

Brain state sets the spatial scale at which cortical geometry constrains human cortical dynamics

Alejandro Salinas-Medina

alejandro.salinas@mail.mcgill.ca

McGill University <https://orcid.org/0009-0009-9548-9315>

Anisleidy Gonzalez-Mitjans

McGill University <https://orcid.org/0000-0001-7205-9704>

Paule-Joanne Toussaint

McGill University <https://orcid.org/0000-0002-7446-150X>

Xue Liu

McGill University

Alan Evans

McGill University <https://orcid.org/0000-0003-3841-6098>

Article

Keywords:

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11120584/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: There is **NO** Competing Interest.

Brain state sets the spatial scale at which cortical geometry constrains human cortical dynamics

Alejandro Salinas-Medina^{*1,2}, Anisleidy González-Mitjans², Paule-Joanne Toussaint², Xue
Liu^{1,3}, and Alan Evans²

¹School of Computer Science, McGill University, Montreal, Quebec

²McGill Centre for Integrative Neuroscience (MCIN), Montreal, Quebec

³Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, United Arab
Emirates

September 23, 2026

Abstract

How anatomy constrains large-scale cortical dynamics is contested. Two lines of evidence about what shapes large-scale cortical activity appear to conflict. One holds that geometric eigenmodes of the cortical surface explain brain function more parsimoniously than connectome-derived modes; another holds that long-range connections shape dynamics over and above geometry. We show that neither is categorically right, and that the disagreement dissolves once spatial scale is resolved: connectivity accounts for the broadest cortical patterns and histological geometry for the finer ones. Using cortico-hippocampal neural field models built on ultrahigh-resolution histological geometry and tested against human intracranial recordings and task fMRI, we identify a second constraint. Anatomical pathways are commonly treated as fixed contributions that operate similarly across brain states. Instead, the same hippocampal-entorhinal relay degrades cortical theta organisation during sleep yet supplies it during spatial navigation. Brain state changes the effective spatial scale at which cortical geometry is read out, so anatomy constrains cortical function only jointly with state.

1 INTRODUCTION

The anatomy of the brain constrains its function at multiple scales. At the microscale, laminar and cytoarchitectural structure shape local excitability [1, 2, 3, 4]; at the mesoscale, canonical microcircuits, cytoarchitectonic areal boundaries, and laminar projection motifs organise local circuits into regions [5, 6, 7]; at the macroscale, cortical folding, surface geometry, and white-matter connectivity constrain how activity spreads [8, 9, 10, 11, 12]. These anatomical constraints raise a modelling question: how should the brain structure be represented to explain the spatial organisation of neural activity? Much of modern neuroscience systems has addressed this question through the network side of the problem: regions are nodes, tractography or functional coupling defines edges, and activity is modeled as propagation over a graph [13, 14, 15, 16]. Network models capture how distant regions communicate, but they can obscure the fact that the cortex and hippocampus are continuous folded structures [1, 17, 18]. Neural activity is also constrained by the geometry of these

*Corresponding author: alejandro.salinas@mail.mcgill.ca

surfaces, which determines the natural spatial patterns activity can express independently of how the brain is parcellated [10, 9]. BigBrain provides an ultrahigh-resolution histological reference for this anatomical substrate, and BigBrainWarp enables its integration with multimodal neuroimaging spaces [1, 17].

Neural field theory (NFT) provides a natural framework for continuous-field modeling of brain activity on anatomical surfaces. Rather than modeling activity only as interactions among discrete populations or regions, neural fields describe activity as a spatially continuous pattern over tissue [19, 20, 21], and on a surface these patterns can be represented in the geometric eigenmodes of the Laplace–Beltrami operator [21, 10, 9].

Two groups have reached opposite conclusions about what these representations imply. Pang and colleagues reported that geometric eigenmodes reconstruct human brain activity more parsimoniously than connectome-derived modes, and argued that geometry is the dominant constraint [9]. Vohryzek and colleagues reported that rare long-range connections shape brain dynamics over and above cortical geometry, and argued that connectivity contributes independently [11]. Royer and colleagues subsequently reported graded contributions of local geometry and network topography [12], which is compatible with either position depending on where the grading is evaluated. Both of the first two claims rest on scale-aggregated summary statistics, so the disagreement is underdetermined by the evidence offered for it: a scalar averaged over spatial scales can be made to favour either basis by changing the mixture of scales in the evaluation set. We show below that neither categorical position survives a scale-resolved comparison.

A second, largely unexamined question runs alongside the first. Structural network models generally treat anatomical connectivity as a fixed substrate that constrains the dynamics expressed by the system [22, 13, 14, 15, 23, 16]. Whether the functional contribution of a particular anatomical route is itself invariant across brain states is less clear. Human electrophysiology suggests that this question is especially relevant for the hippocampal–cortical route. Medial-temporal theta occurs in bouts during real-world ambulatory movement and navigation [24], hippocampal theta dynamics structure real-world and imagined navigation and memory-guided behaviour [25, 26, 27, 28], and hippocampal–entorhinal theta/alpha synchronization is engaged during memory maintenance [29], with entorhinal cortex acting as a principal hippocampal–neocortical interface. Cortico-hippocampal neural-field models generate intrinsic cortical alpha and hippocampal theta on anatomical substrates [30]; this spectral asymmetry matters here because our primary electrophysiological validation targets cortical theta power. Theta is the appropriate band in which to test this, because human hippocampal and neocortical theta are separately generated and dynamically regulated by brain state [31], making theta-band cortical power the signal in which a state-dependent hippocampal contribution should be visible. Together, these findings motivate testing whether hippocampal–cortical coupling is a fixed anatomical constraint or a state-dependent drive.

Contributions. We therefore address three questions in turn: whether a fixed anatomical route makes a fixed functional contribution, whether high-resolution histological geometry provides an empirically supported substrate rather than merely serving as an implementation detail, and whether geometry or connectivity is the dominant constraint. We address these questions across four analyses, organized so that the most general claim leads. First, in a pre-specified state-by-architecture test, we show that **the same anatomical coupling axis reverses sign across states**: it reduces alignment with empirical sleep-theta power, yet on the finest cortical mesh examined improves alignment with empirical navigation-theta power across two independent navigation cohorts [32, 33]. Second, we show that the state-dependent coupling result is grounded in an empirically supported geometric substrate: **the uncoupled BigBrain geometric surface aligns with intracranial**

sleep-theta in slow-wave sleep across two cohorts differing in age: the predominantly pediatric ds004551 [34, 35] and adult N3 sleep in Penn SIEVE [36], whereas a lower-resolution reference surface is sign-reversed or near-null, with controls showing that the effect is not explained by vertex density, pseudo-replication, or generic spatial autocorrelation, and that the positive alignment is concentrated in cortical regions validly sampled by BigBrain but not adequately reached by the lower-resolution surface. Thus, the data support neither the strong claim that histological resolution improves local prediction nor the null claim that population-average surfaces are sufficient. Third, we show that **this alignment is not a property of the surface alone**: a standard Wilson–Cowan neural mass model (NMM) run on the identical geometry is near-null in both states. Finally, using Human Connectome Project task-fMRI data [37, 38, 39], we show that **geometry and connectivity dominate different spatial scales**, so that neither [9] nor [11] is categorically right and both can be defended with correct evidence drawn from different scale ranges. The useful scientific framework is therefore not geometry versus connectivity, but geometry as the field coordinate system, with connectivity and coupling acting as scale- and state-dependent forcing constraints.

2 RESULTS

2.1 A hippocampal–cortical relay reverses sign across brain states

We tested whether the functional contribution of a single hippocampal–cortical relay changes across brain states using a pre-specified state $\times$ architecture design. In the high-resolution BigBrain model, the relay reduced theta alignment during sleep but improved it during navigation, and this effect was associated with a shift in the effective spatial scale at which the histological substrate was read out. Hippocampal drive enters the cortical field through a coupling operator whose strength is set by a single scalar cc , with reciprocal cortex-to-hippocampus drive fixed at $ch = 0.2cc$ ($cc = 0$ is the uncoupled model; Methods, Sec. 4.4). Among the anatomically informed operators compared in sleep (Methods, Sec. 4.6), an entorhinal relay with row-wise drive normalization across candidates (EC-refrow) gave the strongest subject-level theta-map alignment among the coupled models. We therefore carried EC-refrow forward as the single coupling axis to test. Before computing any navigation alignment, we froze that axis, the per-subject slope of AP-partial ρ against cc , the confirmatory fixed- cc points, and the subject- and coverage-feasibility gates.

The slope analysis in Fig. 1a,c was estimated across the full pre-specified EC-refrow coupling-strength sweep:

$$cc \in \{0, 0.002, 0.005, 0.008, 0.010, 0.015, 0.020, 0.030, 0.050, 0.080\}.$$

We applied these unchanged choices to two independent navigation-memory iEEG datasets (ds005522, $n = 49$; ds005523 retrieval, $n = 19$), with predominantly pediatric sleep (ds004551, $n = 114$) as the subject-level anchor, and evaluated adult Penn SIEVE N3 as an independent contact-level sleep replication (datasets and validation pipelines in Methods, Secs. 4.7–4.8). The relay’s effect on theta alignment inverted between states. Across the full EC-refrow sweep shown in Fig. 1a,c, the per-subject slope of AP-partial ρ against coupling strength was negative in the ds004551 sleep anchor (-3.76), while the Penn SIEVE adult N3 pooled-contact full sweep showed the same negative direction (linear slope -2.30). In contrast, per-subject slopes were positive in both navigation cohorts (ds005522: $+1.16$; ds005523: $+2.85$). The navigation-minus-sleep contrast was highly significant in both datasets (one-sided Mann–Whitney $p = 1.7 \times 10^{-10}$ and $p = 6.8 \times 10^{-8}$), and each navigation cohort independently showed a slope significantly greater than zero. **The same frozen axis that pulls the model away from empirical theta in sleep pulls it toward empirical**

theta in navigation. This sign reversal was not a sample-size artifact of the larger ds004551 sleep anchor. When ds004551 subjects were randomly downsampled to match ds005522 ($n = 49$), the sleep slope remained strongly negative (mean -3.761 , 95% interval $[-4.723, -2.793]$). The same held when ds004551 was downsampled to match ds005523 ($n = 19$; mean -3.760 , 95% interval $[-5.608, -1.868]$). Thus the sleep-side decrease is not explained by having more sleep subjects than navigation subjects.

In the Penn SIEVE adult N3 full-sweep analysis shown in Fig. 1a,b,c, pooled-contact AP-controlled alignment decreased from $+0.090$ at $cc = 0$ to -0.113 at $cc = 0.030$, giving a pooled endpoint change of $\Delta\rho_{AP} = -0.203$. Subject-level Penn SIEVE slopes were computed in the complete-sweep subset ($n = 70$), providing a subject-level sensitivity check alongside the pooled contact-level result. The seeded clinically normal-like endpoint check is reported in Table 1, and the separate Penn SIEVE substrate check is reported in Fig. 2g.

The two states differed in which operating point produced alignment along the same coupling axis. In sleep the uncoupled full-resolution BigBrain surface is already a strong theta substrate (pooled AP-partial $\rho = +0.153$ at $cc = 0$; subject-mean $+0.171$), and coupling degrades it monotonically, crossing zero near $cc = 0.008$ and reaching -0.167 at $cc = 0.030$. Coupling therefore organizes the model’s posterior topography while moving it away from the theta-specific target, dissociating topographic organization from empirical theta validity. Navigation shows the identical picture with the roles exchanged: the uncoupled full-resolution surface is negative (-0.002 and -0.075), and coupling supplies the alignment it lacks, rising to subject-mean $+0.134$ and $+0.209$, comparable to the sleep substrate’s own subject-mean alignment ($+0.171$; pooled $+0.153$), so the state-dependent operating-point contrast holds at matched aggregation (Fig. 1b). Where sleep begins aligned and coupling changes it, navigation begins near zero and coupling supplies alignment. To ask whether this navigation gain was specific to coupling or reflected a preferred spatial scale, we also evaluated uncoupled BigBrain maps across a downsampling series (30k–164k vertices, same anatomy and coordinate frame; Methods, Sec. 4.9). Positive navigation alignment emerged from the uncoupled BigBrain 30k nodes surface ($+0.139$ and $+0.134$), comparable to the coupled full-resolution result. Thus coupling and coarsening both move the navigation prediction toward an effective spatial scale favored by the empirical data, without implying that they are the same mechanism. Across the uncoupled downsamples, navigation alignment was strongest at the coarse 30k surface and was lower at every finer resolution than at 30k (ds005522: $+0.139, +0.065, +0.038, -0.002$; ds005523: $+0.134, -0.009, -0.008, -0.075$), the exact inverse of sleep, which peaks at intermediate-to-full resolution. Every contact was retained identically across all four resolutions (9457/9457 in ds005522 and 6085/6085 in ds005523), so the decline reproduces on identical contacts. The substrate thus supports theta alignment at a state-dependent spatial scale (fine in sleep, coarse in navigation, as if the navigating brain reads the same cortical geometry through a coarser lens) with EC-refrow coupling on the full-resolution surface acting as one route to the coarse scale that navigation favours.

