

Supplementary Information

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

Alejandro Salinas-Medina, Anisleidy González-Mitjans,
Paule-Joanne Toussaint, Xue Liu and Alan Evans

Supplementary Methods

1 Experimental sequence and inferential scope

The project proceeded through four experimental phases. The first phase built a BigBrain cortico-hippocampal NFT model and tested whether an AP-topographic hippocampal-cortical drive could reproduce the broad posterior loading of empirical sleep theta. This phase established that the model could generate the correct AP sign but failed detailed map alignment and pairwise target-coupling validation.

The second phase asked whether the failure reflected the coupling architecture. It compared the original K-nearest-neighbor operator, shuffle controls, row-normalized variants, direct anatomical coupling, and the entorhinal relay. EC-refrow was the strongest among coupled models in subject-level ds004551 validation, but it did not beat the uncoupled BigBrain substrate in pooled sleep theta validation.

The third phase therefore isolated the substrate question: before coupling is added, does BigBrain geometry itself provide an empirically useful field basis? The uncoupled BigBrain benchmark produced positive sleep-theta alignment under the frozen validation pipeline, while the lower-resolution reference surface was wrong-sign or coverage-limited. Subject-paired comparisons, spatial nulls, density-matched BigBrain controls, and within-BigBrain downsampling all supported the practical substrate claim.

The fourth phase extended the scale question to task fMRI. HCP task maps were reconstructed in fsLR and BigBrain space using geometric and EDR-connectome modes. This phase showed that geometry and connectivity dominate different spatial wavelengths rather than forming a single fixed hierarchy.

The state-by-architecture navigation experiment was added after the sleep coupling sweep but before empirical navigation alignment. Its purpose was to test whether the EC-refrow coupling axis that harms intrinsic sleep theta alignment becomes favorable during navigation-memory states in which hippocampal drive should be behaviorally engaged.

2 Surface preparation and quality control

The cortical mesh was used as a connected left-hemisphere surface to preserve a single spatial domain for eigenmode computation. The hippocampal mesh required additional alignment because the BigBrain HippUnfold output is not natively in the same orientation as the cortical MNI surface. Rigid ICP was used rather than nonlinear warping so that the hippocampal surface metric was not deformed in a way that would confound the eigenmode calculation. The Taubin smoothing step was limited to five iterations and was verified by mean displacement and triangle quality metrics.

The important methodological point is that the BigBrain model is not simply a denser copy of the reference surface. It combines a single-donor histological cortex, a BigBrain-derived hippocampal surface aligned into the same MNI frame, and mass-weighted eigenmodes computed directly on these surfaces. This is why the reference-surface comparison cannot be interpreted as isolating vertex count alone; surface source, coordinate support, electrode retention, and local metric structure differ.

3 Eigenmode computation details

For each surface, the cotangent Laplacian was assembled from the triangular mesh. The lumped mass matrix encodes vertex-associated surface area and is needed both for physical eigenvalue units and for correct projection. Modes were sorted by ascending eigenvalue. The first mode is the constant mode and was retained in the basis but interpreted separately from nonconstant spatial patterns. The first nonconstant hippocampal mode followed the hippocampal AP axis, validating that the hippocampal surface basis captured the expected long axis after alignment.

For an eigenvalue λ_n on a two-dimensional surface, an approximate wavelength scale can be assigned by the usual inverse-square relationship between eigenvalue and spatial scale. This calibration was used in the fMRI analysis to label BigBrain mode counts by approximate wavelengths. The calibration is approximate because cortical surfaces are folded, finite, and nonuniform, but it gives the correct ordering from broad long-wavelength modes to finer short-wavelength modes.

4 Dynamic simulation implementation

The simulation state contains modal coefficients and their first-order partners for cortical and hippocampal wave fields, together with vertex-space auxiliary variables for synaptodendritic filters and firing-rate transforms. The integrator uses an epoch wrapper. At the first epoch, initial conditions are created from the requested mode counts and vertex counts. At later epochs, the final state from the previous epoch is copied into the delay buffer so that delayed local and cross-region terms can access past values.

At each time step the deterministic derivative is computed before the stochastic increment is applied. The deterministic step unpacks the state, retrieves delayed cortical and hippocampal variables, reconstructs vertex-space fields, applies sigmoid nonlinearities, injects cross-region coupling if *cc* or *ch* is nonzero, evaluates modal wave equations, and returns the full derivative vector. Euler-Maruyama then advances the state as

$$Y_{t+\Delta t} = Y_t + F(Y_t, t)\Delta t + G \Delta W_t,$$

where G is the diagonal stochastic gain vector. The stochastic gain vector is preassembled so that only the intended state blocks receive noise.

