}}{\lambda_{\max}(L_S^{\mathrm{ind}})}. \tag{88}$$

The total internal graph weight is retained for functional coherence:

$$W_S = \frac{1}{2}\sum_{i,j\in S} W_{ij} = \sum_{i<j\in S} W_{ij} > 0. \tag{89}$$

7.3 Local spatial and temporal activity

The canonical internal spatial activity is

$$E_{S,\mathrm{int}}^{(S)}(t) = \frac{1}{n_S}\mathbf{c}_S(t)^\top \widehat{L}_S \mathbf{c}_S(t) = \frac{1}{n_S}\widetilde{\mathbf{c}}_S(t)^\top \widehat{L}_S \widetilde{\mathbf{c}}_S(t). \tag{90}$$

Its edgewise representation is

$$E_{S,\mathrm{int}}^{(S)}(t) = \frac{1}{2n_S\lambda_{\max}(L_S^{\mathrm{ind}})}\sum_{i,j\in S} W_{ij}[c_i(t) - c_j(t)]^2, \tag{91}$$

and in DeltaPhi coordinates,

$$E_{S,\mathrm{int}}^{(S)}(t) = \frac{1}{2n_S\lambda_{\max}(L_S^{\mathrm{ind}})\phi_\star^2}\sum_{i,j\in S} W_{ij}[\delta\Phi_i(t) - \delta\Phi_j(t)]^2. \tag{92}$$

Because $\lambda_{\max}(\widehat{L}_S) = 1$,

$$0 \le E_{S,\mathrm{int}}^{(S)}(t) \le V_S(t). \tag{93}$$

For $\tau_S > 0$, local temporal activity is

$$E_T^{(S)}(t) = \frac{\tau_S^2}{n_S}\|\dot{\mathbf{c}}_S(t)\|_2^2, \tag{94}$$

where the derivative can be evaluated from the primary dynamics when a continuous-time latent-state estimate is available,

$$\dot{\mathbf{c}}_S(t) = \left[-\kappa_G \widehat{L}\mathbf{c}(t) - \mu\mathbf{c}(t) + B\mathbf{u}(t)\right]_S. \tag{95}$$

7.4 Functional coherence and local persistence

Let $T_C > 0$ denote the temporal coherence window and $\varepsilon_C > 0$ a regularization parameter. Unless an initial history is supplied on $[-T_C, 0]$, functional coherence is defined only for

$$t \geq T_C. \tag{96}$$

For $i \in S$, define the temporally centred trace

$$x_i^{(t)}(s) = c_i(t-s) - \frac{1}{T_C}\int_0^{T_C} c_i(t-r)\,dr, \qquad s \in [0, T_C]. \tag{97}$$

For distinct $i, j \in S$, define the regularized correlation

$$r_{ij}^{(C)}(t) = \frac{\int_0^{T_C} x_i^{(t)}(s)x_j^{(t)}(s)\,ds}{\sqrt{\left[\int_0^{T_C} x_i^{(t)}(s)^2\,ds + \varepsilon_C^2\right]\left[\int_0^{T_C} x_j^{(t)}(s)^2\,ds + \varepsilon_C^2\right]}}. \tag{98}$$

By Cauchy–Schwarz, $-1 \leq r_{ij}^{(C)}(t) \leq 1$. The weighted functional coherence is

$$C_S(t) = \frac{1}{W_S}\sum_{i<j\in S} W_{ij}\frac{1 + r_{ij}^{(C)}(t)}{2}, \qquad 0 \leq C_S(t) \leq 1. \tag{99}$$

A value $C_S(t) = 1/2$ corresponds to weighted average correlation equal to zero. Therefore, a threshold intended to require positive functional coherence must satisfy

$$\frac{1}{2} < C_\star < 1. \tag{100}$$

The local regularized structural correlation is

$$r_{\mathrm{str}}^{(S)}(t,s) = \frac{n_S^{-1}\widetilde{\mathbf{c}}_S(t)^\top \widetilde{\mathbf{c}}_S(t-s)}{\sqrt{[V_S(t) + \varepsilon_R^2][V_S(t-s) + \varepsilon_R^2]}}, \qquad s \in [0, T_R], \tag{101}$$

and the local persistence score is

$$R_{\mathrm{str}}^{(S)}(t) = \frac{1}{2}\left[1 + \int_0^{T_R} K_{T_R,\tau_R}(s) r_{\mathrm{str}}^{(S)}(t,s)\,ds\right]. \tag{102}$$