This scale interpretation survives the main design checks. The contrast is not an artifact of electrode coverage: it survives common-support and nearest-neighbour matching on contact count and AP span (Methods, Sec. 4.8), with every navigation slope still exceeding sleep (Fig. 1d). The matched sleep anchor was sensitive to the matching scheme, with its slope ranging from -3.76 to -0.50 (Fig. 1e). Because the ds005523 matching produced the sleep slope closest to zero, and therefore the smallest navigation–sleep separation, we treat this as the conservative matched comparison. The contrast was also not explained by a generic anterior–posterior gradient: an AP-preserving spatial null was exceeded for both the slope and fixed- cc delta statistics in both navigation datasets, with all p -values at the attainable minimum of $1/1001$ under 1,000 permutations. Finally, the effect is specific to the BigBrain substrate rather than a generic property of neural-field maps: a lower-resolution reference map did not supply navigation alignment, but it was coverage-limited rather

172 than conclusive (only $n = 10$ and $n = 4$ subjects survived the pre-specified gate, pooled $\rho = -0.067$
173 and -0.061), so the coupling-specific claim rests on the within-BigBrain 164k uncoupled-versus-
174 coupled contrast.

175 As an exploratory anatomical-transfer check, we repeated the fixed-parameter analysis using ho-
176 mologous right-hemisphere BigBrain and HippUnfold assets (Supplementary Methods, “Exploratory
177 right-hemisphere lateralization sensitivity analysis”). The available contact-level sleep checks pre-
178 served the negative coupling direction, whereas navigation transfer was dataset-dependent. Al-
179 though the two hemispheric eigenspectra were strongly homologous, the frozen $cc = 0.030$ operating
180 point occupied different cortical dominant-frequency regimes. Because right-sided ds004551 cover-
181 age was insufficient for the primary subject-level sleep statistic, this analysis establishes pipeline
182 portability and bounds transfer across anatomical realizations.

183 Together these results identify brain state, rather than anatomy alone, as the variable that sets
184 the sign of a hippocampal–cortical relay: a single frozen axis is detrimental to intrinsic sleep-theta
185 alignment yet, at full resolution, provides a route to navigation-theta alignment, and what it supplies
186 is a shift toward the coarser effective geometry that the navigating brain appears to favour.




Figure 1: Hippocampal–cortical coupling is state-gated and reverses sign across sleep and navigation. **a**, AP-partial Spearman correlation between model and empirical theta power as a function of coupling strength cc , for the ds004551 sleep anchor and both navigation-memory cohorts. **b**, Operating-point comparison across states. Navigation alignment emerges either through EC-refrow coupling at full resolution or through a coarser uncoupled substrate. **c**, Subject-level slope distributions for the ds004551 sleep anchor ($n = 114$), Penn SIEVE N3 complete-sweep subset ($n = 70$), and the two navigation cohorts (ds005522, $n = 49$; ds005523, $n = 19$). **d**, Coverage-matched state contrast under common-support and nearest-neighbour matching. **e**, Matched sleep anchor for the contrast in **d**; the ds005523 comparison is conservative because this anchor is scheme-sensitive. Statistics are reported in Sec. 2.1.

2.2 The state-gating sign reversal replicates across simulator realizations

The state-gating claim reduces to a single directional prediction: strengthening the entorhinal relay should move the model away from empirical theta in sleep and toward it in navigation. A prediction this sharp is worth only as much as its reproducibility, and the sign reversal reported in Section 2.1 was drawn from one stochastic realization of the simulator. We therefore asked whether the reversal is a property of the coupled model or an accident of that particular draw.

Holding the simulator and the analysis pipeline frozen (Methods, Secs. 4.5, 4.7–4.8), we recomputed the uncoupled ($cc = 0$) and coupled ($cc = 0.03$, $ch = 0.2cc$) operating points under three independent random seeds, and measured the coupling response:

$$\Delta\rho_{AP} = \rho_{AP}(cc = 0.03) - \rho_{AP}(cc = 0),$$

predicting a negative response in sleep and a positive response in navigation.

The predicted sign held in every realization. Sleep-theta alignment fell under all three seeds, and navigation-theta alignment rose across both independent cohorts and every analysis contrast, so that all twelve seed-by-dataset comparisons preserved the state-specific direction of the original experiment (Table 1). The reversal is thus a feature of the model, not of a single stochastic draw.

This is the same effect measured before. Under identical empirical contact gates, seed 73924 reproduced the original single-realization endpoints exactly (ds005522, -0.002 to $+0.134$; ds005523 retrieval, -0.075 to $+0.209$), matching the $+0.136$ all-task contrast of Section 2.1. The additional seeds therefore vary only the simulator realization, leaving the empirical sampling and the analysis gate fixed, so the seeded experiment is a stochastic superset of the original contrast rather than a separate analysis.

Coupling also separated two quantities that are easily conflated: simulator-level reproducibility and empirical alignment. Strengthening coupling made the simulated theta topography markedly more reproducible across draws, raising the mean pairwise map correlation from 0.667 to 0.982 and reducing normalized dispersion from 0.0560 to 0.0085. Coupling therefore regularizes the map even as it moves that map away from the sleep target.

The state-specific direction was preserved across seeds, but the magnitude of the coupling response was not. Relative across-seed variability exceeded the pre-specified 25% criterion in both states (sleep 34.3%; navigation 55.5–92.8%), so we do not claim a seed-invariant effect magnitude. The reproducible result is directional: across all stochastic realizations, coupling decreased spatial-map alignment with empirical theta during sleep and increased it during navigation.

Across independent stochastic realizations, the same frozen coupling axis consistently decreased sleep-theta alignment and increased navigation-theta alignment, establishing the direction of the state-gating dissociation as a reproducible property of the coupled system.

Table 1: Seeded validation of the directional state-gating prediction. Values are changes in subject-level AP-partial theta-map alignment from $cc = 0$ to $cc = 0.03$. Coupling decreased ds004551 sleep alignment and increased navigation alignment in every stochastic realization. The Penn SIEVE row is a clinically normal-like adult N3 endpoint sensitivity; it was directionally negative across all three seeds.

State and dataset	Seed 73924	Seed 104729	Seed 221017	Mean $\pm$ SD
Sleep, ds004551	-0.357	-0.284	-0.173	-0.271 ± 0.093
Navigation, ds005522 all-task	$+0.136$	$+0.072$	$+0.007$	$+0.072 \pm 0.065$
Navigation, ds005522 retrieval	$+0.131$	$+0.071$	$+0.004$	$+0.069 \pm 0.064$
Navigation, ds005523 retrieval	$+0.284$	$+0.213$	$+0.076$	$+0.191 \pm 0.106$
Adult sleep, Penn SIEVE N3 clinically normal-like	-0.050	-0.009	-0.000	-0.020 ± 0.026

2.3 High-resolution histological geometry is an empirical sleep-theta substrate

The state-gating result is only as meaningful as the baseline it perturbs. If the uncoupled BigBrain surface were an arbitrary map, its response to coupling would carry no information about brain

state. Before introducing any coupling, we therefore asked whether high-resolution histological geometry tracks empirical sleep theta at all, and we show that it does. We further show that its advantage over a lower-resolution reference surface [30] is not what the resolution difference would suggest. It does not predict theta better where the two surfaces sample the same tissue; it supplies positive alignment in geometrically valid cortex that the reference surface cannot reach (Fig. 2a).

Uncoupled baseline model. The baseline is the corticothalamic neural-field system evaluated on the geometric (Laplace–Beltrami) eigenmodes of the BigBrain left-hemisphere cortical surface (Methods Section 4.4); the parameters used in this NFT model are reported in Table 3. We evaluated the simulated cortical 4–8 Hz theta-power map against intracranial sleep iEEG (ds004551) through the frozen validation pipeline (nearest-surface assignment within 15 mm; Methods, Sec. 4.7). Within that threshold, all 8,985 BigBrain coordinate rows lay inside the 15 mm gate; their median nearest-surface distance was 1.39 mm, and 8,948 had matching empirical bandpower, whereas the reference surface retained 1,379 (15.4%) at a median retained distance of 11.31 mm (Fig. 2a). Over all contact coordinates before gating, the median nearest-surface distance is 1.39 mm for BigBrain and 43.98 mm for the reference surface, so the disparity is in coverage rather than in retained proximity and the two surfaces do not sample the same cortex.

The uncoupled map aligned positively with empirical sleep theta (Fig. 2c). The pooled Spearman correlation was raw $\rho = +0.133$ ($n = 8,948$) and rose to partial $\rho = +0.153$ after controlling for MNI-Y/AP position, indicating that the alignment extends beyond the empirical anterior–posterior theta gradient. The effect is small: $\rho^2 \approx 0.018$, taken as a descriptive rank-correlation heuristic rather than a regression variance-explained, and the extreme p -values (to 6.3×10^{-48}) reflect the large pooled n rather than a large effect, since statistical detectability and predictive magnitude are distinct quantities. Pooling also hides substantial between-subject spread: per-subject correlations averaged raw $+0.140$ and partial $+0.171$ (medians $+0.181$ and $+0.203$; $SD \approx 0.245$), were positive in 88/114 and 87/114 subjects, and a 10,000-resample subject bootstrap placed the mean raw ρ at 95% CI $[+0.096, +0.184]$ and partial at $[+0.126, +0.215]$; thus indicating a consistent group-level alignment.

Penn SIEVE provided an independent adult contact-level substrate check (Fig. 2g). In clinically normal-like N3 contacts, the uncoupled 164k BigBrain map showed positive AP-controlled alignment (AP-partial $\rho = +0.067$, $n = 1,880$; AP-bin $p = 0.036$), whereas the 30k BigBrain control was negative (-0.012) and the lower-resolution reference surface was near-null ($+0.013$, $n = 1,875$; two-sided AP-bin $p = 0.634$). Across all Penn SIEVE N3 contacts, the same ordering strengthened: 164k BigBrain reached AP-partial $\rho = +0.115$ ($n = 5,388$; AP-bin $p = 9.999 \times 10^{-5}$), whereas the reference surface was sign-reversed (-0.032 , $n = 5,361$; two-sided AP-bin $p = 0.035$). Thus Penn SIEVE supports the narrower adult contact-level substrate claim: the uncoupled 164k BigBrain map carries positive AP-controlled N3 theta alignment, while the lower-resolution reference surface is near-null or wrong-sign.

Decomposing the sample by where each surface is geometrically valid locates this alignment precisely, and locates it away from the explanation a resolution difference would invite (Fig. 2c). On the $n = 1,379$ shared-analytic contacts (1,385 were geometrically valid on both surfaces; 1,379 survived within-set standardization), both maps are negative (BigBrain raw $\rho = -0.121$, partial -0.117 ; reference raw $\rho = -0.116$, partial -0.100), so there is no evidence that high-resolution histology improves local theta prediction where both surfaces sample the same tissue. The positive pooled effect is concentrated entirely in the $n = 7,563$ contact-sessions valid only on BigBrain (partial $\rho = +0.181$); with the 6 reference-valid non-analytic rows, the three sets sum to 8,948 ($1,379 + 7,563 + 6$). What the high-resolution geometry supplies is expanded geometrically valid

coverage together with positive alignment across that additional cortex, not superior dynamics at shared locations.

This analysis also constrains the interpretation of the nominally matched comparison in (Fig. 2b). Sampling the reference surface at the identical BigBrain-retained rows leaves BigBrain positive (raw $\rho = +0.133$, partial $+0.153$) and the reference sign-reversed (raw -0.128 , partial -0.104), but most of the 8,948 coordinates lie beyond the reference surface’s 15 mm validity threshold. That comparison extrapolates the reference map into cortex it cannot validly represent, so it too reflects coverage rather than a like-for-like contrast of local dynamics. Threshold sensitivity showed that the headline BigBrain correlation was stable across distance cutoffs. Across 5, 10, 15, 20, 25, and 30 mm nearest-surface thresholds, BigBrain contributed 8,832/8,947/8,948 analysable rows at 5/10/15 mm and 8,948 thereafter, from 8,869/8,984/8,985 coordinate rows at those thresholds. Its full-sample raw Spearman correlation was already positive at 5 mm ($\rho \approx +0.127$) and was approximately $+0.133$ from 10–30 mm. By contrast, reference-surface retention grew from only 39 geometrically valid rows at 5 mm to 2,857 rows at 30 mm, and its estimate was correspondingly support-dependent.

The advantage does not reduce to mesh density either (Fig. 2f). At reference-comparable resolution (BigBrain 30k), sampled at the identical BigBrain-retained electrodes, the 30k surface remained above the reference (raw $\rho = +0.030$, partial $+0.017$ versus raw -0.128 , partial -0.104), with a subject-paired advantage in 86/114 subjects (mean $\Delta = +0.177$). The pooled AP-controlled effect at 30k is itself weak, consistent with the within-BigBrain result that alignment is non-monotonic in vertex count and peaks at intermediate resolution. Across the BigBrain family, AP-partial alignment was positive at all four resolutions but non-monotonic in vertex count ($+0.017$, $+0.159$, $+0.138$, and $+0.153$ at 30k, 60k, 100k, and 164k), peaking at the intermediate 60k surface. If additional resolution had improved alignment monotonically, resolution would behave as a quality parameter, and the representation would be an implementation detail to be maximized rather than a modelling decision to be justified. It does not, and the state-dependent resolution effects reported in Sec. 2.1 make the same point from the other direction: sleep favours fine resolution, navigation the coarse 30k surface.