Projection loss is computed by projecting the final epoch firing-rate fields into the retained mode basis and reconstructing them back to vertex space. Large loss would indicate that the retained basis is not representing the nonlinear field well. This diagnostic is used for numerical fidelity, not for parameter learning.

Simulation state representation. At each integration time t , the simulator state $\mathbf{Y}(t)$ collects all dynamical variables required to advance the coupled cortical-hippocampal system. Schematically,

$$\mathbf{Y}(t) = \left[\zeta^c(t), \dot{\zeta}^c(t), \zeta^h(t), \dot{\zeta}^h(t), \mathbf{z}_{\text{aux}}^c(t), \mathbf{z}_{\text{aux}}^h(t) \right], \quad (\text{S1})$$

where ζ^c and ζ^h are the retained cortical and hippocampal modal coefficients, their dotted counterparts are the first-order partners required by the second-order wave equations, and $\mathbf{z}_{\text{aux}}^c$ and $\mathbf{z}_{\text{aux}}^h$ collect the vertex-space auxiliary variables required by the local synaptodendritic and firing-rate dynamics. With 300 retained modes in each anatomical system, the modal first-order component has dimension $2(300 + 300) = 1200$. Vertex-space fields are reconstructed from the modal coefficients only when nonlinear firing-rate, local synaptodendritic, or cross-region coupling terms must be evaluated, after which the resulting quantities are projected back into the retained modal basis. Quantities obtained algebraically from the state, such as firing rates produced by the sigmoid nonlinearity, need not be stored as independent modal state variables. Fixed meshes, eigenmodes, eigenvalues, projectors, coupling operators, and model parameters determine the evolution but are not themselves dynamical state variables.

5 AP-topographic coupling

The original coupling operator was designed to test a strong spatial hypothesis: hippocampal AP position should map onto cortical AP position. Each target vertex received sparse input from a fixed number of source vertices, with weights decreasing with distance in the selected coupling coordinate system. The hippocampus-to-cortex operator can be summarized by row sums

$$s_i = \sum_j W_{C \leftarrow H}(i, j),$$

and by the row-normalized hippocampal AP target

$$\bar{y}_H(i) = \frac{\sum_j W_{C \leftarrow H}(i, j) y_H(j)}{\sum_j W_{C \leftarrow H}(i, j)}.$$

The AP target relation was strong, and the resulting model produced a posterior cortical theta map. However, the same target structure did not predict empirical pairwise theta connectivity in the expected direction. This is the first major boundary on the mechanistic coupling claim.

6 Shuffle and row-structure controls

Shuffle controls tested whether the first AP-topographic model depended on structured hippocampal target organization. Column shuffling preserved the amount of drive each cortical row received while disrupting which hippocampal positions contributed to that drive. Random shuffling and AP-bin shuffling destroyed or reversed the posterior model theta pattern. This showed that the posterior model gradient was not a trivial consequence of row sums alone.

These controls do not validate the original coupling as biological anatomy. They only show that the simulated AP effect depends on structured coupling architecture. That distinction motivated the anatomical direct and EC-relay operators.

7 Anatomical direct and EC-refrow operators

Anatomical candidates were derived from group structural connectivity. The direct model allowed hippocampal labels to project directly to cortical targets. The EC-mediated model routed hippocampal drive through entorhinal relay labels before cortical projection; parahippocampal cortex is a retained cortical target rather than part of the relay stage. Both were converted to vertex-level sparse operators by assigning hippocampal and cortical vertices to source and target labels and then distributing target drive over eligible vertices.

The direct model produced a stable but weak AP effect. The EC-mediated relay produced the strongest subject-level improvement among the coupled models, consistent with the biological role of entorhinal cortex as a hippocampal-neocortical interface. However, the later coupling sweep showed that this is not a monotonic empirical improvement over the uncoupled substrate in sleep. EC-refrow is therefore a plausible route-selection model, not a completed validation of hippocampal-cortical connectivity.