A threshold intended to require positive local structural alignment must satisfy

$$\frac{1}{2} < R_\star^{\mathrm{loc}} < 1. \tag{103}$$

For each coalition, define

$$\mathbf{f}_S(t) = \begin{pmatrix} A_{\mathrm{hom},S}(t) \\ E_{S,\mathrm{int}}^{(S)}(t) \\ E_T^{(S)}(t) \\ C_S(t) \\ R_{\mathrm{str}}^{(S)}(t) \end{pmatrix}. \tag{104}$$

8 Margin-Based Activation and Multiscale Operational Condition

8.1 Standardized margins

Thresholds are model parameters to be fixed by calibration, synthetic-data analysis, nested training-set optimization, or preregistration. They are not universal biological constants. For a scalar feature with admissible interval $q_{\min} \le q(t) \le q_{\max}$ and scale $s_q > 0$, define

$$\mathrm{mar}_{[q_{\min},q_{\max}]}(q(t); s_q) = \min\left\{\frac{q(t) - q_{\min}}{s_q}, \frac{q_{\max} - q(t)}{s_q}\right\}. \tag{105}$$

For a lower-bound constraint $q(t) \ge q_\star$, define

$$\mathrm{mar}^{-}_{q_\star}(q(t); s_q) = \frac{q(t) - q_\star}{s_q}. \tag{106}$$

A margin is nonnegative if and only if the corresponding constraint is satisfied.

8.2 Global and local margins

Let the global threshold parameters satisfy

$$\begin{aligned}
&0 \le d_{\min} < d_{\max} < 1,\\
&0 < A^{\mathrm{glob}}_{\min} < A^{\mathrm{glob}}_{\max} \le 1,\\
&0 \le E^{\mathrm{glob}}_{S,\min} < E^{\mathrm{glob}}_{S,\max},\\
&0 \le E^{\mathrm{glob}}_{T,\min} < E^{\mathrm{glob}}_{T,\max},\\
&\frac{1}{2} < R^{\mathrm{glob}}_\star < 1.
\end{aligned} \tag{107}$$

The global margin is

$$\begin{aligned}
m_\Phi(t) = \min\Big\{ &\mathrm{mar}_{[d_{\min},d_{\max}]}(d_\Phi(t); s_d),\\
&\mathrm{mar}_{[A^{\mathrm{glob}}_{\min},A^{\mathrm{glob}}_{\max}]}(A_{\mathrm{hom}}(t); s_A),\\
&\mathrm{mar}_{[E^{\mathrm{glob}}_{S,\min},E^{\mathrm{glob}}_{S,\max}]}(E^{\mathrm{glob}}_S(t); s_{E_S}),\\
&\mathrm{mar}_{[E^{\mathrm{glob}}_{T,\min},E^{\mathrm{glob}}_{T,\max}]}(E^{\mathrm{glob}}_T(t); s_{E_T}),\\
&\mathrm{mar}^{-}_{R^{\mathrm{glob}}_\star}(R^{\mathrm{glob}}_{\mathrm{str}}(t); s_R)\Big\}.
\end{aligned} \tag{108}$$

The global condition is

$$\Gamma_\Phi(t) = \mathbf{1}[m_\Phi(t) \ge 0]. \tag{109}$$

For each $S \in \Sigma^{\mathrm{str}}_\Phi$, let

$$\begin{aligned}
&0 < A^{\mathrm{loc}}_{\min} < A^{\mathrm{loc}}_{\max} \le 1,\\
&0 \le E^{\mathrm{loc}}_{S,\min} < E^{\mathrm{loc}}_{S,\max},\\
&0 \le E^{\mathrm{loc}}_{T,\min} < E^{\mathrm{loc}}_{T,\max},\\
&\frac{1}{2} < C_\star < 1,\\
&\frac{1}{2} < R^{\mathrm{loc}}_\star < 1.
\end{aligned} \tag{110}$$

The local margin is

$$
\begin{aligned}
m_S(t) = \min\Big\{ & \mathrm{mar}_{[A^{\mathrm{loc}}_{\min}, A^{\mathrm{loc}}_{\max}]}(A_{\mathrm{hom},S}(t); s^{\mathrm{loc}}_A), \\
& \mathrm{mar}_{[E^{\mathrm{loc}}_{S,\min}, E^{\mathrm{loc}}_{S,\max}]}(E^{(S)}_{S,\mathrm{int}}(t); s^{\mathrm{loc}}_{E_S}), \\
& \mathrm{mar}_{[E^{\mathrm{loc}}_{T,\min}, E^{\mathrm{loc}}_{T,\max}]}(E^{(S)}_{T}(t); s^{\mathrm{loc}}_{E_T}), \\
& \mathrm{mar}^{-}_{C_\star}(C_S(t); s_C), \\
& \mathrm{mar}^{-}_{R^{\mathrm{loc}}_\star}(R^{(S)}_{\mathrm{str}}(t); s^{\mathrm{loc}}_R)\Big\}.
\end{aligned} \tag{111}
$$