Two classes of trivial explanation are excluded directly. Across 114 subjects BigBrain exceeded the reference in 86/114 for both raw and partial ρ (mean paired $\Delta = +0.283$ and $+0.286$; Wilcoxon $p = 4.3 \times 10^{-11}$; Fig. 2d), so the pooled effect is not carried by a minority of subjects. And the alignment is neither spatial autocorrelation nor an AP-only gradient: a global-shuffle null gave a 95% interval of $[-0.021, +0.020]$ and the stricter AP-bin-shuffle null $[+0.041, +0.067]$, both one-sided $p = 0.0005$ against the observed $+0.153$ (Fig. 2e). Holm family-wise-error and Benjamini–Hochberg false-discovery-rate corrections were applied separately to three pre-defined nine-test families: pooled raw Spearman alignment, pooled AP-partial Spearman alignment, and subject-level raw alignment against zero; the saturated $cc = 0.080$ operating point was excluded from these multiplicity families. Positive support after correction was restricted to weak coupling ($cc \in \{0, 0.002, 0.005\}$; canonical $cc = 0$ values raw $\rho = +0.133$, AP-partial $\rho = +0.153$). Several higher-coupling operating points remained significant after correction, but only as negative alignments. A 10,000-permutation within-subject omnibus test confirmed an overall operating-point effect ($Q = 107.692$, $p = 9.999 \times 10^{-5}$), without changing the interpretation that the sleep sweep is exploratory and that stronger coupling does not improve theta-map validity.

Together, these results support the uncoupled high-resolution BigBrain geometry as an empirical sleep-theta substrate and locate its advantage in geometrically valid coverage rather than superior local prediction. This uncoupled condition therefore provides the empirical baseline perturbed in the state-gating analysis.




Figure 2: The uncoupled high-resolution substrate aligns with empirical sleep theta in ds004551 and adult Penn SIEVE N3, and its advantage is geometric coverage rather than superior local dynamics. **a**, Coverage at the 15 mm nearest-surface threshold. BigBrain retained all 8,985 coordinate rows within threshold (8,948 analysable; median 1.39 mm); the reference surface [30] placed 1,385/8,985 (1,379 analysable; median retained distance 11.31 mm; median over all coordinates before gating 43.98 mm). **b**, Same-row contrast on the BigBrain-retained rows, extrapolating the reference surface beyond its validity threshold. **c**, Coverage decomposition into all rows ($n = 8,948$), the shared-valid set ($n = 1,379$), and the BigBrain-only set ($n = 7,563$). **d**, Subject-paired advantage in AP-partial ρ across 114 subjects. **e**, Observed AP-partial alignment against global-shuffle and AP-bin-preserving nulls. **f**, Within-BigBrain resolution series (30k, 60k, 100k, 164k) with the density-matched reference. **g**, Adult Penn SIEVE N3 contact-level check, in the clinically normal-like stratum ($n = 1,880$) and across all N3 contacts ($n = 5,388$). Values and tests are reported in Sec. 2.3.

2.4 The state-gated alignment is not a property of the surface alone

A positive alignment on a distinctive surface invites an obvious deflation: perhaps any dynamics simulated on this geometry would align, so the surface, not the neural field model, does the work. We show instead that a generic NMM simulator [16] on the identical geometry is null in both states, placing the alignment in the model class rather than in the surface.

We simulated a standard stochastic Wilson–Cowan NMM on the identical BigBrain geometry (30k surface; global coupling 0.014; additive noise 1×10^{-5} ; 60 s with 2 s burn-in) and passed it through the same frozen pipelines (Fig. 3a). The baseline produced a non-degenerate theta map (theta-power SD 1.2×10^{-5} ; 4,113 unique sampled values), so a null result reflects a genuine null rather than a dead simulation. All models in this section were compared on the common Wilson–Cowan/BigBrain 15 mm intersection, slightly more restrictive than the BigBrain-only intersection of Sections 2.1 and 2.3, so pooled values differ from those sections by small amounts.

In sleep, the NMM baseline was null (partial $\rho = +0.001$, $p = 0.93$), whereas the uncoupled neural field at 164k was positive (+0.153) and the reference surface [30] was wrong-sign (-0.104 on the same contacts; Fig. 3b). The neural field exceeded the baseline subject-paired at full resolution (86/114; Wilcoxon $p = 1.6 \times 10^{-10}$; Fig. 3d) but not robustly at matched 30k resolution (66/114; $p \approx 0.10$ raw, 0.16 partial), so the decisive sleep margin combines model class with resolution. The reference surface fell below even the null baseline (the NMM exceeded it in 87/114 subjects for raw ρ), so it is not merely uninformative but actively inverted relative to a generic simulator.

The same baseline was near-null against navigation theta in every phase and dataset (AP-partial ρ from -0.005 to -0.006 ; $p = 0.56$ to 0.69). Positive navigation alignment came from exactly two neural field operating points on this intersection (Fig. 3c): the coarse uncoupled 30k surface and the full-resolution $cc = 0.030$ coupled model. In the pooled contact-level TVB/BigBrain intersection, these values were +0.107 and +0.213 for the 30k surface (ds005522 all-task and ds005523 retrieval), and +0.097 and +0.212 for the $cc = 0.030$ model. The corresponding subject-level means used for paired inference were +0.139/ +0.134 for the 30k surface and +0.134/ +0.209 for the $cc = 0.030$ model, matching the state-gating aggregation. Each operating point exceeded the baseline subject-paired (30k: Wilcoxon $p = 1.6 \times 10^{-4}$ and 1.4×10^{-3} ; $cc = 0.030$: $p = 1.5 \times 10^{-3}$ and 8.0×10^{-3} ; Fig. 3d), whereas the weaker $cc = 0.005$ model and the uncoupled 164k model did not.

Across both states, a NMM baseline on identical geometry is near-null, placing the state-gated alignment in the neural field model class rather than in the surface alone or in any simulator placed on it. We bound this claim deliberately: the tested Wilson–Cowan configuration (one coupling value, one noise setting, one seed) is null here.




Figure 3: **The state-gated alignment is specific to the neural field model class.** **a**, Contact-sampled theta-power maps for the Wilson-Cowan baseline, two neural field operating points (30k uncoupled and $cc = 0.030$), and their model-class difference ($cc = 0.030$ minus Wilson-Cowan), on identical BigBrain geometry and identical coordinate sampling. **b**, Sleep AP-partial alignment across model class, resolution, and the reference surface [30]. **c**, Navigation AP-partial alignment across model class and operating point in both cohorts. **d**, Subject-paired baseline tests: percentage with neural field > Wilson-Cowan against $-\log_{10}$ Wilcoxon p . All models are compared on the common Wilson-Cowan/BigBrain 15 mm intersection. Values are reported in Sec. 2.4.

2.5 Geometry and connectivity dominate different spatial scales in task fMRI

Geometric eigenmodes are taken to provide a more parsimonious account of cortical activity than connectivity-based modes [9], implying a fixed hierarchy in which geometry is the dominant constraint. We show instead that geometry and connectivity do not form a fixed hierarchy: their dominance reverses across spatial scale, with connectivity better reconstructing the longest-wavelength patterns and high-resolution geometry better reconstructing finer ones, with a crossover near ≈ 93 mm.


We reproduced and extended a prior geometric-basis HCP task-fMRI reconstruction [9] to

BigBrain-resolution geometry and distance-dependent (exponential-distance-rule, EDR) connectome modes, in a 2×2 design crossing data resolution and basis type (Methods, Sec. 4.11; Fig. 4a). Each basis was projected under the metric in which it is orthonormal: the BigBrain geometric basis under the lumped surface-mass metric, and the EDR basis under the identity metric of its graph Laplacian, for which the reduced inverse is the transpose. We additionally repeated the comparison with the EDR basis projected under the BigBrain cortical mass metric, leaving the geometric cell unchanged, to separate the basis comparison from the choice of projector. Reconstruction accuracy rose with mode count in every condition, reaching mean $r \approx 0.53$, 0.87 , and 0.93 at $N = 10$, 100 , and 200 modes across 47 task contrasts (Fig. 4b), and BigBrain geometric modes exceeded the per-contrast BigBrainWarp round-trip-fidelity benchmark in 43/47 contrasts at $N = 200$ modes (benchmark median $r = 0.928$, range 0.802 – 0.948 ; median residual $+0.020$; $p = 7.8 \times 10^{-13}$; Fig. 4e), so the high-resolution basis is not limited by surface-registration error. The round-trip scalar is a mean summary of a frequency-dependent operation: the same infrastructure test records attenuation of short-wavelength content and a small amplitude reduction under transfer. Because the scale-resolved result concerns the short-wavelength regime, the geometric advantage at $N = 100$ and $N = 200$ is established on task maps that have themselves passed through a transfer step known to attenuate fine spatial structure. This bounds the magnitude of the short-wavelength effect rather than its direction, since both bases reconstruct the same transferred maps.


The two bases did not form a fixed hierarchy. Their advantage reversed across spatial scale (Fig. 4c). EDR-connectome modes better reconstructed broad, long-wavelength patterns from ≈ 312 to ≈ 148 mm (median geometry–connectivity = -0.048 , -0.014 , and -0.005 at $N = 5$, 10 , and 20 ; Holm-corrected $p = 1.4 \times 10^{-6}$, 2.0×10^{-4} , and 0.013), the two bases were statistically indistinguishable at ≈ 93 mm ($N = 50$: median -0.004 , $p = 0.88$), and BigBrain geometric modes dominated finer patterns at ≈ 66 mm and below ($N = 100$: median $+0.010$, $p = 3.0 \times 10^{-6}$; $N = 200$: median $+0.005$, $p = 4.6 \times 10^{-11}$). The sign of the difference therefore changes once, between ≈ 93 and ≈ 66 mm. The reversal was consistent across task contrasts (Fig. 4d) and persisted when both bases were projected under the same BigBrain cortical mass metric (median -0.048 to -0.005 at $N = 5$ – 20 , $+0.011$ and $+0.006$ at $N = 100$ and $N = 200$; Holm-corrected $p < 0.05$ at every mode count except the crossover). Because the geometric reconstruction is identical between the two runs by construction, the reversal does not reduce to a projection-normalization choice.

It is therefore more accurate to say that geometry and connectivity dominate different spatial scales than that either is categorically dominant: connectivity carries the long-wavelength, brain-wide component of task activity, whereas high-resolution geometric modes better reconstruct the finer-scale component.


a Task map and basis-specific reconstructions




b 2 x 2 reconstruction curves



c Scale-dependent reversal



d Per-contrast geometry - EDR split



e BigBrain exceeds warp ceiling




Figure 4: Geometry and connectivity dominate different spatial scales in task fMRI. **a**, A representative task map (social theory-of-mind) with $N = 10$ -mode geometry and EDR-connectivity reconstructions and their residual, on identical BigBrain geometry. **b**, Mean reconstruction accuracy across 47 contrasts for the 2×2 conditions (fsLR and BigBrain $\times$ geometry and EDR) against mode count N , with the corresponding BigBrain wavelength on the upper axis. **c**, Scale-dependent reversal of the median geometry–connectivity difference. Open markers show the matched-metric replication, in which the connectome basis is projected under the BigBrain cortical mass metric ($*p < 0.05$, $**p < 0.01$, $***p < 0.001$, Holm-corrected across the six evaluated mode counts). **d**, Per-contrast geometry–connectivity difference across mode count. **e**, BigBrain geometric reconstruction against the per-contrast BigBrainWarp round-trip-fidelity benchmark at $N = 200$ modes. Values are reported in Sec. 2.5.

3 DISCUSSION

These results address three linked questions about how anatomy constrains cortical dynamics: whether a fixed anatomical pathway makes a fixed functional contribution, whether high-resolution histological geometry provides an empirically supported substrate rather than merely an implementation detail, and whether geometric or connectivity-derived modes provide the better account of cortical activity. The answers differ in form: the relay contribution is state-dependent, the substrate advantage is conditional on empirical coverage, and the relative contribution of geometry and connectivity depends on spatial scale. The last question also bears directly on the apparent disagreement between geometric and connectome-based accounts of cortical organisation [9, 11].

On the first question, our results do not support a state-invariant functional contribution of the tested anatomical relay. The same anatomically informed hippocampal–cortical relay had opposite effects across brain states: it degraded intrinsic sleep-theta alignment, but supplied positive navigation-theta alignment in the tested full-resolution model. This dissociation was observed in two independent navigation cohorts, survived AP-preserving nulls and coverage-matched controls, and replicated in direction across three independent stochastic realizations. Hippocampal relay coupling is therefore a state-gated drive whose contribution to empirical theta alignment depends on brain state.