8 Reference-rowwise scaling

Reference-rowwise scaling was introduced because raw anatomical operators differ in row sums as well as target locations. Without scaling, a model comparison would confound anatomical route with total drive magnitude. For each candidate, the rowwise magnitude distribution was matched to the original K-nearest-neighbor reference while preserving the candidate’s target pattern. In matrix terms, if s_i^{ref} is the reference row sum and s_i^{cand} is the candidate row sum, the scaled row is

$$W_{i\cdot}^{\text{scaled}} = \frac{s_i^{\text{ref}}}{s_i^{\text{cand}} + \epsilon} W_{i\cdot}^{\text{cand}}.$$

Rows with zero candidate support were left at zero. This procedure makes subject-level comparisons more interpretable because the models differ mainly in target structure rather than global rowwise input strength.

9 ds004551 validation details

The sleep iEEG validation samples a model map at empirical contact coordinates. The core row unit is a contact-session, because a subject may contribute contacts across sessions or runs. Coordinates are transformed into the model space, assigned to the nearest cortical surface vertex, and retained if the distance is no more than 15 mm. For the BigBrain full-resolution surface, retention was high and nearest-surface distance was small; for the lower-resolution reference surface, retention was much lower. This difference is itself part of the substrate result, because a surface that cannot validly sample empirical contacts cannot support the same detailed electrode-level validation.

The validation reports both pooled and subject-level quantities. Pooled correlations describe the global contact-session relationship but do not replace subject-level inference. Subject-level correlations are computed within each subject, and paired tests compare models across the same subjects. The AP-partial statistic is obtained by residualizing empirical and model ranks with respect to MNI-Y and correlating the residuals. This removes the broad AP gradient and asks whether finer spatial structure remains aligned.

10 Pairwise target-coupling validation

The target-coupling validation asked whether the model-defined hippocampal-cortical target structure predicted empirical theta connectivity between electrode pairs. If a coupling operator is a mechanistic map of hippocampal drive to cortex, then electrode pairs that are well matched by that operator should show stronger empirical theta relationships than mismatched pairs. This was tested using model-defined target mismatch measures and empirical pairwise theta connectivity.

The result did not support the expected direction for any coupled model tested. This is why the manuscript does not claim that EC-refrow closes the mechanistic hippocampal-cortical connectivity problem. The supported claim is narrower: among coupled models with matched dynamics and eigenbasis, EC-mediated relay is the strongest route tested, but pairwise target validation remains

unresolved. The rank-one structure of the relay prior bounds this directly: an operator whose source regions share a single target profile cannot express the source-specific pairwise structure that this test interrogates.

11 Reference-surface benchmark and coverage decomposition

The uncoupled BigBrain benchmark compared full-resolution BigBrain and the lower-resolution reference surface under the same ds004551 validation. The headline result is that BigBrain produces positive empirical sleep-theta alignment while the reference surface is wrong-sign or coverage-limited. This claim has three support regimes.

First, on each surface’s own valid rows, BigBrain retains far more contacts and is positive, whereas the reference surface retains far fewer rows and is negative. Second, when the reference surface is sampled at the identical BigBrain-retained rows, it remains sign-reversed; however, that comparison is extrapolative because most of those contacts are beyond the reference surface’s 15 mm validity threshold. Third, on rows valid for both surfaces, both maps are negative. Therefore the BigBrain advantage is not a claim of better local dynamics on mutually accessible contacts. It is a practical empirical-interface claim: BigBrain validly samples additional cortical regions where its neural-field theta map is positively aligned with sleep iEEG.

This coverage decomposition is essential. It prevents overclaiming the reference comparison as a pure surface-resolution experiment. The stronger causal control for vertex count is the within-BigBrain downsampling analysis, where surface source and coordinate frame are held fixed.

12 Within-BigBrain downsampling

Downsampled BigBrain variants were generated at approximately 30k, 60k, 100k, and 164k vertices. Each variant preserved the BigBrain anatomical source and MNI coordinate frame, and each received its own generalized Laplace-Beltrami eigenbasis and uncoupled simulation. Under ds004551 sleep validation, all four variants showed positive raw and AP-controlled alignment, with the strongest AP-controlled result at the intermediate 60k surface rather than the full 164k surface. The supported sleep claim is therefore robustness across the BigBrain surface family, not monotonic dependence on vertex count.

For navigation, the same variants were evaluated after the state-gating result. The navigation trend was strongest at the coarse 30k map and declined toward the full-resolution uncoupled map. Because contact retention can change with surface resolution, the fixed-intersection audit was decisive: all retained contacts were present under all four variants in both navigation datasets. Thus the navigation resolution trend is not a retained-contact artifact. The interpretation remains bounded: the analysis shows a BigBrain-family resolution dependence on a fixed contact set, not that vertex count alone is the mechanistic cause.