8.3 Active coalitions and the multiscale state

The active coalition family is

$$
\Sigma^{\mathrm{act}}_\Phi(t) = \Big\{ S \in \Sigma^{\mathrm{str}}_\Phi : m_S(t) \geq 0 \Big\}. \tag{112}
$$

Define the dominant local margin by

$$
m_\Sigma(t) = \begin{cases} \max\limits_{S \in \Sigma^{\mathrm{str}}_\Phi} m_S(t), & \Sigma^{\mathrm{str}}_\Phi \neq \varnothing, \\ -\infty, & \Sigma^{\mathrm{str}}_\Phi = \varnothing. \end{cases} \tag{113}
$$

Then

$$
\Gamma_\Sigma(t) = \mathbf{1}[m_\Sigma(t) \geq 0] = \mathbf{1}[\Sigma^{\mathrm{act}}_\Phi(t) \neq \varnothing]. \tag{114}
$$

Without an initialized history, the multiscale condition is evaluated only for

$$
t \geq T_H := \max\{T_C, T_R\}. \tag{115}
$$

The multiscale margin and operational condition are

$$
m_{\mathrm{TUGN}}(t) = \min\{m_\Phi(t), m_\Sigma(t)\}, \qquad \Gamma_{\mathrm{TUGN}}(t) = \mathbf{1}[m_{\mathrm{TUGN}}(t) \geq 0]. \tag{116}
$$

Equivalently,

$$
\Gamma_{\mathrm{TUGN}}(t) = \Gamma_\Phi(t)\Gamma_\Sigma(t). \tag{117}
$$

9 Observation Model, Reduced-State Representation, and Inference

9.1 Observation model and likelihood

Let $\mathbf{y}_k \in \mathbb{R}^p$ denote an observed neural measurement vector at $t_k = k\Delta t$, $k = 0, \ldots, K$. The canonical observation model is

$$
\mathbf{y}_k = \mathbf{b}_0 + \phi_\star M \mathbf{c}_k + \boldsymbol{\varepsilon}_k, \tag{118}
$$

where $\mathbf{b}_0 \in \mathbb{R}^p$, $M \in \mathbb{R}^{p \times N}$, and

$$
\boldsymbol{\varepsilon}_k \overset{\mathrm{iid}}{\sim} \mathcal{N}(\mathbf{0}, R), \qquad R = R^\top \succ 0. \tag{119}
$$

Equivalently,

$$
\mathbf{y}_k = \mathbf{b}_0 + M\delta\boldsymbol{\Phi}_k + \boldsymbol{\varepsilon}_k. \tag{120}
$$

Conditional on $\boldsymbol{\theta} = (\kappa_G, \mu)^\top$, $\mathbf{c}_0$, B, $\mathbf{u}_{0:K-1}$, M, and R, the deterministic state recursion defines $\mathbf{c}_1, \ldots, \mathbf{c}_K$. The conditional Gaussian log-likelihood is

$$\begin{aligned}\ell(\boldsymbol{\theta}, \mathbf{c}_0) = & -\frac{K+1}{2}\left[p \log(2\pi) + \log|R|\right] \\ & -\frac{1}{2}\sum_{k=0}^{K}\left[\mathbf{y}_k - \mathbf{b}_0 - \phi_\star M \mathbf{c}_k\right]^\top R^{-1}\left[\mathbf{y}_k - \mathbf{b}_0 - \phi_\star M \mathbf{c}_k\right].\end{aligned} \tag{121}$$

The factorization over time in (121) relies on the independent temporal error assumption in (119). An autocorrelated-error extension requires a joint covariance model and is not part of the present primary specification.