The spectral regime of the coupled system provides an important bound on this interpretation. In the parent cortico-hippocampal neural-field model [30], the cortical system intrinsically expresses alpha-band dynamics while the hippocampal system expresses theta-band dynamics. The EC-reflow simulations reproduced this qualitative cortical-alpha/hippocampal-theta regime after substitution onto the BigBrain geometry, showing that the principal spectral organisation of the parent model was retained rather than lost under the new anatomical realization (Supplementary Methods, “EC-reflow sweep spectral-regime validation”). The state-gating result therefore arose within the qualitative cortico-hippocampal spectral regime of the parent NFT model rather than from a loss of its basic spectral organisation.

The exploratory right-hemisphere benchmark sharpened this interpretation in a direction consistent with a distinction previously established in NFT. Gabay and Robinson showed that eigenmode spatial orientations are constrained by gross cortical shape, whereas eigenvalues (and hence eigenfrequencies and the resulting dynamical properties) depend more sensitively on finer-scale cortical structure [10]. The two hemispheric realizations of this donor were consistent with that distinction. Their cortical eigenvalue spectra were strongly homologous after accounting for surface area (log-spectrum correlation 0.99994), and their dynamics were nearly indistinguishable through $cc = 0.020$, yet the upper cortical dominant-frequency transition fell in different coupling intervals. The frozen $cc = 0.030$ operating point therefore sampled different cortical operating regimes on the two sides while hippocampal peak frequencies remained nearly identical (Supplementary Methods, “Exploratory right-hemisphere lateralization sensitivity analysis”). Spectrally homologous cortical Laplacians therefore participate in hemisphere-specific coupled systems with different dynamical transition intervals. This extends to realizations within one brain (an observation previously made across individuals): that the coupling value at which dynamics transition depends on the anatomical realization [40]; here donor, model parameters, simulator, random seed, and analysis pipeline are all held fixed.

Empirical analysis across hemispheres showed partial preservation of the state-dependent coupling pattern. The available sleep-side checks preserved the negative coupling direction, whereas the navigation result varied across datasets. In ds005522, the right-hemisphere subject-level slope reversed sign, while in ds005523 it retained the original positive direction at attenuated magnitude (Supplementary Methods, “Exploratory right-hemisphere lateralization sensitivity analysis”). Be-

cause cortical geometry, hippocampal geometry, surface masses, projected coupling operators, and empirical electrode coverage vary together between hemispheres, this comparison does not identify which anatomical component produced the difference and should not be interpreted as a formal test of biological lateralization. The benchmark therefore demonstrates sensitivity to anatomical realization.

Because coupling changed the model’s spectral regime, we also tested whether the sleep-side decrease could be removed by simple cortical retuning. This control builds on a characteristic property of the parent cortico-hippocampal NFT model: increasing coupling shifts and sharpens the dominant cortical and hippocampal spectral peaks [30]. The same qualitative behaviour was reproduced on this BigBrain implementation, where increasing EC-refrow coupling shifted and sharpened the dominant spectral peaks through the moderate-coupling range (Supplementary Methods, “EC-refrow sweep spectral-regime validation”). Parameter selection used simulated spectra and stability diagnostics only; empirical alignment was not used during selection. Across uniform synaptic-rate, beta/alpha-ratio, and recurrent-gain families, no parameter family produced the complete theta-preserving coupling grid (Methods, Sec. 4.12). Under uniform synaptic-rate scaling, the dominant cortical peak shifted directly from the subtheta branch to the alpha branch at the weak-coupling values rather than passing continuously through theta. Within these three bounded one-parameter directions, the sleep-side decrease is therefore not explained by a trivial frequency-offset effect arising from one mistuned cortical parameter. Multidimensional retuning and explicitly state-dependent coupling remain future extensions.

On the second question, whether high-resolution histological geometry provides an empirically supported substrate rather than merely an implementation detail, the result is more nuanced. The uncoupled BigBrain histological surface provides an empirically supported sleep-theta substrate, so the null position that the histological substrate adds nothing is not supported. But it does not predict theta better than the population reference where both surfaces validly sample the same tissue, so the stronger claim that high-resolution histology improves local prediction is not supported either. What the BigBrain histological surface supplies in this comparison is valid sampling of cortical regions that the population reference surface does not adequately reach.

The substrate and model-class controls locate this advantage more precisely. The uncoupled BigBrain neural-field model produced positive group-level alignment with empirical sleep theta in ds004551 and showed the same positive AP-controlled direction in adult Penn SIEVE N3, whereas the lower-resolution reference surface was sign-reversed in ds004551 and near-null or wrong-sign in Penn SIEVE. This advantage should be read carefully. It was not evidence that BigBrain improves local theta-map prediction on the subset of contacts both surfaces validly sample; on that subset, both surfaces were negative. The positive effect was concentrated in contacts validly sampled only by BigBrain. The advantage is therefore coverage-mediated: the BigBrain histological surface improves the empirical interface between model and electrodes by extending valid cortical coverage that the lower-resolution surface does not adequately reach. Coverage is thus part of the empirical validity of the model–data interface, but it should be distinguished from local predictive accuracy at shared locations.

At the same time, a Wilson–Cowan NMM baseline on identical BigBrain geometry was near-null in both states, showing that the alignment is not a property of the surface alone. In sleep, that baseline exceeded the population reference surface in 87/114 subjects for raw ρ , placing the reference below even the null comparator. The result therefore depends on the dynamical formulation operating on the geometry, although the Wilson–Cowan comparison was configuration-specific and broader model-class comparisons will be required to establish how specific the effect is to NFT models.

On the third question, whether geometry or connectivity is the dominant constraint, neither

position is categorically correct. Geometric eigenmodes were reported to be more parsimonious than connectome modes [9]; long-range connections were reported to contribute over and above geometry [11]. Our scale-resolved comparison shows that the connectome basis reconstructs the longest-wavelength task patterns, whereas histological geometry better reconstructs the finer ones. The apparently conflicting conclusions can therefore be reconciled once spatial scale is resolved, and the graded picture reported subsequently [12] is consistent with such a scale-resolved account rather than requiring a fixed hierarchy between geometry and connectivity.

The fMRI analysis places this result in a broader multiscale account. Distance-dependent connectome modes better reconstructed broad, long-wavelength task patterns, whereas BigBrain geometric modes dominated finer spatial patterns. This scale reversal provides a common interpretation of the electrophysiological and fMRI results: high-resolution geometry provides the field substrate, while connectivity and relay coupling shape how activity on that substrate is expressed across spatial scales and states. It also explains how two apparently opposed positions in the literature can both be defended with correct evidence. A scale-aggregated summary can obscure this reversal because its value depends on the spatial scales represented in the evaluation. Scale dependence can therefore reconcile apparently opposed conclusions drawn from different scale ranges without requiring either basis to be generally superior.

The electrophysiological resolution controls connect this scale dependence back to the state-gating result. Within the uncoupled BigBrain family, sleep favoured intermediate-to-fine resolution, whereas navigation alignment was strongest on the coarse 30k surface. Positive navigation alignment also emerged through EC-reflow coupling on the full-resolution substrate. Coupling and coarsening therefore provide distinct routes to positive navigation alignment without implying that they are mechanistically identical. Taken together with the fMRI crossover, these findings support a common scale-dependent interpretation: the effective spatial scale at which the cortical substrate is read out is state-dependent. The resulting picture is therefore relational rather than categorical: substrate, scale, and drive jointly shape empirical alignment.

Several limitations define the next tests. First, the model represents tangential surface geometry through Laplace–Beltrami eigenmodes but omits cortical depth, hippocampal subfields, layer-specific connectivity, cytoarchitectural gradients, and depth-dependent current generators. Because iEEG signals are depth-sensitive, agreement here should be read as macroscopic spatial alignment, not as a laminar mechanism. A natural extension is a multilayer or volumetric field model built from BigBrain-derived laminar surfaces.

Second, the geometric substrate is a single 65-year-old donor, not a population-average or participant-specific template. Registration into a common space does not remove donor-specific folding and surface-area effects. The exploratory right-hemisphere benchmark demonstrates portability to a second within-donor anatomical realization and exposes sensitivity to that realization, but it does not address between-donor generalization. Moreover, unmatched right-hemisphere electrode coverage in ds004551 prevents replication of the primary subject-level sleep statistic and therefore precludes a formal biological lateralization claim (Supplementary Methods, “Exploratory right-hemisphere lateralization sensitivity analysis”). Generalization will require multiple histological templates or participant-specific surfaces.

Third, the EC-reflow relay improved theta-map alignment but did not pass the pairwise target-coupling validation. Bridging macroscopic map alignment to empirical pairwise connectivity will likely require dynamic coupling weights, state-dependent delays, and additional relays such as thalamic or septal pathways.

Fourth, the Wilson–Cowan control was configuration-specific. A single tested neural-mass configuration cannot establish that the observed alignment is unique to neural-field models, and broader comparisons across dynamical model classes and parameterizations will be required to determine

the specificity of the result.

Finally, the navigation effects were estimated from modest cohort sizes ($n = 49$ and $n = 19$). Although downsampling the ds004551 sleep anchor to matched subject counts preserved the negative sleep slope, the incomplete right-hemisphere transfer shows that the navigation effect is also sensitive to anatomical realization. Replication in larger navigation cohorts with broader and bilaterally matched coverage will therefore be important for establishing how robustly the state-dependent coupling effect generalizes across subjects, datasets, and anatomical realizations.

The most defensible conclusion is therefore layered. The BigBrain histological geometry provides an empirically supported field substrate with broader valid electrode coverage than the population reference surface. Connectivity-derived and geometric modes preferentially reconstruct different spatial scales of task-evoked activity. Hippocampal relay coupling provides a state-gated drive that decreases intrinsic sleep-theta alignment but increases navigation-theta alignment in the tested full-resolution model. The relevant scientific unit is therefore not geometry versus connectivity, or coupling versus no coupling, but the state-dependent interaction between substrate, scale, and drive.

4 METHODS

4.1 Study design and analysis layers

The study was organized around three questions: whether the functional contribution of a fixed anatomical relay changes across brain states, whether high-resolution histological geometry provides an empirically supported substrate for cortical dynamics, and whether geometry and connectivity account for different spatial scales of brain activity. These questions were addressed through four linked analyses.

First, we tested state-dependent hippocampal–cortical coupling using cortico-hippocampal neural-field models constructed on BigBrain-resolution cortical and hippocampal surfaces. An anatomically informed entorhinal relay selected from the preceding sleep architecture comparison was carried forward as a frozen coupling axis and evaluated across predominantly pediatric sleep, adult N3 sleep, and two independent navigation-memory iEEG cohorts. Second, we tested the geometric substrate itself by comparing uncoupled BigBrain predictions with a population reference surface, within-BigBrain downsampling controls, common-support and coverage analyses, and spatial nulls. Third, we tested whether the resulting alignment was a property of the surface alone using a Wilson–Cowan neural mass model on the identical BigBrain geometry. Finally, we tested the geometry-versus-connectivity question in task fMRI by comparing geometric and EDR connectome eigenmodes across spatial scales.

The electrophysiological validation was map-based. For each model condition, a cortical spatial map of band-limited power was sampled at empirical electrode contact coordinates. Raw Spearman correlation quantified model–empirical theta-map alignment, and AP-partial Spearman correlation quantified alignment after controlling for the dominant anterior–posterior coordinate. The state-gating analysis used the per-subject slope of AP-partial ρ against coupling strength cc as its primary sweep statistic, together with pre-specified fixed- cc contrasts. Subject-level inference used one summary statistic per subject, so that each subject contributed a single observation to group-level tests rather than individual contacts being treated as independent samples. For the primary sweep analysis, the resulting per-subject slope distributions are shown in Fig. 1c.

4.2 BigBrain cortical and hippocampal surfaces

Cortical surface. We used the BigBrain left-hemisphere white-matter surface registered to MNI152NLin2009bSym space [1]. The cortical mesh contained 163,842 vertices and 327,680 triangular faces. The verified left-surface area was approximately 134,239 mm² (1342.4 cm²). This histological surface preserves single-donor folding and metric structure that are smoothed in population-average surfaces [41], while remaining in a coordinate frame suitable for sampling independent iEEG and fMRI data [17].

Hippocampal surface. For the BigBrain cortico-hippocampal simulations we used a HippUnfold-derived [42] left hippocampal surface from the BigBrain specimen (**sub-bbhist**), aligned into MNI152NLin2009bSym space. The final hippocampal mesh contained 32,004 vertices and 63,250 triangular faces. Because the native HippUnfold output is in **corobl** orientation, we rigidly aligned the hippocampal surface to the MNI152NLin2009bSym reference hippocampal midthickness surface using iterative closest point registration. The final alignment had symmetric mean surface distance 1.119 mm, RMS distance 1.692 mm, p95 distance 3.032 mm, and approximate Hausdorff distance 11.530 mm. Five Taubin smoothing iterations reduced sliver triangles prior to eigendecomposition, with mean displacement 0.052 mm. Smoothing reduced triangles with minimum angle $< 5^\circ$ from 36 to 4, and triangles with minimum angle $< 15^\circ$ from 4,299 to 3,048.