13 State-by-architecture preregistered navigation test

The navigation analysis was designed to test a state-dependent prediction after the sleep EC-reflow sweep. In sleep, increasing EC-reflow coupling decreases AP-partial theta-map alignment. If coupling is a state-gated drive rather than a globally good or bad anatomical addition, the same axis should become favorable when hippocampal drive is behaviorally engaged during navigation-memory tasks.

The analysis choices were recorded before empirical model alignment: model family EC-reflow; primary statistic per-subject slope of AP-partial ρ versus cc across the full EC-reflow sweep (0, 0.002, 0.005, 0.008, 0.010, 0.015, 0.020, 0.030, 0.050, 0.080); primary contrast navigation slope greater

than sleep slope; confirmatory fixed points $cc = 0.005$ and $cc = 0.030$; hemisphere route `both_mirrored`; bipolar acquisition; minimum 30 contacts per subject; minimum 60 mm AP span; 15 mm maximum surface distance; and minimum 8 usable subjects.

The resulting analysis produced the stronger form of the state dissociation: the uncoupled navigation baseline was approximately zero in ds005522 and negative in ds005523, while EC-reflow coupling moved both datasets into positive alignment. Coverage matching and AP-binned spatial nulls addressed the main cross-cohort and AP-gradient objections.

14 Wilson-Cowan model-class control

The Wilson-Cowan baseline tests whether the empirical alignment follows from placing any dynamical simulator on the BigBrain surface. It does not. The tested TVB-class model was near-null in sleep and navigation, whereas specific neural-field operating points aligned with the empirical data. In sleep, the full-resolution uncoupled NFT is the decisive model-class contrast. In navigation, the positive operating points are the coarse 30k uncoupled NFT and the full-resolution EC-coupled NFT at $cc = 0.030$. The interpretation is bounded to the tested Wilson-Cowan configuration; no claim is made that every possible neural-mass parameterization would be null.

15 BigBrainWarp and fMRI reconstruction infrastructure

The fMRI phase began by validating BigBrainWarp infrastructure. A synthetic round-trip from fsLR-32k to BigBrain and back produced mean Pearson correlation approximately 0.986. This scalar is a synthetic surface-resampling fidelity benchmark. It measures overall round-trip infrastructure fidelity rather than per-contrast reconstruction accuracy. The per-contrast fMRI benchmark used for the 43/47 comparison is stored separately in `figure_4/outputs/warp_ceiling_per_task.csv`; its benchmark values have median $r = 0.928$ and range 0.802–0.948 across the 47 task contrasts. Because these per-contrast values are lower than the scalar round-trip summary, they also illustrate the wavelength-dependent transfer attenuation described in the main article, Results, “Geometry and connectivity dominate different spatial scales in task fMRI”.

The Pang fsLR reconstruction baseline was reproduced on 47 HCP task contrasts. Task maps were then warped to BigBrain space and stacked in a common left-hemisphere vertex array. BigBrain geometric modes came from the cortical Laplace-Beltrami basis. BigBrain EDR modes came from a distance-dependent graph whose weights decay with inter-vertex distance. The graph Laplacian was $L_G = D - W$, where D is the degree matrix. The first 200 EDR modes were used for reconstruction comparisons.

16 fMRI statistical comparisons and scale calibration

For each task and mode count, reconstruction accuracy was the Pearson correlation between empirical and reconstructed maps. The central contrast was BigBrain geometry minus BigBrain EDR (Cell C minus Cell D). Negative values indicate stronger EDR reconstruction; positive values indicate stronger geometric reconstruction. C–D values are median across-task differences. Across the reported mode counts, EDR dominated broad long-wavelength patterns, the bases were statistically indistinguishable at $N = 50$ ($p = 0.88$), and geometry dominated at finer scales. The full table of native-metric and matched-metric medians with Holm-corrected p -values is released with the reproduction archive. Wavelength labels were derived from the BigBrain geometric eigenvalue spectrum and the observed scale ratio between BigBrain and fsLR surfaces.

17 Compute and reproducibility notes

The GPU implementation enables full-resolution BigBrain simulations that would be impractical with the original CPU implementation. Full-resolution EC-reflow simulations required 1.2 million integration steps for a 60 s run, completing in approximately 42 minutes on an NVIDIA A100. Peak allocation of the main state and stochastic-noise buffers ($Y + dW$) was 18.84 GB; total device-resident peak was not recorded.