9.2 Reduced-state representation

Let $U_r \in \mathbb{R}^{N \times r}$, $r < N$, have orthonormal columns, $U_r^\top U_r = I_r$. Define

$$P_r = U_r U_r^\top, \qquad \mathbf{e}_\perp(t) = (I_N - P_r)\mathbf{c}(t), \qquad \mathbf{z}(t) = U_r^\top \mathbf{c}(t). \tag{122}$$

Then

$$\mathbf{c}(t) = U_r \mathbf{z}(t) + \mathbf{e}_\perp(t). \tag{123}$$

Assume that U_r is formed by eigenvectors of $\widehat{L}$. The projected dynamics are

$$\dot{\mathbf{z}}(t) = -A_r \mathbf{z}(t) + B_r \mathbf{u}(t), \qquad A_r = \kappa_G U_r^\top \widehat{L} U_r + \mu I_r, \qquad B_r = U_r^\top B. \tag{124}$$

The residual dynamics are

$$\dot{\mathbf{e}}_\perp(t) = -A\mathbf{e}_\perp(t) + (I_N - P_r)B\mathbf{u}(t). \tag{125}$$

Thus, the reduced representation is exact for all time only if

$$\mathbf{e}_\perp(0) = \mathbf{0}, \qquad (I_N - P_r)B\mathbf{u}(t) = \mathbf{0} \quad \forall t. \tag{126}$$

A sufficient condition for the second equality for every admissible input is

$$\mathrm{Range}(B) \subseteq \mathrm{Range}(U_r). \tag{127}$$

9.3 Discrete-time observability and dynamic identifiability

Define

$$F_r = e^{-A_r \Delta t}, \qquad G_r = A_r^{-1}(I_r - F_r)B_r. \tag{128}$$

Then

$$\mathbf{z}_{k+1} = F_r \mathbf{z}_k + G_r \mathbf{u}_k. \tag{129}$$

The reduced observation matrix is $M_r = \phi_\star M U_r$. The primary sampled-data observability criterion uses the discrete Gramian

$$W_{\mathrm{obs},K}^{(r)} = \sum_{k=0}^{K} (F_r^k)^\top M_r^\top R^{-1} M_r F_r^k. \tag{130}$$

The reduced state is observable from sampled measurements on $\{t_0, \ldots, t_K\}$ if and only if

$$W_{\mathrm{obs},K}^{(r)} \succ 0. \tag{131}$$

The continuous-time Gramian is a separate theoretical object for continuous observation and is not the primary empirical criterion here.

For dynamic parameter identifiability, rearrange the primary dynamics as

$$B\mathbf{u}(t) - \dot{\mathbf{c}}(t) = \kappa_G \widehat{L}\mathbf{c}(t) + \mu\mathbf{c}(t). \tag{132}$$

Define

$$X(t) = \begin{bmatrix} \widehat{L}\mathbf{c}(t) & \mathbf{c}(t) \end{bmatrix} \in \mathbb{R}^{N\times 2}, \qquad \mathcal{G}_{\mathrm{id}} = \int_0^T X(t)^\top X(t)\,dt. \tag{133}$$

Proposition 9.1 (State-level dynamic identifiability)**.** *Suppose* $\mathbf{c}(t)$*,* $\dot{\mathbf{c}}(t)$*,* B*, and* $\mathbf{u}(t)$ *are known on* $[0,T]$*. Then* $(\kappa_G,\mu)^\top$ *is uniquely identifiable from* (132) *if and only if*

$$\mathcal{G}_{\mathrm{id}} \succ 0. \tag{134}$$

If $\mathbf{c}(t) = \sum_{m=1}^N a_m(t)\mathbf{v}_m$ and $w_m = \int_0^T a_m(t)^2\,dt$, then

$$\det(\mathcal{G}_{\mathrm{id}}) = \frac{1}{2}\sum_{m=1}^N\sum_{n=1}^N w_m w_n(\lambda_m - \lambda_n)^2. \tag{135}$$

Thus identifiability requires at least two excited graph modes with distinct eigenvalues. A purely uniform trajectory cannot identify κ_G, and a trajectory confined to a single eigenspace identifies only $\mu + \kappa_G\lambda_m$.

10 Mathematical Propositions and Conditional Feasibility

10.1 Continuity of descriptors

Proposition 10.1 (Continuity of TUC-N descriptors)**.** *Suppose* $\mathbf{c} \in C^1([-T_H,T];\mathbb{R}^N)$*, where* $T_H = \max\{T_C, T_R\}$*. Then all global and local descriptors used in* $m_\Phi(t)$*,* $m_S(t)$*, and* $m_{\mathrm{TUC\text{-}N}}(t)$ *are continuous on* $[0,T]$*.*

The result follows from continuity of matrix multiplication, norms, exponentials, finite sums, and bounded-window integrals, together with $\varepsilon_R > 0$ and $\varepsilon_C > 0$. The thresholded indicators Γ_Φ, Γ_Σ, and $\Gamma_{\mathrm{TUC\text{-}N}}$ need not be continuous because they can change when a margin crosses zero.