Reference surface. The lower-resolution reference implementation [30] used a population cortical mesh with 29,708 vertices and 59,412 faces, and a population HippUnfold hippocampal mesh with approximately 7,246 vertices, retaining 110 cortical and 25 hippocampal modal coefficients in its uncoupled configuration. We refer to it throughout as the lower-resolution reference surface. It differs from BigBrain simultaneously in surface source, vertex density, coordinate support, and electrode-retention properties, so it serves as an alternative-substrate comparator and implementation reference rather than isolating the contribution of any single factor (Table 2).

Table 2: Surface and modal comparison for the reference implementation and the BigBrain neural-field models.

Quantity	Reference implementation	BigBrain model	Ratio
Cortex vertices	29,708	163,842	5.5×
Cortex faces	59,412	327,680	5.5×
Cortex modes retained	110	300	2.7×
Hippocampal vertices	$\approx 7,246$	32,004	4.4×
Hippocampal modes retained	25	300	12.0×
Modal first-order state dimension	270	1,200	4.4×
Cortex source	population surface	BigBrain donor	–
Hippocampus source	population HippUnfold	BigBrain HippUnfold	–
Implementation	MATLAB/CPU	Python/CuPy/GPU	–

4.3 Laplace-Beltrami eigenmodes and mass-weighted projection

We represented cortical and hippocampal fields in eigenmodes of the Laplace-Beltrami operator on each anatomical surface. For each mesh we solved the generalized eigenproblem

$$L\mathbf{v}_n = \lambda_n M\mathbf{v}_n, \quad (1)$$

where L is the cotangent stiffness matrix and M is the lumped mass matrix, with each diagonal entry equal to one third of the adjacent face areas. This formulation gives eigenvalues in mm^{-2} and preserves the physical spatial scale that appears in the neural-field wave equation. A standard unweighted eigenproblem, $L\mathbf{v} = \lambda\mathbf{v}$, would instead absorb mesh sampling density into the eigenvectors and produce dimensionless eigenvalues inappropriate for comparing surfaces or downsampled variants [43].

The full-resolution BigBrain left cortical and hippocampal bases retained 300 modes each and were generated with `torch.linalg` in double precision (generalized mass matrix, seed 42, tolerance 10^{-6}). The Figure 2 cortical resolution bases were generated with SciPy ARPACK smallest-magnitude solves for 30k and 60k (`which=SM`, tolerance 10^{-6}), and a shift-invert solve for 100k ($\sigma = 0$, tolerance 10^{-7}). The 164k variant reused the full-resolution basis. Eigenmodes, eigenvalues, and reduced inverse operators were stored as NumPy arrays. In general, vertex-space fields $y \in \mathbb{R}^N$ are projected into a retained mode basis by

$$P = \left(\Phi^\top M \Phi \right)^{-1} \Phi^\top M, \quad a = Py, \quad \hat{y} = \Phi a. \quad (2)$$

The simplification $P = \Phi^\top M$, equivalently $a_n = \sum_j \psi_{jn} M_{jj} y_j$, is valid only when the retained basis is exactly M -orthonormal ($\Phi^\top M \Phi = I$). At full resolution, the BigBrain reduced inverse operators use this mass-weighted projector: projectors recomputed from the meshes agree with the stored arrays to relative Frobenius errors of 1.74×10^{-13} for cortex and 1.80×10^{-11} for hippocampus. This projection is used both for dynamical simulation and for reconstructing task-fMRI maps in BigBrain space. The resolution controls used three downsampled BigBrain cortical meshes (30k, 60k, and 100k) and the unchanged 164k mesh, all in the BigBrain coordinate frame.

4.4 Neural-field dynamical model

We simulated cortico-hippocampal activity with a stochastic neural-field model projected onto coupled cortical and hippocampal Laplace–Beltrami eigenbases. The model follows the cortico-hippocampal NFT structure of [30], but replaces the cortical and hippocampal geometry with BigBrain-resolution surface eigenmodes. Linear field propagation is advanced in modal coordinates; nonlinear firing rates, local synaptodendritic transforms, and inter-surface coupling are evaluated in vertex space when needed and then projected back into mode space. This division is what makes vertex-resolution nonlinear simulation tractable: the spatial basis carries the surface geometry, while the time-stepping state dimension is controlled by the retained mode count.

Uncoupled baseline model. The uncoupled baseline is the corticothalamic component of the cortico-hippocampal neural-field system, evaluated on the BigBrain cortical Laplace–Beltrami eigenmodes. The excitatory cortical field $\phi_e^c(\mathbf{r}, t)$ obeys a damped wave equation,

$$\left[\frac{1}{\gamma_c^2} \frac{\partial^2}{\partial t^2} + \frac{2}{\gamma_c} \frac{\partial}{\partial t} + 1 - r_c^2 \nabla^2 \right] \phi_e^c(\mathbf{r}, t) = Q_e^c(\mathbf{r}, t), \quad (3)$$

where $\gamma_c = v_c/r_c$ is the damping rate, v_c is the axonal conduction velocity, and r_c is the characteristic cortical spatial range. Expanding the field in cortical eigenmodes,

$$\phi_e^c(\mathbf{r}, t) = \sum_n \zeta_n^c(t) R_n^c(\mathbf{r}), \quad \nabla^2 R_n^c = -\lambda_n^c R_n^c, \quad (4)$$

reduces Eq. 3 to independent modal wave equations,

$$\left[\frac{1}{\gamma_c^2} \frac{d^2}{dt^2} + \frac{2}{\gamma_c} \frac{d}{dt} + 1 + r_c^2 \lambda_n^c \right] \zeta_n^c(t) = \kappa_n^c(t), \quad (5)$$

637 where $\kappa_n^c(t)$ is the modal projection of the source term. Thus BigBrain geometry enters the uncou-
 638 pled cortical dynamics through the eigenfunctions R_n^c and eigenvalues λ_n^c .
 639 Firing rates followed the sigmoid

$$Q_j^c(\mathbf{r}, t) = \frac{Q_{\max}}{1 + \exp[-(V_j^c(\mathbf{r}, t) - \theta_c)/\sigma_c]}, \quad j \in \{e, i, r, s\}, \quad (6)$$

640 with transmembrane potentials generated from presynaptic input by the second-order synaptoden-
 641 dritic filter

$$\left[\frac{1}{\alpha_c \beta_c} \frac{d^2}{dt^2} + \left(\frac{1}{\alpha_c} + \frac{1}{\beta_c} \right) \frac{d}{dt} + 1 \right] V_j^c(\mathbf{r}, t) = P_j^c(\mathbf{r}, t). \quad (7)$$

642 The corticothalamic presynaptic sums were

$$\begin{aligned} P_e^c &= \nu_{ee}^c \phi_e^c + \nu_{ei}^c Q_i^c + \nu_{es}^c Q_s^c(t - \tau_c), \\ P_r^c &= \nu_{rs}^c Q_s^c + \nu_{re}^c \phi_e^c(t - \tau_c), \\ P_s^c &= \nu_{se}^c \phi_e^c(t - \tau_c) + \nu_{sr}^c Q_r^c + \nu_{sn}^c (\phi_{n0} + \eta_c), \end{aligned} \quad (8)$$

643 where τ_c is the corticothalamic delay and η_c is stochastic nonspecific input. In coupled simulations,
 644 hippocampal activity contributes to the cortical excitatory source through

$$Q_e^c(\mathbf{r}, t) \mapsto Q_e^c(\mathbf{r}, t) + cc \int w_{ch}(\mathbf{r}, \mathbf{r}') \phi_e^h(\mathbf{r}', t - \tau_{ch}) d^2 \mathbf{r}', \quad (9)$$

645 with a reciprocal cortex-to-hippocampus term scaled by ch . Here w_{ch} is the hippocampus-to-
 646 cortex operator. In uncoupled simulations $cc = ch = 0$, so both cross-region drive terms vanish
 647 exactly. The cortical theta-power map evaluated against empirical data is therefore generated by
 648 the corticothalamic cortical system alone; the decoupled hippocampo-septal system contributes no
 649 cortical drive.

Table 3: Neural-field model parameters used for the BigBrain cortico-hippocampal simulations. Values are the fixed baseline parameters inherited from the cortico-hippocampal neural-field model [30], with the spatial parameters expressed in the millimetre-consistent convention used by the BigBrain implementation.

Parameter	Symbol	Cortex	Hippocampus	Unit / note
Axonal conduction velocity	v	11600	5.8	mm s ⁻¹
Characteristic spatial range	r_e	100	0.1	mm
Damping rate	$\gamma = v/r_e$	116	58	s ⁻¹
Maximum firing rate	$Q_{\max}$	340	340	s ⁻¹
Firing threshold	θ	12.92	129.2	mV-equivalent
Threshold width	σ	3.4	30.5	mV-equivalent
Synaptodendritic rise rate	α	83.33	83.33	s ⁻¹
Decay-to-rise ratio	β/α	9.23	9.23	unitless
Synaptodendritic decay rate	β	769.2	769.2	s ⁻¹
Delay	τ	0.040	0.005	s
Excitatory–excitatory gain	ν_{ee}	3.03	3.03	mV·s-equivalent
Inhibitory–excitatory gain	ν_{ei}	−6.00	−6.00	mV·s-equivalent
Relay–excitatory gain	ν_{es}	2.06	2.06	mV·s-equivalent
Excitatory–relay gain	ν_{se}	1.65	1.65	mV·s-equivalent
Reticular–relay gain	ν_{sr}	−0.83	−0.83	mV·s-equivalent
Nonspecific–relay gain	ν_{sn}	0.98	0.98	mV·s-equivalent
Excitatory–reticular gain	ν_{re}	0.33	0.33	mV·s-equivalent
Relay–reticular gain	ν_{rs}	0.03	0.03	mV·s-equivalent
Nonspecific input baseline	ϕ_{n0}	1.0	1.0	unitless
Default integration step	Δt	5×10^{-5}		s
Simulated duration	T	60		s
Burn-in excluded at save		2		s
Further burn-in at extraction		5		s
Saved sampling interval		1		ms
Default retained modes	N	300	300	modes
Default random seed		73924		seed
Projector	P_C, P_H	mass-weighted Galerkin projector		default
Default coupling operator		fully modal $PW\Phi$		default
EC-refrow sweep used for slope analysis	cc	0, 0.002, 0.005, 0.008, 0.010, 0.015, 0.020, 0.030, 0.050, 0.080		Fig. 1
Confirmatory fixed points	cc	0.005 and 0.030 versus 0		endpoint checks
Reciprocal coupling	ch	$ch = 0.2 cc$		sweep

Internal simulator values store cortical thresholds and gains in volts rather than mV, e.g. $\theta_C = 12.92/1000$ and $\nu_{ee}^C = 3.03/1000$, while hippocampal values follow the scaling convention of [30]. The table reports reader-facing mV-equivalent values. The millimetre-consistent cortical spatial convention uses $v_C = 11600$ mm s⁻¹ and $r_e^C = 100$ mm, giving the same cortical damping rate $\gamma_C = 116$ s⁻¹ as the original metre-scale notation.

4.5 Numerical integration and simulation outputs

The GPU simulator (Python/CuPy) reads the stored cortical and hippocampal bases and truncates each to the requested 300 modes. The cortical and hippocampal projectors are the mass-weighted operators

$$P_C = \left(\Phi_C^\top M_C \Phi_C \right)^{-1} \Phi_C^\top M_C, \quad P_H = \left(\Phi_H^\top M_H \Phi_H \right)^{-1} \Phi_H^\top M_H, \quad (10)$$

which reduce to $\Phi_C^\top M_C$ and $\Phi_H^\top M_H$ only for exactly M -orthonormal retained bases. To avoid repeated large matrix products during time-stepping, cross-region coupling is preprojected into modal coordinates. We use direction-explicit notation:

$$A_{C \leftarrow H} = P_C W_{C \leftarrow H} \Phi_H, \quad A_{H \leftarrow C} = P_H W_{H \leftarrow C} \Phi_C, \quad (11)$$

where Φ_C and Φ_H are vertex-by-mode basis matrices, and $W_{C \leftarrow H}$ and $W_{H \leftarrow C}$ are sparse vertex-space coupling matrices.

The EC-reflow sweep, the seeded endpoint runs, and the 30k control run were integrated with Euler–Maruyama at $\Delta t = 5 \times 10^{-5}$ s. Simulations are divided into 0.05-s epochs so that delayed state variables can be preserved across epochs without allocating the full trajectory in memory. The first 2 s are excluded from the saved record (initial-condition burn-in) and saved outputs are downsampled to 1 ms resolution. Map extraction discards a further 5 s as a conservative margin before spectral estimation, so band-power maps are computed from 53 of the 60 simulated seconds.