The reproducible figure repository stores compact CSV/NPZ inputs rather than raw iEEG signals. Each figure has a preparation script that refreshes compact inputs from local analysis outputs, a renderer that writes PDF/SVG/PNG, and a verifier that checks the numerical claims used in the figure. This separates scientific provenance from raw-data redistribution and makes the plotted manuscript values auditable.

18 BigBrainWarp surface-warp parameters and provenance

The fMRI reconstruction analysis used BigBrainWarp only as a surface-metric resampling infrastructure step. We did not estimate a new registration, tune a registration loss, or use BigBrainWarp as an empirical atlas source. Its role was to move continuous HCP/Pang group-average task maps from fsLR-32k surface space into BigBrain histological surface space so that the same task maps could be reconstructed using BigBrain geometric and EDR bases.

The software environment was the Docker image `caseypaquola/bigbrainwarp:latest`, with archived image digest `sha256:fdad755a90c865b72c65de827e144d715cdea35bc749351a0120cfd136080614`. The task-map batch was run on 47 left-hemisphere HCP contrasts. The input files were plain-text fsLR-32k surface metrics and the outputs were BigBrain `.shape.gii` files, later stacked into a $163,842 \times 47$ left-hemisphere vertex array. The command family used for each task was

```
bigbrainwarp \
  --in_lh <task_lh.txt> \
  --in_rh <task_lh.txt> \
  --in_space fs_LR \
  --out_space bigbrain \
  --interp linear \
  --desc <task> \
  --wd mount/bigbrain_groupmean
```

BigBrainWarp required both hemisphere arguments for this route, whereas the Pang-style task maps used here were left-hemisphere only. We therefore passed the left-hemisphere map as both `-in_lh` and `-in_rh` and used only the resulting left-hemisphere BigBrain output. This workaround affects the discarded right-hemisphere output only and does not mix hemispheres in the reported left-hemisphere analyses.

The parameter choices were appropriate for continuous surface scalar maps. The input space was explicitly `fs_LR`, the output space was explicitly `bigbrain`, and interpolation was explicitly set to `linear`. For the BigBrainWarp surface route, the output type was the surface default and the resulting files were `tpl-bigbrain_hemi-L_desc-*.shape.gii`. The `-out_den` option was intentionally omitted for the forward fsLR-to-BigBrain warp because that option applies only when the output space is `fs_LR`. The volume-resolution option `-out_res` was not applicable because no volume warp was used. The round-trip control used the same surface-warp route in reverse, with `-in_space bigbrain`, `-out_space fs_LR`, `-interp linear`, and `-out_den 32`.

Output GIFTI provenance confirmed that, for these surface metric maps, BigBrainWarp invoked Connectome Workbench surface resampling rather than an additional fitted transform. The

effective lower-level operation was `wb_command -metric-resample` with `BARYCENTRIC` interpolation through BigBrainWarp’s precomputed spherical surface registrations. The provenance included the fsLR-32k to fsLR-164k resampling step and the subsequent resampling from fsLR-164k spherical coordinates to the BigBrain sphere using the provided multimodal surface-matching transforms. Thus the analysis used BigBrainWarp’s supplied transforms and Workbench barycentric surface-metric resampling; no empirical task-fMRI information entered the registration step.


As an infrastructure check, we ran a synthetic round-trip `fsLR-32k → BigBrain → fsLR-32k`. The forward control used `-in_space fs_LR, -out_space bigbrain, -bb_space histological, -interp linear, and -out_type surface`. The reverse control used `-in_space bigbrain, -out_space fs_LR, -bb_space histological, -interp linear, -out_type surface, and -out_den 32`. The mean round-trip Pearson correlation across hemispheres was 0.986, which we used as a synthetic surface-resampling fidelity benchmark. It is not a task-by-task empirical reconstruction score and was not used to optimize model parameters; the task-specific fMRI comparison uses the archived per-contrast round-trip-fidelity values in `figure_4/outputs/warp_ceiling_per_task.csv`.

19 EC-refrow sweep spectral-regime validation

As a model-regime validation, we regenerated the modal spectral diagnostic for the archived EC-refrow BigBrain sweep used in Fig. 1. The diagnostic was adapted from the original BigBrain sweep-analysis script used to reproduce the coupled cortico-hippocampal spectra of [1], but was run entirely from the reproducibility archive: the ten archived EC-refrow Narval simulations ($cc = 0, 0.002, 0.005, 0.008, 0.010, 0.015, 0.020, 0.030, 0.050, 0.080$) were analysed without empirical data.