10.2 Conditional feasibility of the multiscale operational regime

Let $S \in \Sigma_\Phi^{\mathrm{str}}$, and let $\mathbf{v} \in \mathbb{R}^N$ satisfy

$$\widehat{L}\mathbf{v} = \lambda_v\mathbf{v}, \qquad \lambda_v > 0, \qquad \mathbf{1}_N^\top\mathbf{v} = 0, \tag{136}$$

with $\mathbf{v}_S$ nonconstant. Consider the trajectory

$$\mathbf{c}(t) = a\mathbf{v} + q(t)\mathbf{1}_N, \qquad a > 0, \qquad q(t) = q_0 + \nu t, \qquad \nu \neq 0. \tag{137}$$

Then $\widetilde{\mathbf{c}}(t) = a\mathbf{v}$, and

$$V_{\mathrm{glob}} = \frac{a^2}{N}\|\mathbf{v}\|_2^2, \qquad E_S^{\mathrm{glob}} = \frac{a^2\lambda_v}{N}\|\mathbf{v}\|_2^2, \qquad E_T^{\mathrm{glob}} = \tau_G^2\nu^2. \tag{138}$$

Let $P_{\perp,S} = I_{n_S} - n_S^{-1}\mathbf{1}_{n_S}\mathbf{1}_{n_S}^\top$. Then

$$V_S = \frac{a^2}{n_S}\|P_{\perp,S}\mathbf{v}_S\|_2^2, \qquad E_{S,\mathrm{int}}^{(S)} = \frac{a^2}{n_S}\mathbf{v}_S^\top\widehat{L}_S\mathbf{v}_S, \qquad E_T^{(S)} = \tau_S^2\nu^2. \tag{139}$$

For the trajectory (137), the temporally centred traces are identical across nodes,

$$x_i^{(t)}(s) = -\nu \left(s - \frac{T_C}{2} \right), \qquad i \in S. \tag{140}$$

Define

$$H_C = \int_0^{T_C} \nu^2 \left(s - \frac{T_C}{2} \right)^2 ds = \frac{\nu^2 T_C^3}{12}. \tag{141}$$

Then

$$r_{ij}^{(C)}(t) = \frac{H_C}{H_C + \varepsilon_C^2}, \qquad i \neq j, \tag{142}$$

and, crucially,

$$C_S(t) = \frac{1}{2} \left[1 + \frac{H_C}{H_C + \varepsilon_C^2} \right]. \tag{143}$$

For $1/2 < C_\star < 1$, the condition $C_S(t) \geq C_\star$ is equivalent to

$$H_C \geq \frac{2C_\star - 1}{2(1 - C_\star)} \varepsilon_C^2, \tag{144}$$

or, equivalently,

$$|\nu| \geq \varepsilon_C \sqrt{\frac{6(2C_\star - 1)}{(1 - C_\star) T_C^3}}. \tag{145}$$

The persistence scores are

$$R_{\text{str}}^{\text{glob}}(t) = \frac{1}{2} \left[1 + \frac{V_{\text{glob}}}{V_{\text{glob}} + \varepsilon_R^2} \right], \qquad R_{\text{str}}^{(S)}(t) = \frac{1}{2} \left[1 + \frac{V_S}{V_S + \varepsilon_R^2} \right]. \tag{146}$$

The input required to generate (137) is

$$\begin{aligned} B\mathbf{u}(t) &= \dot{\mathbf{c}}(t) + A\mathbf{c}(t) \\ &= [\nu + \mu q(t)] \mathbf{1}_N + a[\mu + \kappa_G \lambda_v] \mathbf{v}. \end{aligned} \tag{147}$$

Proposition 10.2 (Conditional feasibility)**.** *Suppose*

$$\text{span}\{\mathbf{1}_N, \mathbf{v}\} \subseteq \text{Range}(B), \tag{148}$$

and that there exist $a > 0$, $q_0 \in \mathbb{R}$, $\nu \neq 0$, and a finite nonempty interval $I \subseteq [T_H, \infty)$ such that all global and local margins are strictly positive for every $t \in I$. Then

$$\Gamma_{\text{TUC-N}}(t) = 1, \qquad \forall t \in I. \tag{149}$$

The proposition establishes conditional feasibility, not a universal non-vacuity theorem for arbitrary threshold choices. In this construction, positive functional coherence is generated by the shared time-varying input component $(q_0 + \nu t)\mathbf{1}_N$. The construction therefore demonstrates compatibility of the formal gates, rather than establishing that local coherence is intrinsically generated by the spatial graph mode.