Initial conditions set cortical and hippocampal modal field coefficients to one, with auxiliary state variables initialized to zero. At each step, the solver reconstructs vertex-space fields, applies the sigmoid nonlinearity, injects delayed cross-region drive when coupling is nonzero, computes deterministic derivatives, and advances the state with stochastic forcing. At epoch boundaries, a projection-error diagnostic compares vertex-space firing rates with their reconstruction from the retained eigenmodes.

The final NPZ outputs contain cortical and hippocampal modal time series, projection-error traces, time step, output stride, coupling values, gain traces, and run diagnostics. Cortical theta-power maps were computed from the post-burn-in cortical modal trajectory by reconstructing cortical activity and estimating band-limited power in the 4–8 Hz band. Alpha-band (8–12 Hz) maps were retained for exploratory spectral-regime checks. A separate modal spectral sweep over the archived trajectories reproduced the qualitative cortical-alpha and hippocampal-theta regime reported by [30] (Supplementary Methods, “EC-reflow sweep spectral-regime validation”).

4.6 Coupling architectures

We evaluated several coupling operators while keeping the dynamical equations, surface basis, and empirical validation pipeline fixed.

AP-topographic K-nearest-neighbor coupling. The first coupled model used AP-topographic hippocampus-cortex coupling. Sparse operators $W_{C \leftarrow H} \in \mathbb{R}^{N_C \times N_H}$ and $W_{H \leftarrow C} \in \mathbb{R}^{N_H \times N_C}$ were constructed with distance-decaying K-nearest-neighbor weights. The target map was organized so that hippocampal AP position corresponded to cortical AP position. This first model was useful because it produced a posterior cortical theta gradient, but it failed detailed electrode-level and pairwise target-coupling validation.

Shuffle controls. To test whether the AP-topographic model depended on structured hippocampal-cortical mapping rather than row-sum magnitude alone, shuffled controls permuted hippocampal columns while preserving row sums. Random and AP-bin shuffles destroyed or reversed the posterior theta pattern, showing that the first model’s AP effect was architecture-dependent.

Anatomical direct and relay models. Anatomical coupling candidates used group structural-connectivity information derived from siibra/Julich labels in HCP subjects 000–019. The direct candidate modeled hippocampus-to-cortex projection without an intermediate relay. The EC-reflow

candidate modeled hippocampal drive passing through entorhinal relay labels before reaching cortex; parahippocampal cortex is among the retained cortical target regions rather than part of the relay stage. Vertex-level operators were generated by assigning hippocampal and cortical vertices to their nearest source and target region centres and then applying sparse target weighting. Two properties of the accepted relay operator follow from its construction. The relay prior is formed as the product of total hippocampal-to-entorhinal weight and total entorhinal-to-target weight, so it is rank one over source and target regions: every hippocampal source carries the same cortical target profile up to a scalar. The operator therefore encodes route selection rather than source-specific target selectivity. Each cortical vertex draws from at most 96 hippocampal vertices, and cortical rows are zeroed either when the nearest retained target centre lies beyond 30 mm or when the relay prior for the target region is empty; in the accepted bundle these rules leave 101,782 of 163,842 cortical vertices receiving hippocampal drive. The coupled model is therefore a spatially restricted perturbation of the uncoupled substrate.

Reference-rowwise scaling. Because the candidate operators differ in target structure and row-wise coupling magnitude, we separated “where the coupling goes” from “how much drive each row receives.” Reference-rowwise scaling matched each anatomical candidate’s rowwise magnitude distribution to the original K-nearest-neighbor reference while preserving its target structure. The suffix “refrow” denotes this operation. The main state-gating model is the entorhinal relay with reference-rowwise scaling (EC-refrow).

Coupling sweep. The exploratory EC-refrow sweep simulated ten coupling values,

$$cc \in \{0, 0.002, 0.005, 0.008, 0.010, 0.015, 0.020, 0.030, 0.050, 0.080\}, \quad ch = 0.2cc.$$

The same ten simulations were sampled at electrode coordinates in all three empirical cohorts (ds004551, Penn SIEVE, and the navigation datasets), so model dynamics are identical across states and only the empirical comparison set differs. The full ten-point grid, including $cc = 0.080$, contributed to the descriptive per-subject slope analysis in Fig. 1a,c. The $cc = 0.080$ trajectory exceeded the cortical divergence threshold ($\phi_{\text{RMS}} > 1000$) and was excluded from the separate spectral-regime diagnostic, but not from the full-grid slope. Fixed operating-point checks compared $cc = 0.005$ and $cc = 0.030$ against the uncoupled model; $cc = 0.080$ was not a confirmatory endpoint contrast.

For the exploratory right-hemisphere benchmark, the same ten-point coupling grid, simulator parameters, random seed, and stability criterion were retained while substituting the homologous right cortical, hippocampal, mass, eigenmode, and coupling assets; construction and validation details are reported in Supplementary Methods, “Exploratory right-hemisphere lateralization sensitivity analysis”.

Spectral-regime diagnostic. For each archived EC-refrow trajectory, non-constant modal time series were linearly detrended, Hann-tapered, Fourier transformed, and averaged to obtain cortical and hippocampal modal power spectra at the saved 1000-Hz sampling rate. Band summaries used 4–8 Hz for theta and 8–12 Hz for alpha. Across the coupled mid-range, both the left- and right-hemisphere realizations showed prominent cortical alpha-band and hippocampal theta-band structure, the joint regime empirically characterized in human recordings [44, 45, 29]. This reproduces the qualitative cortico-hippocampal spectral regime reported by [30] on the substituted BigBrain geometry; full spectra and the hemisphere comparison are reported in Supplementary Methods, “EC-refrow sweep spectral-regime validation” and “Exploratory right-hemisphere lateralization sensitivity analysis”.

4.7 Sleep iEEG validation

The sleep validation used the ds004551 iEEG dataset. Model theta maps were sampled at empirical electrode contact coordinates using nearest-surface assignment to the BigBrain cortical mesh. Contacts farther than 15 mm from the sampled model surface were excluded. The 15 mm threshold was a geometric quality-control rule fixed before empirical alignment. For BigBrain, all 8,985 coordinate rows fell within 15 mm (median distance 1.39 mm); 8,948 of these had matching empirical bandpower and entered the frozen statistics. Subject-level analyses included 114 subjects. We reran retention and pooled/subject-level validation at 5, 10, 15, 20, 25, and 30 mm. For each retained contact-session row, empirical theta power was compared with sampled model theta power. The primary pooled statistics were Spearman correlation and AP-partial Spearman correlation controlling MNI-Y position. Subject-level analyses computed correlations within each subject and then used paired tests across subjects for model comparisons. For coupled-model comparisons, the subject-level quantity was the EC-refrow minus comparator difference in raw or AP-partial correlation.

Spatial nulls were used to test whether BigBrain alignment could be explained by generic spatial autocorrelation or by AP-aligned but otherwise unstructured maps. The AP-bin null permuted model values within AP quantile bins, preserving large-scale anterior-posterior structure while destroying finer spatial organization. The global-shuffle null permuted model values without preserving AP structure. For the exploratory EC-refrow sleep coupling sweep, pointwise tests were corrected within three pre-defined nine-test families: pooled raw Spearman alignment, pooled AP-partial Spearman alignment, and subject-level raw alignment against zero. We report both Holm family-wise-error and Benjamini–Hochberg false-discovery-rate corrections. A within-subject omnibus statistic was computed by summing squared standardized subject-level deviations across coupling strengths and comparing the observed value with 10,000 label permutations.

4.7.1 Penn SIEVE adult sleep iEEG validation

Penn SIEVE [36] was used as an independent adult sleep validation cohort, analysed across the full EC-refrow coupling sweep. We analyzed processed sleep-stage EDF clips with released MNI channel coordinates. Empirical theta was recomputed from the EDF clips as theta power divided by total 1–80 Hz power, using 2 s Welch windows with 1 s stride. Model maps were sampled at contact coordinates using the same `both_mirrored` convention as the navigation analyses, $x_{\text{model}} = -|x|$, and the same 15 mm nearest-surface gate.

The primary Penn SIEVE sleep analysis used N3 non-rapid-eye-movement sleep, the deep slow-wave sleep stage most comparable to the ds004551 slow-wave sleep anchor. We report all contacts and a clinically normal-like stratum. The clinically normal-like stratum excluded any bipolar channel with an endpoint in the seizure-onset zone, in resected tissue, or with an endpoint spike rate greater than 1/hour. In N3, the all-contact BigBrain analysis retained 5,388 contacts, and the all-contact reference-surface analysis retained 5,361 contacts. In the clinically normal-like stratum, BigBrain retained 1,880 contacts from 78 subjects, and the reference surface retained 1,875 contacts under the same distance gate. The complete three-seed paired endpoint analysis included 70 subjects.

For Penn SIEVE, the primary statistic was the per-subject slope of AP-partial Spearman ρ against cc across the full EC-refrow sweep, matching the ds004551 and navigation analyses. Pooled contact-level Spearman and AP-partial Spearman correlations are reported as descriptive alignment statistics. Spatial inference used an AP-bin-preserving null that permuted model ranks within 10 equal-count MNI-Y bins. Directional one-sided tests were used for the predicted uncoupled BigBrain positive alignment and EC-coupled negative sleep shift; the reference-surface test was two-sided. Fixed operating-point checks compared $cc = 0.030$ against $cc = 0$ on the same contact set, with

both AP-bin pooled nulls and subject-level signed-rank/sign tests reported. The seeded Penn SIEVE endpoint check repeated the clinically normal-like N3 analysis at $cc = 0$ and $cc = 0.030$ for seeds 73924, 104729, and 221017.

4.8 Navigation iEEG validation

Navigation-memory validation used OpenNeuro datasets ds005522 and ds005523. The analysis was specified before empirical model alignment was computed. The frozen choices were the EC-refrow coupling grid, the 15 mm nearest-surface assignment rule, bipolar acquisition, the `both_mirrored` hemisphere route, the per-subject coupling-sensitivity statistic, minimum subject gates, and confirmatory fixed-coupling contrasts.

Pre-extraction audit. Before reading empirical EDF signals, the pipeline audited BIDS sidecars only: bipolar channel files, electrode-coordinate files, and event metadata. Contacts were transformed into the common left-hemisphere model space using `both_mirrored`, which maps all contacts to $x_{\text{model}} = -|x|$. Each contact was assigned to the nearest BigBrain cortical vertex and excluded if the distance exceeded 15 mm. Subjects were required to have at least 30 retained contacts and at least 60 mm anterior-posterior span. The minimum dataset feasibility gate was 8 subjects. Both datasets passed this gate before empirical alignment was computed: ds005522 had 49 subjects passing contact and AP gates, and ds005523 had 19.

Bandpower extraction. After gates were frozen, DataLad payloads were materialized and EDF signals were read with MNE-Python. The event labels defining navigation phases were: `NAV_LEARN` and `TRAVEL_LEARN` for encoding, `NAV_TEST` and `TRAVEL_TEST` for retrieval, and all four labels for `all_task`. ds005522 was analyzed in all-task, encoding, and retrieval phases. ds005523 was analyzed in retrieval as the confirmatory replication phase. Signals were resampled to 200 Hz when their native sampling rate exceeded that value. Phase windows were concatenated within run; phase-band cells with less than 10 s of total data were skipped. Power spectra were estimated with Welch windows of 4 s and 50% overlap. Bandpower was integrated with the trapezoid rule over theta (4–8 Hz) and alpha (8–12 Hz). Log-transformed power was z-scored within subject, session, acquisition, phase, and band before model comparison.

Primary state-by-architecture statistic. For each subject, phase, band, and coupling value, the sampled model map was correlated with empirical log power. The primary statistic was the per-subject slope of AP-partial Spearman ρ as a function of cc across the full EC-refrow sweep. The confirmatory state contrast tested whether navigation slopes exceeded the sleep slope anchor from ds004551. Fixed operating-point checks compared $cc = 0.005$ and $cc = 0.030$ against $cc = 0$. The analysis thereby tests whether the same EC-refrow coupling axis that harms sleep theta alignment becomes favorable during navigation-memory states.

Robustness analyses. Coverage matching compared sleep and navigation subjects after accounting for log contact count and AP span. Two checks were used: common-support filtering and greedy nearest-neighbor matching without replacement. A navigation spatial-null test shuffled model theta values within AP bins before recomputing subject-level partial correlations, preserving AP structure while destroying finer spatial organization. Resolution-variant analyses sampled uncoupled BigBrain maps at 30k, 60k, 100k, and 164k vertices. A fixed-intersection audit confirmed that all retained navigation contacts were present across all four variants in both datasets (9,457/9,457 in ds005522 and 6,085/6,085 in ds005523).

4.9 Reference-surface and within-BigBrain resolution controls

The reference-surface benchmark asked whether the uncoupled BigBrain result was a generic consequence of running an NFT model on any cortical surface. The lower-resolution reference surface and BigBrain surface were validated under the same ds004551 sleep pipeline. Because the two surfaces retain very different sets of electrodes under the 15 mm distance rule, the analysis distinguishes three support sets: the BigBrain-retained set, the reference-retained set, and the overlap of rows valid on both surfaces. Sampling the reference surface at the BigBrain-retained set is explicitly extrapolative, because most of those contacts lie beyond the reference surface’s 15 mm validity threshold.