The resulting spectra matched the expected qualitative regime of the coupled cortico-hippocampal NFT model. Cortical modal power showed dominant alpha-band structure, with the coupled traces peaking near 11–12 Hz, while hippocampal modal power remained theta-band dominated, with peaks in the approximately 3–5 Hz range. Increasing coupling shifted and sharpened the dominant spectral peaks through the moderate-coupling range, consistent with the spectral behaviour described by [1]. The highest coupling point, $cc = 0.080$, exceeded the divergence threshold ($\phi_{\text{RMS}} > 1000$) and was excluded from the diagnostic, matching the interpretation that excessive coupling leaves the stable spectral regime.


This check is not an additional empirical validation and is not used to compute the sleep or navigation alignment statistics. Its role is narrower: it verifies that the archived EC-refrow simulations used for the state-gating analysis retain the expected cortical-alpha and hippocampal-theta model regime before their cortical theta maps are sampled at empirical electrodes.



Supplementary Figure 1: **Spectral-regime validation of the archived EC-refrow sweep.** Modal PSDs and coupling-by-frequency heatmaps show cortical alpha-band and hippocampal theta-band structure across non-divergent EC-refrow simulations. The diagnostic was regenerated from the archived simulation NPZs without empirical data.

20 Exploratory right-hemisphere lateralization sensitivity analysis

As an exploratory lateralization sensitivity analysis, we prepared homologous right-hemisphere BigBrain cortical and hippocampal assets using the same numerical specification and downstream processing stages as the left-hemisphere production analysis [2, 3]. The right hippocampal surface was Taubin-smoothed, rigidly aligned by iterative closest point registration to the right MNI152NLin2009bSym hippocampal reference, decomposed into 300 mass-weighted eigenmodes, and used with the hemisphere-specific cortical basis to rebuild the EC-refrow coupling bundle before simulation on Narval. The right-hemisphere analysis therefore represents a second anatomical realization rather than mirrored left-hemisphere data.



Supplementary Figure 2: **Left and right cortico-hippocampal mesh assets.** The BigBrain cortical surfaces have matched mesh cardinality (163,842 vertices and 327,680 faces), as do the hippocampal surfaces (32,004 vertices and 63,250 faces). The hippocampal surfaces were processed through the same smoothing, alignment, eigenmode, and coupling-asset workflow.

Supplementary Table 1: **Left–right anatomical and spectral asset comparison.** The two realizations have matched mesh and modal dimensions while preserving their hemisphere-specific geometry.

Asset or processing stage	Left hemisphere	Right hemisphere
Cortical vertices	163,842	163,842
Cortical faces	327,680	327,680
Cortical surface area	134,238.8 mm ²	136,148.8 mm ²
Hippocampal vertices	32,004	32,004
Hippocampal faces	63,250	63,250
Hippocampal surface area	1,057.90 mm ²	1,038.47 mm ²
Cortical eigenmode array	163,842 $\times$ 300	163,842 $\times$ 300
Hippocampal eigenmode array	32,004 $\times$ 300	32,004 $\times$ 300
Sum of retained cortical eigenvalues	4.14	4.09
Sum of retained hippocampal eigenvalues	492	499
Hippocampal Taubin smoothing	5 iterations	5 iterations
ICP symmetric mean distance	1.119 mm	1.051 mm
ICP RMS distance	1.692 mm	1.589 mm
ICP p95 distance	3.032 mm	2.534 mm
Approximate Hausdorff distance	11.530 mm	10.380 mm
$C \leftarrow H$ matrix dimensions	163,842 $\times$ 32,004	163,842 $\times$ 32,004
$C \leftarrow H$ nonzero entries	9,771,072	9,895,008
$H \leftarrow C$ matrix dimensions	32,004 $\times$ 163,842	32,004 $\times$ 163,842
$H \leftarrow C$ nonzero entries	9,771,072	9,895,008

Spectral homology and mode organization. The right cortical eigenmode array was distinct from the accepted left array by SHA-256 and was loaded through the hemisphere-specific simulation path. Nevertheless, the cortical spectra were closely homologous. Under uniform area scaling, the 1.42% larger right cortical surface predicts a right/left eigenvalue ratio of 0.98597. The observed median ratio was 0.98440 over the first 50 nonconstant modes and 0.98709 over all 299, with a log-spectrum correlation of 0.99994. These values are consistent with a homologous right cortical basis derived from the slightly larger surface.