11 Computational Verification Protocol

The computational program has two distinct levels. First, unit tests verify the deterministic formulas directly from known synthetic trajectories $\mathbf{c}^{\text{true}}(t)$. Second, end-to-end tests evaluate recovery from noisy observations generated by

$$\mathbf{y}_k = \mathbf{b}_0 + \phi_\star M \mathbf{c}_k^{\text{true}} + \boldsymbol{\varepsilon}_k. \tag{150}$$

The inferential procedure estimates $\widehat{\mathbf{c}}_k$, $\widehat{\kappa}_G$, and $\widehat{\mu}$ before computing $\widehat{\Gamma}_{\text{TUC-N}}(t_k)$. The two levels must be reported separately: unit tests verify formula implementation, whereas end-to-end tests assess recoverability from noisy observations.

The minimum synthetic suite includes six conditions:

SI. Structured persistent driven state. A trajectory satisfying the conditional-feasibility construction.

SII. Uniform static state. $\mathbf{c}(t) = q_0 \mathbf{1}_N$, yielding zero global and local internal spatial activity.

SIII. Static structured state. $\mathbf{c}(t) = a\mathbf{v}$, with $\widehat{L}\mathbf{v} = \lambda_v \mathbf{v}$ and $\lambda_v > 0$, yielding positive spatial activity but zero temporal activity.

SIV. Temporally active uniform state. $\mathbf{c}(t) = q(t)\mathbf{1}_N$, with $\dot{q}(t) \neq 0$, yielding temporal activity but no graph-relative spatial activity.

SV. Low-persistence structured state. A trajectory with adequate instantaneous activity but $R^{\text{glob}}_{\text{str}}(t) < R^{\text{glob}}_\star$.

SVI. Weakly coherent local activity. A state with adequate global descriptors and local spatial activity but $C_S(t) < C_\star$.

Let $\Gamma^{\text{true}}_{\text{TUC-N}}(t_k)$ be the known synthetic label and $\widehat{\Gamma}_{\text{TUC-N}}(t_k)$ the estimated label. Accuracy is

$$\text{Acc}_{\text{TUC-N}} = \frac{1}{K+1} \sum_{k=0}^{K} \mathbf{1}\left[\widehat{\Gamma}_{\text{TUC-N}}(t_k) = \Gamma^{\text{true}}_{\text{TUC-N}}(t_k)\right]. \tag{151}$$

Define

$$N_- = \sum_{k=0}^{K} \mathbf{1}[\Gamma^{\text{true}}_{\text{TUC-N}}(t_k) = 0], \qquad N_+ = \sum_{k=0}^{K} \mathbf{1}[\Gamma^{\text{true}}_{\text{TUC-N}}(t_k) = 1]. \tag{152}$$

When $N_- > 0$,

$$\text{FPR}_{\text{TUC-N}} = \frac{\sum_{k=0}^{K} \mathbf{1}[\widehat{\Gamma}_{\text{TUC-N}}(t_k) = 1,\ \Gamma^{\text{true}}_{\text{TUC-N}}(t_k) = 0]}{N_-}. \tag{153}$$

When $N_+ > 0$,

$$\text{FNR}_{\text{TUC-N}} = \frac{\sum_{k=0}^{K} \mathbf{1}[\widehat{\Gamma}_{\text{TUC-N}}(t_k) = 0,\ \Gamma^{\text{true}}_{\text{TUC-N}}(t_k) = 1]}{N_+}. \tag{154}$$

When the relevant denominator is zero, the corresponding rate must be reported as NA, rather than zero.

Across R_{sim} replicates, parameter and state recovery are measured by

$$\text{RMSE}_{\boldsymbol{\theta}} = \left[\frac{1}{R_{\text{sim}}} \sum_{r=1}^{R_{\text{sim}}} \|\widehat{\boldsymbol{\theta}}^{(r)} - \boldsymbol{\theta}^{\text{true},(r)}\|_2^2\right]^{1/2}, \tag{155}$$

and

$$\text{RMSE}_{\mathbf{c}} = \left[\frac{1}{K+1} \sum_{k=0}^{K} \|\widehat{\mathbf{c}}_k - \mathbf{c}_k^{\text{true}}\|_2^2\right]^{1/2}. \tag{156}$$

No Brier score is included in the primary protocol because the canonical model currently outputs a binary operational indicator, not a fully specified posterior probability $\mathbb{P}[\Gamma_{\text{TUC-N}}(t) = 1 \mid \mathbf{y}_{0:K}]$.