Within-BigBrain downsampling tested whether the BigBrain advantage was merely a vertex-count effect. We compared approximately 30k, 60k, and 100k-vertex downsampled cortical meshes with the unchanged 164k mesh, preserving the BigBrain anatomical source and coordinate frame. Cortical bases and uncoupled NFT trajectories were generated for the resolution controls and evaluated with the same ds004551 validation pipeline. The archived 164k resolution-control and substrate-benchmark products share an uncoupled modal trajectory but retain distinct historical theta-map extractions: the resolution extraction supplies the resolution series, whereas the benchmark extraction supplies the substrate comparison. Navigation analyses used fixed contact intersections across variants.

4.10 Wilson-Cowan model-class baseline

As a model-class control, we simulated a stochastic Wilson-Cowan NMM in The Virtual Brain (TVB) on the BigBrain 30k surface. The tested configuration used global coupling 0.014, additive noise amplitude 10^{-5} , a 60 s simulation, and a 2 s burn-in. The resulting theta map was non-degenerate, with theta-power standard deviation 1.2×10^{-5} and 4,113 unique sampled values. It was passed through the same frozen sleep and navigation validation pipelines on matched contact sets. This analysis is a post hoc model-class robustness check: it shows that the tested Wilson-Cowan configuration is near-null under these validations, but it does not claim that all possible NMM parameterizations must fail.

4.11 Task-fMRI reconstruction

Task-fMRI reconstruction used 47 HCP task contrasts. We reproduced the Pang-style fsLR-32k geometric reconstruction baseline [9], warped HCP task maps to BigBrain space using [17], computed BigBrain geometric eigenmodes, and computed BigBrain EDR-connectome modes. The main design crossed data resolution and basis type:

A :	fsLR data	fsLR geometric modes,
B :	fsLR data	fsLR EDR modes,
C :	BigBrain data	BigBrain geometric modes,
D :	BigBrain data	BigBrain EDR modes.

BigBrain geometric reconstruction used mass-weighted projection

$$a_N = \Psi_{1:N}^\top M y, \quad \hat{y}_N = \Psi_{1:N} a_N. \quad (12)$$

The EDR modes are the eigenvectors of a graph Laplacian carrying no surface-area weighting, so their mass matrix is the identity and the corresponding reduced inverse is $a_N = \Psi_{1:N}^\top y$. The primary comparison therefore projects each basis under its native metric, which is how each is conventionally

constructed and used. The fsLR conditions use unit-weighted least squares, following the reference implementation [9]; fsLR-32k vertex areas are close to uniform, whereas BigBrain vertex areas span roughly 0.1–2.3 mm², which is why the BigBrain geometric cell uses the mass-weighted projector rather than least squares. As a control, the comparison was repeated with the EDR basis projected under the BigBrain lumped cortical mass metric using the full weighted projector

$$a_N = \left(\Psi_{1:N}^\top M \Psi_{1:N} \right)^{-1} \Psi_{1:N}^\top M y, \quad (13)$$

with the geometric cell unchanged. Because the geometric coefficients are already computed under that metric, the geometric reconstruction is identical between the two runs and any change in the contrast is attributable to the connectome projector alone. Reconstruction accuracy was measured by Pearson correlation between empirical and reconstructed task maps over valid cortical vertices. We report two registration-related quantities. The synthetic round-trip scalar ($r \approx 0.986$) summarizes overall BigBrainWarp surface-resampling fidelity. The fMRI reconstruction benchmark is task-specific: for each HCP contrast, BigBrain geometric reconstruction at $N = 200$ modes was compared with the per-contrast round-trip-fidelity stored values (median $r = 0.928$, range 0.802–0.948). The central scale-reversal statistic was Cell C minus Cell D (BigBrain geometry minus BigBrain EDR), reported as the median across the 47 task contrasts. Paired tests across contrasts were two-sided Wilcoxon signed-rank tests, Holm-corrected across the six evaluated mode counts.

4.12 Spectral-retuning control

Coupling shifts the model’s cortical spectrum [30], so the sleep-side decrease has a competing explanation: coupling could move power out of the 4–8 Hz window without changing the spatial map at all. If that were true, retuning the cortex to hold its dominant peak in theta would restore sleep alignment, and the state-gating result would reduce to a choice of measurement band. We tested this directly.

Parameter selection used simulated spectra, finiteness, RMS, and saturation diagnostics only; empirical iEEG alignment was never read. The selector writes a confirmation manifest only when all five coupling values ($cc \in \{0, 0.005, 0.010, 0.020, 0.030\}$) carry a stable theta candidate, so empirical validation cannot be run on a partial grid.

Three parameter families were evaluated across 223 simulations. Uniform scaling of cortical synaptic rates used a three-stage search: a 50-run pilot over the ten coupling values by five rate scales, a 42-run adaptive refinement of the observed subtheta-to-alpha transition, and a 36-run ultra-fine search inside the four remaining brackets at step sizes down to 0.001 in rate scale (128 runs, 123 stable, six theta-eligible). Two further families were tested independently and never combined into a post hoc multi-parameter search: scaling of the cortical beta/alpha-ratio with alpha and all gains held at their submitted values (25 runs, all stable), and signed scaling of cortical recurrent gain with its ratio and all synaptic time constants fixed (25 runs plus 45 refinements of every observed alpha-to-subtheta bracket, including both nonmonotonic brackets, all stable). A run was theta-eligible only if it remained stable and its dominant cortical spectral peak fell in the theta band.

The retuning analysis identified a consistent constraint on the cortical spectral regime. Under uniform synaptic-rate scaling, the dominant cortical peak transitioned directly between subtheta and alpha branches at every weak-coupling value, including at the finest parameter-search resolution tested. Across all three parameter families, no theta-eligible candidate was found at $cc \in \{0, 0.005, 0.010, 0.020\}$, preventing construction of a complete theta-preserving grid spanning $cc \in \{0, 0.005, 0.010, 0.020, 0.030\}$. These results establish that, within the evaluated ranges, maintaining a dominant cortical theta peak across coupling strengths requires changes beyond the tested one-parameter adjustments, motivating investigation of coordinated parameter changes.

4.13 Statistical reporting and reproducibility

Pooled electrode-level correlations are reported as descriptive alignment statistics; subject-level paired tests are used for inference wherever the comparison is between models within the same empirical cohort. One-sided tests were used only when a directional contrast was pre-specified, such as navigation slope greater than sleep slope. Spatial-null tests used permutation distributions with finite-sample floors reported explicitly.

The reproduction archive (described in the Code availability section) invokes per-figure analysis and renderer scripts through `run_all_figures.sh`. It reprocesses the required empirical recordings and rebuilds the reported statistics and figures from checksum-verified included simulation trajectories, modal bases, and other specified inputs. Regenerating the neural-field trajectories and BigBrain EDR basis requires HPC. Each figure writes values-used and verification reports; numerical acceptance uses declared tolerances rather than requiring historical byte identity. The right-hemisphere and spectral-regime analyses are separate optional modules, outside the four-figure principal run. Access information is provided in the Data availability and Code availability sections. The spectral-retuning control is an auxiliary analysis outside the principal four-figure workflow. The *theta_preservation_results/* directory contains per-simulation job manifests for the 128 uniform synaptic-rate simulations.

5 CONCLUSION

This paper asked whether anatomical constraints on cortical dynamics are fixed, or instead depend on brain state and spatial scale. Three questions followed (Sec. 1):

1. Whether a fixed anatomical route makes a fixed functional contribution: the EC-refrow relay reverses sign between sleep and navigation (Methods, Sec. 4.6; Sec. 2.1).
2. Whether high-resolution histological geometry is an empirically supported substrate rather than an implementation detail: the uncoupled surface aligns with sleep theta through expanded valid coverage rather than superior local prediction (Methods, Sec. 4.9; Sec. 2.3), and that alignment depends on the neural-field model class rather than on the surface alone, since an NMM on identical geometry is near-null in both states (Methods, Sec. 4.10; Sec. 2.4).
3. Whether geometry or connectivity is the dominant constraint: neither is uniformly dominant, and each provides the better account over a different range of spatial scales (Methods, Sec. 4.11; Sec. 2.5).

Two categorical framings inherited from the literature therefore hold only conditionally: that an anatomical route makes a fixed functional contribution across brain states, as is implicit in treating structural connectivity as a state-invariant constraint [22, 14, 15], and that either geometry [9] or connectivity [11] is uniformly dominant. Instead, the results support a layered account in which the functional contribution of anatomical coupling depends on brain state, while the relative contribution of geometry and connectivity depends on spatial scale, consistent with the graded picture reported subsequently [12].

A drive layer, hippocampal-cortical coupling, is state-gated (Sec. 2.1). On the full-resolution BigBrain substrate, increasing EC-refrow coupling decreases theta alignment during sleep but increases it during navigation. Coupling is therefore neither globally beneficial nor globally harmful: its effect on empirical theta alignment changes sign with brain state.

A substrate layer, high-resolution histological geometry, defines the field basis and controls electrode-to-surface correspondence (Sec. 2.3). The uncoupled BigBrain geometry is a genuine sleep-theta substrate; its advantage lies in expanded geometrically valid coverage rather than superior local prediction, and the alignment depends on the neural-field model class rather than on the surface alone (Sec. 2.4), since the tested conventional surface-based neural-mass baseline [16] on identical BigBrain geometry is near-null in both states.

A scale layer connects these state-dependent effects to spatial representation (Sec. 2.5). Within the uncoupled BigBrain family, sleep favours intermediate-to-fine resolution whereas navigation favours the coarse 30k surface. At full resolution, EC-reflow coupling produces positive navigation alignment, while an uncoupled coarse 30k surface does so independently. The same broader dependence on spatial scale appears independently in task fMRI: connectivity better reconstructs the longest-wavelength patterns, whereas high-resolution geometry better reconstructs finer ones, with a crossover near ≈ 93 mm. We interpret this as a shared scale dependence rather than as evidence that coupling and coarsening are mechanistically identical.

Across the analyses, anatomy does not impose a single fixed constraint. Coupling changes its functional contribution with brain state, while the relative contribution of geometry and connectivity changes with spatial scale. What determines alignment is therefore not anatomy alone, but anatomy read out at a state-dependent spatial scale.

Acknowledgements

A.S.-M. gratefully acknowledges the Secretaría de Ciencia, Humanidades, Tecnología e Innovación (SECIHTI, formerly CONACYT/CONAHCYT), Mexico, for supporting his doctoral studies. This research was enabled in part by support provided by Calcul Québec (<https://www.calculquebec.ca>) and the Digital Research Alliance of Canada (<https://alliancecan.ca>).

Author contributions

A. Salinas-Medina originated the research ideas and formulated the research questions and hypotheses; designed and implemented the neural-field models, coupling operators, and validation pipeline; executed all simulations and analyses; interpreted the results; produced the figures; and wrote the manuscript. A. González-Mitjans contributed through scientific discussions and responses to scientific and methodological questions arising during the research. P.-J. Toussaint contributed through scientific discussions and responses to questions concerning the neuroimaging and neuroanatomical aspects of the research. X. Liu provided academic supervision and guidance on the research scope, methodological development, and scientific presentation. A. Evans provided academic supervision and guidance on the neuroscientific context, interpretation, and scientific presentation.

Competing interests

The authors declare no competing interests.

Funding

A.S.-M. discloses support for his doctoral studies, during which this work was carried out, from the Secretaría de Ciencia, Humanidades, Tecnología e Innovación (SECIHTI, formerly CONAHCYT),

Mexico [scholarship number 834021]. A.G.-M., P.-J.T., X.L., and A.E. declare no relevant funding.

Data availability

The study used OpenNeuro datasets ds004551 (version 1.0.6; sleep iEEG) [34], ds005522 (version 1.0.0) [32], and ds005523 (version 1.0.1; navigation-memory iEEG) [33]. Penn SIEVE dataset 414 was used for adult sleep/wake iEEG validation [36]. The BigBrain [1], BigBrainWarp [17], HippUnfold [42], and HCP task-fMRI resources [37, 38, 39] are described in Methods (Secs. 4.2 and 4.11).

Raw OpenNeuro and Penn SIEVE recordings are not redistributed in the reproduction archive. The archive documentation provides the pinned OpenNeuro versions, retrieval commands, Penn SIEVE access instructions, and expected directory layout.

The derived inputs used for the reported analyses, including checksum-pinned neural-field trajectories, modal bases, and other specified HPC products, are deposited in the accompanying Zenodo reproduction archive (record 22880060; reserved DOI: 10.5281/zenodo.22880060). During peer review, the private draft is available through this private Zenodo reviewer-access link. The reserved DOI will become resolvable when the Zenodo record is published.

These checksum-verified archived products define the analysis inputs used for reproduction. Regenerating the neural-field trajectories and BigBrain EDR basis requires suitable high-performance computing resources and may not produce bitwise-identical outputs across numerical solvers, hardware, and software environments.