The hippocampal spectra were similarly concordant. The surface-area difference predicts a right/left eigenvalue ratio of 1.01872, compared with an observed median ratio of 1.01479 across the 299 nonconstant modes and a log-spectrum correlation of 0.999914. Individual hippocampal eigenfunctions nevertheless showed mode mixing when the homologous geometries were compared. This behavior has a direct precedent in cortical perturbation theory, where modes that lose one-to-one correspondence can remain related through rotation within the same low-dimensional eigenspace [4].

Median absolute same-index correlation declined from 0.870 over the first 10 hippocampal modes to 0.719 over the first 20 and 0.337 over the first 50. Mass-weighted principal-angle analysis, however, showed substantial overlap between the corresponding subspaces: mean $\cos^2 \theta$ was 0.911, 0.805, and 0.741 over the first 10, 50, and 299 modes, respectively. Allowing modest index tolerance captured 0.917 of the first 20 left modes within the first 30 right modes and 0.885 of the first 50 within the

first 70. The two bases are therefore spectrally homologous and occupy strongly related subspaces while retaining hemisphere-specific spatial organization.

Coupling-operator validation. Reference-rowwise scaling produced closely matched sparse coupling magnitudes across hemispheres (Supplementary Table 2). The total $C \leftarrow H$ weight differed by only 0.39%, its median nonzero row sum by 0.04%, and the total reciprocal $H \leftarrow C$ weight by 2.80%. Sparse coupling magnitude is therefore well controlled in the comparison.

Supplementary Table 2: **Coupling-operator comparison after reference-rowwise scaling.** Sparse coupling magnitudes are closely matched, while projection through the hemisphere-specific modal systems preserves differences in effective drive.

Quantity	Left	Right	Right versus left
$C \leftarrow H$ total weight	1,646,581.9	1,653,065.1	+0.39%
$C \leftarrow H$ median nonzero row sum	16.3333	16.3263	-0.04%
$C \leftarrow H$ nonzero cortical rows	101,782	103,073	+1.27%
$H \leftarrow C$ total weight	893,044.4	918,003.9	+2.80%
$\ P_C W_{C \leftarrow H} \Phi_H\ _F$	319.26	242.31	-24.1%
$\ P_H W_{H \leftarrow C} \Phi_C\ _F$	3.984	3.801	-4.62%

After modal projection, the cortex-receiving operator differed more clearly: $\|P_C W_{C \leftarrow H} \Phi_H\|_F$ was 319.26 on the left and 242.31 on the right, whereas the reciprocal operator was 3.984 and 3.801. Thus, the sparse-domain normalization succeeded, while projection through the hemisphere-specific bases and mass-weighted projectors produced a measurable difference in effective modal drive. This separation between broad spectral homology and greater sensitivity of projected dynamics is consistent with neural-field theory, which predicts that gross shape stabilizes the orientation of low-order modes while finer geometric differences can affect eigenvalues, eigenfrequencies, and dynamical behavior [4].

Right-hemisphere simulation benchmark. We simulated the right-hemisphere bundle on the unchanged ten-point coupling grid, with $ch = 0.2cc$, 300 + 300 modes, seed 73924, time step 5×10^{-5} s, and duration 60 s. All ten simulations completed. As in the left-hemisphere diagnostic, $cc = 0.080$ exceeded the prespecified $\phi_{\text{RMS}} > 1000$ threshold and was excluded from the spectral-regime diagnostic only ($\phi_{\text{RMS}} = 27,200$ left and 20,368 right).

The two realizations produced closely matched dynamics through $cc = 0.020$. At that value, the cortical global spectral maximum was 11.85 Hz in both hemispheres and field amplitudes differed by less than 2%. Across the coupled mid-range, both realizations reproduced the prominent cortical alpha-band and hippocampal theta-band structure expected from the cortico-hippocampal model [1]. The cortical alpha feature should be distinguished from the global spectral maximum at the lowest coupling values, which remained at 1.81 Hz through $cc \leq 0.010$.

At $cc = 0.030$, hippocampal peaks remained nearly identical (5.22 Hz left and 5.15 Hz right), whereas the cortical global maximum was 12.33 Hz on the left and 2.93 Hz on the right. Cortical theta fraction remained 0.189 and 0.163, respectively, demonstrating redistribution of spectral power rather than loss of theta-band activity.

The coarse single-seed sweeps therefore bracketed the upper cortical dominant-frequency transition between $cc = 0.020$ and 0.030 on the right and between $cc = 0.030$ and 0.050 on the left. These intervals describe a reproducible difference in the sampled realizations but do not estimate a precise transition threshold.