12 Empirical Predictions and Model Evaluation

The empirical contribution of TUC-N must be evaluated through predictions that exceed those of prespecified baseline models. The primary endpoints are an access-related binary outcome $Y_{i,k}^{\text{acc}} \in \{0, 1\}$ and a categorical reported-content outcome $Y_{i,k}^{\text{cont}} \in \{1, \dots, Q\}$. Their operational definitions must be declared before model fitting.

Let $\mathbf{x}_{i,k}^{\text{std}}$ denote prespecified standard covariates. The baseline access model is

$$\operatorname{logit} \Pr[Y_{i,k}^{\text{acc}} = 1] = \alpha_i + \mathbf{x}_{i,k}^{\text{std}\top} \boldsymbol{\beta}_{\text{std}}. \tag{157}$$

The global TUC-N access model is

$$\operatorname{logit} \Pr[Y_{i,k}^{\text{acc}} = 1] = \alpha_i + \mathbf{x}_{i,k}^{\text{std}\top} \boldsymbol{\beta}_{\text{std}} + \mathbf{f}_\Phi(t_k)^\top \boldsymbol{\beta}_\Phi, \tag{158}$$

where

$$\mathbf{f}_\Phi(t) = \begin{pmatrix} d_\Phi(t) \\ A_{\text{hom}}(t) \\ E_S^{\text{glob}}(t) \\ E_T^{\text{glob}}(t) \\ R_{\text{str}}^{\text{glob}}(t) \end{pmatrix}. \tag{159}$$

For reported content, the coalition summary uses active coalitions only:

$$\mathbf{z}_\Sigma(t) = \begin{cases} \dfrac{1}{|\Sigma_\Phi^{\text{act}}(t)|} \displaystyle\sum_{S \in \Sigma_\Phi^{\text{act}}(t)} \mathbf{f}_S(t), & \Sigma_\Phi^{\text{act}}(t) \neq \varnothing, \\ \mathbf{0}, & \Sigma_\Phi^{\text{act}}(t) = \varnothing. \end{cases} \tag{160}$$

For $q = 1, \dots, Q$, the multiclass model is

$$\Pr[Y_{i,k}^{\text{cont}} = q] = \frac{\exp(\eta_{q,i,k})}{\sum_{q'=1}^{Q} \exp(\eta_{q',i,k})}, \tag{161}$$

with

$$\eta_{q,i,k} = \alpha_{q,i} + \mathbf{x}_{i,k}^{\text{std}\top} \boldsymbol{\beta}_{q,\text{std}} + \mathbf{f}_\Phi(t_k)^\top \boldsymbol{\beta}_{q,\Phi} + \mathbf{z}_\Sigma(t_k)^\top \boldsymbol{\beta}_{q,\Sigma}. \tag{162}$$

For a generic endpoint Y, define held-out log predictive density by

$$\text{LPD}_Y(\mathcal{M}) = \sum_{(i,k) \in \mathcal{D}_{\text{test}}} \log p_{\mathcal{M}} \left(Y_{i,k} \mid \mathcal{D}_{\text{train}}\right). \tag{163}$$

The access comparison is

$$\Delta \, \text{LPD}_{\text{glob}-\text{std}}^{\text{acc}} = \text{LPD}_{Y^{\text{acc}}}(\mathcal{M}_{\text{glob}}) - \text{LPD}_{Y^{\text{acc}}}(\mathcal{M}_{\text{std}}), \tag{164}$$

and the content comparison is

$$\Delta \, \text{LPD}_{\text{multi}-\text{glob}}^{\text{cont}} = \text{LPD}_{Y^{\text{cont}}}(\mathcal{M}_{\text{multi}}) - \text{LPD}_{Y^{\text{cont}}}(\mathcal{M}_{\text{glob}}). \tag{165}$$

The corresponding empirical predictions are

$$H_1: \quad \Delta \, \text{LPD}_{\text{glob}-\text{std}}^{\text{acc}} > 0, \qquad H_2: \quad \Delta \, \text{LPD}_{\text{multi}-\text{glob}}^{\text{cont}} > 0. \tag{166}$$

Non-confirmation occurs if the relevant held-out predictive improvement is nonpositive, if the reduced-state system is not observable, if the dynamical parameters are not identifiable, or if synthetic verification reveals unacceptable recovery error under the declared protocol.