Code availability

The analysis, simulation, figure-generation, and verification code is included in the same Zenodo reproduction archive (record 22880060; reserved DOI: 10.5281/zenodo.22880060). During peer review, access is provided through the private Zenodo reviewer-access link given in the Data availability section.

The principal reproduction workflow is invoked through `run_all_figures.sh`. It rebuilds the four main figures and reported statistics from the required empirical recordings and checksum-verified derived HPC products. The archive also provides workflows for the three-seed endpoint analysis, the spectral-regime diagnostic, and the optional right-hemisphere sensitivity analysis.

Verification uses declared numerical tolerances rather than requiring historical byte identity. Console transcripts and timestamped logs document the principal and optional runs. The spectral-retuning control is an auxiliary analysis outside the principal four-figure workflow; the scope of its archived provenance is stated in Methods (Sec. 4.13).

Human research ethics

This study comprised secondary analyses of previously collected human neuroimaging and intracranial electrophysiology data. The authors did not recruit participants, perform interventions, influence clinical care, collect new biospecimens, or have access to direct participant identifiers. Under Article 2.4 of the Canadian Tri-Council Policy Statement (TCPS2), research relying exclusively on the secondary use of anonymous information does not require research ethics board review. All data were used in accordance with the terms of their respective data-sharing agreements.

For OpenNeuro dataset ds004551, the original study was approved by the Institutional Review Board of Wayne State University. Written informed consent was obtained from participants' legal

guardians, and written assent was obtained from participants aged 13 years or older. For Open-Neuro datasets ds005522 and ds005523, participants provided written informed consent, and the multisite studies were approved by the institutional review boards at the participating institutions, the University of Pennsylvania Institutional Review Board as the data-coordinating site, and the Human Research Protection Official at the Space and Naval Warfare Systems Center Pacific.

For Penn SIEVE dataset 414, the source study received ethical approval from the Institutional Review Board of the University of Pennsylvania, and the source investigators confirmed that the required participant consent had been obtained and archived. The WU–Minn Human Connectome Project protocol was approved by the Institutional Review Board of Washington University in St. Louis (IRB no. 201204036), and all participants provided written informed consent.

The BigBrain post-mortem tissue donor provided written informed consent for the use of the donated tissue for research and education. Its use was approved by the Ethics Committee of the Faculty of Medicine at Heinrich Heine University Düsseldorf (approval no. 4863). HippUnfold-derived surfaces and the remaining anatomical resources were used as previously published digital reference materials and did not involve new participant recruitment or tissue collection.

References

- [1] Amunts, K., et al. BigBrain: an ultrahigh-resolution 3D human brain model. *Science* 340, 1472–1475 (2013).
- [2] Wagstyl, K., et al. BigBrain 3D atlas of cortical layers: cortical and laminar thickness gradients diverge in sensory and motor cortices. *PLoS Biol.* 18, e3000678 (2020).
- [3] Demirtaş, M., et al. Hierarchical heterogeneity across human cortex shapes large-scale neural dynamics. *Neuron* 101, 1181–1194 (2019).
- [4] Wang, X.-J. Macroscopic gradients of synaptic excitation and inhibition in the neocortex. *Nat. Rev. Neurosci.* 21, 169–178 (2020).
- [5] Douglas, R. J. and Martin, K. A. C. Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.* 27, 419–451 (2004).
- [6] Zilles, K. and Amunts, K. Centenary of Brodmann’s map — conception and fate. *Nat. Rev. Neurosci.* 11, 139–145 (2010).
- [7] García-Cabezas, M. Á., Zikopoulos, B., and Barbas, H. The Structural Model: a theory linking connections, plasticity, pathology, development and evolution of the cerebral cortex. *Brain Struct. Funct.* 224, 985–1008 (2019).
- [8] Van Essen, D. C. A tension-based theory of morphogenesis and compact wiring in the central nervous system. *Nature* 385, 313–318 (1997).
- [9] Pang, J. C., et al. Geometric constraints on human brain function. *Nature* 618, 566–574 (2023).
- [10] Gabay, N. C. and Robinson, P. A. Cortical geometry as a determinant of brain activity eigenmodes. *Phys. Rev. E* 96, 032413 (2017).
- [11] Vohryzek, J., Sanz-Perl, Y., Kringelbach, M. L., and Deco, G. Human brain dynamics are shaped by rare long-range connections over and above cortical geometry. *Proceedings of the National Academy of Sciences of the United States of America*, 122(1):e2415102122, 2025. doi:10.1073/pnas.2415102122.

- [12] Royer, J., Paquola, C., Rodriguez-Cruces, R., Auer, H., Ngo, A., Sahlas, E., Arafat, T., Abdallah, C., Mansilla de la Torre, D., Pana, R., Hall, J., Leech, R., Smallwood, J., Frauscher, B., and Bernhardt, B. C. Human cortical dynamics reflect graded contributions of local geometry and network topography. *Nature Communications*, 17:1002, 2026. doi:10.1038/s41467-025-67740-2.
- [13] Bullmore, E. and Sporns, O. Complex brain networks. *Nat. Rev. Neurosci.* 10, 186–198 (2009).
- [14] Sporns, O. Contributions and challenges for network models in cognitive neuroscience. *Nat. Neurosci.* 17, 652–660 (2014).
- [15] Fornito, A., Zalesky, A., and Bullmore, E. *Fundamentals of Brain Network Analysis*. Academic Press (2016).
- [16] P. Sanz Leon et al., “The Virtual Brain: a simulator of primate brain network dynamics,” *Front. Neuroinform.*, vol. 7, p. 10, 2013.
- [17] Paquola, C., et al. The BigBrainWarp toolbox. *eLife* 10, e70119 (2021).
- [18] DeKraker, J., Cabalo, D. G., Royer, J., Khan, A. R., Karat, B., Benkarim, O., Rodriguez-Cruces, R., Frauscher, B., Pana, R., Hansen, J. Y., and Misic, B. HippoMaps: Multiscale cartography of human hippocampal organization. *bioRxiv*, 2024. doi:10.1101/2024.02.23.581734. Preprint.
- [19] Wilson, H. R. and Cowan, J. D. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.* 12, 1–24 (1972).
- [20] Deco, G., et al. The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Comput. Biol.* 4, e1000092 (2008).
- [21] Robinson, P. A., et al. Eigenmodes of brain activity: neural field theory predictions and comparison with experiment. *NeuroImage* 142, 79–98 (2016).
- [22] Honey, C. J., et al. Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci. USA* 106, 2035–2040 (2009).
- [23] Atasoy, S., Donnelly, I., and Pearson, J. Human brain networks function in connectome-specific harmonic waves. *Nat. Commun.* 7, 10340 (2016).
- [24] Aghajan, Z. M. et al. Theta oscillations in the human medial temporal lobe during real-world ambulatory movement. *Curr. Biol.* **27**, 3743–3751.e3 (2017). doi:10.1016/j.cub.2017.10.062.
- [25] Stangl, M. et al. Boundary-anchored neural mechanisms of location-encoding for self and others. *Nature* **589**, 420–425 (2021). doi:10.1038/s41586-020-03073-y.
- [26] Maoz, S. L. L. et al. Dynamic neural representations of memory and space during human ambulatory navigation. *Nat. Commun.* **14**, 6643 (2023). doi:10.1038/s41467-023-42231-4.
- [27] Seeber, M. et al. Human neural dynamics of real-world and imagined navigation. *Nat. Hum. Behav.* **9**, 781–793 (2025). doi:10.1038/s41562-025-02119-3.
- [28] Stangl, M., Maoz, S. L. & Suthana, N. Mobile cognition: imaging the human brain in the “real world”. *Nat. Rev. Neurosci.* **24**, 347–362 (2023). doi:10.1038/s41583-023-00692-y.

- [29] Li, J., Cao, D., Yu, S., Wang, H., Imbach, L., Stieglitz, L., Sarnthein, J. & Jiang, T. Theta–alpha connectivity in the hippocampal–entorhinal circuit predicts working memory load. *J. Neurosci.* **44**, e0398232023 (2024). doi:10.1523/JNEUROSCI.0398-23.2023.
- [30] Phogat, R., Behler, A., Sonkusare, S., Pang, J. C., Koussis, N., Roberts, J. A., DeKraker, J., Shine, J. M., Fornito, A., Robinson, P. A., and Breakspear, M. A unified model of cortico-hippocampal interactions through neural field theory. *bioRxiv*, 2025. doi:10.1101/2025.09.07.674721. Preprint.
- [31] Cantero, J. L., Atienza, M., Stickgold, R., Kahana, M. J., Madsen, J. R., and Kocsis, B. Sleep-dependent theta oscillations in the human hippocampus and neocortex. *J. Neurosci.* **23**, 10897–10903 (2003). doi:10.1523/JNEUROSCI.23-34-10897.2003.
- [32] Herrema, H. G. and Kahana, M. J. Spatial Navigation Memory of Object Locations. *OpenNeuro*, dataset ds005522, version 1.0.0, 2024. doi:10.18112/openneuro.ds005522.v1.0.0.
- [33] Herrema, H. G. and Kahana, M. J. Spatial Memory of Object Locations with Open-Loop Stimulation at Encoding. *OpenNeuro*, dataset ds005523, version 1.0.1, 2024. doi:10.18112/openneuro.ds005523.v1.0.1.
- [34] Sakakura, K., Kuroda, N., Sonoda, M., Mitsuhashi, T., Firestone, E., Luat, A. F., Marupudi, N. I., Sood, S., and Asano, E. iEEG on children during slow wave sleep. *OpenNeuro*, dataset ds004551, version 1.0.6, 2023. doi:10.18112/openneuro.ds004551.v1.0.6.
- [35] Sakakura, K., Kuroda, N., Sonoda, M., Mitsuhashi, T., Firestone, E., Luat, A. F., Marupudi, N. I., Sood, S., and Asano, E. Developmental atlas of phase-amplitude coupling between physiologic high-frequency oscillations and slow waves. *Nature Communications*, **14**:6435, 2023. doi:10.1038/s41467-023-42091-y.
- [36] Pattnaik, A. R., Ong, I. Z., Bernabei, J. M., et al. Normative intracranial EEG highlights epileptic abnormalities across wakefulness and sleep. *medRxiv* (2025). doi:10.1101/2025.09.10.25335533. Preprint.
- [37] Van Essen, D. C., et al. The WU-Minn Human Connectome Project. *NeuroImage* **80**, 62–79 (2013).
- [38] Barch, D. M., et al. Function in the human connectome: task-fMRI. *NeuroImage* **80**, 169–189 (2013).
- [39] Glasser, M. F., Smith, S. M., Marcus, D. S., Andersson, J. L. R., Auerbach, E. J., Behrens, T. E. J., Coalson, T. S., Harms, M. P., Jenkinson, M., Moeller, S., Robinson, E. C., Sotiropoulos, S. N., Xu, J., Yacoub, E., Ugurbil, K., and Van Essen, D. C. The Human Connectome Project’s neuroimaging approach. *Nature Neuroscience*, **19**(9):1175–1187, 2016. doi:10.1038/nn.4361.
- [40] Bansal, K., Medaglia, J. D., Bassett, D. S., Vettel, J. M. & Muldoon, S. F. Data-driven brain network models differentiate variability across language tasks. *PLoS Comput. Biol.* **14**, e1006487 (2018). doi:10.1371/journal.pcbi.1006487.
- [41] Van Essen, D. C., Donahue, C. J., Coalson, T. S., Kennedy, H., Hayashi, T., and Glasser, M. F. Cerebral cortical folding, parcellation, and connectivity in humans, nonhuman primates, and mice. *Proc. Natl. Acad. Sci. USA* **116**, 26173–26180 (2019). doi:10.1073/pnas.1902299116.

- 1139 [42] DeKraker, J., Haast, R. A. M., Yousif, M. D., Karat, B., Lau, J. C., Köhler, S. & Khan,
1140 A. R. Automated hippocampal unfolding for morphometry and subfield segmentation with
1141 HippUnfold. *eLife* **11**, e77945 (2022). doi:10.7554/eLife.77945.
- 1142 [43] Reuter, M., Wolter, F.-E., and Peinecke, N. Laplace–Beltrami spectra as ‘Shape-DNA’ of
1143 surfaces and solids. *Computer-Aided Design* **38**, 342–366 (2006). doi:10.1016/j.cad.2005.10.011.
- 1144 [44] Nunez, P. L. *Neocortical Dynamics and Human EEG Rhythms*. Oxford University Press, New
1145 York (1995).
- 1146 [45] Buzsáki, G. Theta oscillations in the hippocampus. *Neuron* **33**, 325–340 (2002).
1147 doi:10.1016/S0896-6273(02)00586-X.
- 1148 [46] Cofré, R. et al. Whole-brain models to explore altered states of consciousness from the bottom
1149 up. *Brain Sci.* **10**, 626 (2020). doi:10.3390/brainsci10090626.
- 1150 [47] Wan, B. et al. Microstructural asymmetry in the human cortex. *Nat. Commun.* **15**, 10124
1151 (2024). doi:10.1038/s41467-024-54243-9.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [supplementary.pdf](#)