Whole-brain models commonly explore global coupling over a range and select an operating point according to agreement between simulated and empirical observables [5]. Retuning the model separately for each hemisphere could therefore absorb this difference into the fitted parameter. Instead, retaining the left-defined operating point made transfer across anatomical realizations directly testable. This choice is also supported by personalized network modeling showing that global-coupling transition values can vary with the underlying anatomical realization [6].

Exploratory empirical transfer. The right-hemisphere empirical analyses were treated as exploratory because electrode coverage was not bilaterally matched. Nevertheless, the sleep-side direction was preserved in two complementary contact-level checks. In clinically normal-like Penn SIEVE N3 contacts, the $cc = 0.030$ minus $cc = 0$ AP-partial alignment difference was -0.180 on the left and -0.143 on the right. In ds004551, the released coordinates contained no laterally placed right-hemisphere contacts: the maximum reported x coordinate was $+1.49$ mm. After empirical-data filtering, 18 analysable near-midline positive- x contact-session observations remained across eight subjects, with at most three observations per subject. Their pooled AP-partial slope was negative (-2.058), but we report this for completeness rather than as an independent right-hemisphere check.

Navigation transfer was dataset-dependent. The right-hemisphere subject-level slope was -0.419 in ds005522 ($n = 34$, one-sided $p = 0.734$), compared with $+1.16$ on the left. In ds005523 retrieval, the right slope remained positive but was attenuated to $+0.605$ ($n = 14$, $p = 0.251$), compared with $+2.85$ on the left. Thus, the sleep-side direction generalized across the available right-hemisphere checks, whereas the left-hemisphere navigation benefit did not transfer consistently to the second realization.

Scope. This exploratory benchmark demonstrates that the same anatomical specification, coupling construction, simulation configuration, and empirical interface can be applied to a second hemisphere-specific realization. It also shows that holding the operating point fixed exposes informative sensitivity to the underlying realization.

Independent histological evidence provides anatomical motivation for treating the two BigBrain hemispheres as related but nonidentical substrates. A layer-resolved analysis of the same donor found higher left-hemisphere staining intensity anteriorly and lower intensity posteriorly, with an anterior-posterior asymmetry in superficial layers and an inferior-superior pattern in deeper layers [7]. The post-mortem asymmetry map also corresponded to population in-vivo microstructural asymmetry in the Human Connectome Project ($N = 1101$, $r = 0.482$), indicating that the observed pattern was not restricted to this donor [7]. This is particularly relevant because the empirical alignment statistic explicitly conditions on anterior-posterior position.

The present comparison is not, however, a formal test of biological lateralization. Cortical and hippocampal geometry, surface masses, eigenfunctions, projected coupling operators, and empirical coverage vary together. Moreover, ds004551 does not provide sufficient right-hemisphere coverage for the primary subject-level sleep statistic. Isolating individual anatomical contributions will require component-swap simulations, while a formal bilateral empirical test will require matched electrode coverage. Within those bounds, the right-hemisphere benchmark establishes technical portability, provides convergent directional support for the sleep-side finding, and identifies a clear path for testing how state-dependent coupling transfers across anatomical realizations.

Supplementary References

- [1] Phogat, R. et al. A unified model of cortico-hippocampal interactions through neural field theory. Preprint at <https://doi.org/10.1101/2025.09.07.674721> (2025).
- [2] Amunts, K. et al. BigBrain: an ultrahigh-resolution 3D human brain model. *Science* **340**, 1472–1475 (2013).
- [3] DeKraker, J. et al. Automated hippocampal unfolding for morphometry and subfield segmentation with HippUnfold. *eLife* **11**, e77945 (2022). <https://doi.org/10.7554/eLife.77945>.
- [4] Gabay, N. C. & Robinson, P. A. Cortical geometry as a determinant of brain activity eigenmodes. *Phys. Rev. E* **96**, 032413 (2017).
- [5] Cofré, R. et al. Whole-brain models to explore altered states of consciousness from the bottom up. *Brain Sci.* **10**, 626 (2020). <https://doi.org/10.3390/brainsci10090626>.
- [6] 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). <https://doi.org/10.1371/journal.pcbi.1006487>.
- [7] Wan, B. et al. Microstructural asymmetry in the human cortex. *Nat. Commun.* **15**, 10124 (2024). <https://doi.org/10.1038/s41467-024-54243-9>.