13 Discussion

This article defines the canonical mathematical architecture of TUC-N as a graph-dynamical multiscale model of candidate neural states. The framework is organized around a neural reference state $\mathbf{\Phi}_0$, a reference-relative deviation field $\delta\mathbf{\Phi}(t)$, and a normalized latent state $\mathbf{c}(t)$. The primary dynamics combine graph-structured relaxation, nodewise relaxation, and explicit input. The global component quantifies distance from the reference, structural homogeneity, graph-relative spatial activity, temporal activity, and persistence. The local component defines a stable repertoire of graph-connected candidate coalitions and distinguishes that repertoire from its momentary active subset.

The model is designed so that its components can fail independently. A state may be distant from the reference but spatially uniform; it may be spatially differentiated but temporally inactive; or it may be temporally active but lack persistent structural organization. A coalition may show local graph energy without functional coherence, or functional coherence without adequate structural persistence. The margin-based construction preserves these distinctions rather than collapsing them into an unexamined scalar.

The canonical linear dynamics are intentionally restrictive. Their function is to provide a transparent starting point in which stability, bounded-input response, spectral behaviour, observability, and identifiability can be analysed explicitly. Nonlinear and stochastic extensions should be considered only when they provide a measurable advantage in preregistered model comparison, parameter recovery, or empirical prediction.

Several modelling decisions require explicit sensitivity analysis: graph construction, reference estimation, the observation operator, the coalition construction algorithm, threshold selection, and calibration scales. Coalition stability and functional coherence also require controls for measurement artefacts, common input, spatial leakage, and preprocessing-induced dependencies. Dedicated mathematical work remains necessary to establish broader stability results, necessary and sufficient conditions for identifiability, and robustness of coalition selection under perturbations. Simulation-based calibration is needed to quantify recovery under controlled noise, graph error, and observation-model mismatch. Modality-specific observability analysis is needed to determine which state components can be reconstructed from EEG, MEG, fMRI, and intracranial recordings.

14 Conclusion

TUC-N provides a graph-dynamical state-space framework for representing reference-relative neural activity across global and local scales. The model defines a normalized latent field evolving under stable graph-based dynamics, a set of global descriptors of neural organization, a structural repertoire of candidate local coalitions, and a margin-based multiscale operational condition.

The framework generates explicit requirements for observation, state reconstruction, discrete-time observability, dynamic identifiability, simulation-based verification, and empirical prediction. Its central empirical claims are evaluable through held-out comparison with prespecified baseline models. The present work establishes the formal specification required for subsequent theoretical, computational, and empirical developments.

A Pre-analysis declaration of model components

Before confirmatory empirical analysis, the following quantities must be fixed, calibrated independently, or estimated on training data only: $\mathbf{\Phi}_0$, $\phi_\star$, W, B, the admissible input representation $\mathbf{u}(t)$, the observation operator M, the observation covariance R, $\rho_\star$, $V_\star$, $V_{\star,S}$, ε_R, ε_C, T_R, T_C, τ_R, all threshold intervals, all margin scales, π_{set}, and the coalition-construction procedure.

Thresholds affected by the canonical energy normalization must be calibrated in the normalized scale used in (66) and (90).

B Logical status of the claims

The mathematical statements in this article have distinct logical roles. Equations that introduce quantities are *definitions*. Stability, continuity, observability, and state-level identifiability statements are *mathematical propositions* conditional on their stated assumptions. The construction in Section 10 is a *conditional feasibility result*; it does not establish universal non-vacuity for arbitrary threshold choices. The access and reported-content statements in Section 12 are *empirical predictions* whose support depends on held-out data and prespecified comparisons.

References

[1] Melloni, L., Mudrik, L., Pitts, M., et al. (2023). An adversarial collaboration protocol for testing contrasting predictions of global neuronal workspace and integrated information theory. *PLOS ONE*, *18*(2), e0268577.

[2] Ferrante, O., Melloni, L., Pitts, M., et al. (2025). Adversarial testing of global neuronal workspace and integrated information theories of consciousness. *Nature*, *642*, 133–142.

[3] Paninski, L., Ahmadian, Y., Ferreira, D. G., et al. (2010). A new look at state-space models for neural data. *Journal of Computational Neuroscience*, *29*, 107–126.

[4] Linderman, S. W., Johnson, M. J., Miller, A. C., Adams, R. P., Blei, D. M., & Paninski, L. (2017). Bayesian learning and inference in recurrent switching linear dynamical systems. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 914–922.

[5] Faskowitz, J., Esfahlani, F. Z., Jo, Y., Sporns, O., & Betzel, R. F. (2020). Edge-centric functional network representations of human cerebral cortex reveal overlapping system-level architecture. *Nature Neuroscience*, *23*, 1644–1654.