Spectral graph theory of brain oscillations

Ashish Raj[1*], Chang Cai[1], Xihe Xie[2], Eva Palacios[1], Julia Owen[3], Pratik Mukherjee[1,4], Srikantan Nagarajan[1,4*]

[1]Department of Radiology and Biomedical Imaging, University of California, San Francisco, CA

[2]Department of Neuroscience, Weill Graduate School of Medicine, Weill Cornell Medicine, New York, NY

[3]Department of Radiology, University of Washington, Seattle, WA

[4]Department of Bioengineering and Therapeutic Sciences, University of California, San Francisco, CA

[*]Address Correspondence to:

Ashish Raj, PhD
Department of Radiology and Biomedical Imaging
UCSF Box 0946
University of California, San Francisco
185 Berry St, Suite 370, San Francisco, CA 94143
Email: ashish.raj@ucsf.edu

Srikantan Nagarajan, PhD
Department of Radiology and Biomedical Imaging
University of California at San Francisco
Email: sri@ucsf.edu

Keywords: spectral graph theory, connectomes, magnetoencephalography, brain activity, alpha rhythm

Short title: Linear graph spectral model of brain oscillations

Abstract

The relationship between the brain's structural wiring and the functional patterns of neural activity is of fundamental interest in computational neuroscience. We propose a linear, hierarchical graph spectral model of brain activity at mesoscopic and macroscopic scales that accurately predicts spatial and spectral features of neural oscillations across the brain. This novel model yields an elegant closed-form solution of the structure-function problem specified by the graph Laplacian spectrum of the structural connectome with simple, universal rules of dynamics specified by few unknown parameters. This parsimony stands in contrast to conventional complex numerical simulations of coupled non-linear lumped neural mass models (NMM). The model was highly successful in reproducing empirical spatial and spectral patterns of activity measured by scalp magneto-encephalography (MEG) after source localization, in contrast to NMM. The model may represent an important step towards understanding the fundamental relationship between network topology and the macroscopic whole-brain dynamics.

Introduction

The Structure-Function Problem in Neuroscience

It is considered paradigmatic in neuroscience that the brain's structure at various spatial scales is critical for determining its function. In particular, the relationship between the brain's *structural wiring* and the *functional* patterns of neural activity is of fundamental interest in computational neuroscience. Brain structure and function at the scale of macroscopic networks, i.e. amongst identifiable GM regions and their long-range connections through WM fiber bundles, can be adequately measured using current non-invasive measurement techniques. Fiber architecture can be measured from diffusion tensor imaging (DTI) followed by tractography algorithms [1,2]. Similarly, brain function manifested in neural oscillations can be measured non-invasively using magnetoencephalography (MEG) and reconstructed across whole-brain networks. Does structure constrain functional activity patterns that arise on the macroscopic network or graph, whose nodes represent gray matter regions, and whose edges have weights given by the structural connectivity (SC) of white matter fibers between them? We address this critical open problem here, as the structural and functional networks estimated at various scales are not trivially predictable from each other [3].

Although numerical models of single neurons and local microscopic neuronal assemblies, ranging from simple integrate-and-fire neurons to detailed multi-compartment and multi-channel models [4–8] have been proposed, it is unclear if these models can explain structure-function coupling at meso- or macroscopic scales. At one extreme, the Blue Brain Project [9,10] seeks to model in detail all 10^{11} neurons and all their connections in the brain. Indeed spiking models linked up via specified synaptic connectivity and spike timing dependent plasticity rules were found to produce regionally and spectrally organized self-sustaining dynamics, as well as wave-like propagation similar to real fMRI data [11]. However, it is unclear whether such efforts will succeed in providing interpretable models at whole-brain scale [12].

Therefore the traditional computational neuroscience paradigm at the microscopic scale does not easily extend to whole-brain macroscopic phenomena, as large neuronal ensembles exhibit emergent properties that can be unrelated to individual neuronal behavior [13–18], and are instead largely governed by long-range connectivity [19–22]. At this scale, graph theory involving network statistics can phenomenologically capture structure-function relationships [23–25], but do not explicitly embody any details about neural physiology [14,15]. Strong correlations are known between functional and structural connections at this scale [3,18,26–31], and important graph properties are shared by both SC and FC networks, such as small worldness, power-law degree distribution, hierarchy, modularity, and highly connected hubs [24,32].

A more detailed accounting of the structure-function relationship requires that we move beyond statistical descriptions to mathematical ones, informed by computational models of neural activity. Numerical simulations are available of mean field [17,33,34] and neural mass [22,35] approximations of the dynamics of neuronal assemblies. By coupling many such neural mass models (NMMs) using anatomic connectivity information, it is possible to generate via large-scale stochastic simulations a rough picture of how the network modulates local activity at the global scale to allow the emergence of coherent functional networks [22]. However, simulations are unable to give an analytical (i.e. closed form) encapsulation of brain dynamics and present an interpretational challenge in that behavior is only deducible indirectly from thousands of trial runs of time-consuming simulations. Consequently, the essential minimal rules of organization and dynamics of the brain remain unknown. Furthermore, due to their nonlinear and stochastic nature, model parameter inference is ill-posed, computationally demanding and manifest with inherent identifiability issues [36].

How then do stereotyped spatiotemporal patterns emerge from the structural substrate of the brain? How will diseases perturb brain structure, thereby impacting its function? While stochastic simulations are

powerful and useful tools, they provide limited neuroscientific insight, interpretability and predictive power, especially for the practical task of inferring macroscopic functional connectivity from long-range anatomic connectivity. Therefore, there is a need for more direct models of structural network-induced neural activity patterns – a task for which existing numerical modeling approaches, whether for single neurons, local assemblies, coupled neural masses or graph theory, are not ideally suited.

A hierarchical, analytic, low-dimensional and linear spectral graph theoretic model of brain oscillations

Here we present a new linear graph model capable of reproducing empirical macroscopic spatial and spectral properties of neural activity. We are interested specifically in the transfer function induced by the macroscopic structural connectome, rather than in the behavior of local neural masses. Therefore we seek an explicit formulation of the frequency spectra induced by the graph, using the eigen-decomposition of the structural graph Laplacian, borrowing heavily from **spectral graph theory** used in diverse contexts including clustering, classification, and machine learning [37–40]. This theory conceptualizes brain oscillations as a linear superposition of eigenmodes. These eigen-relationships arise naturally from a biophysical abstraction of fine-scaled and complex brain activity into a simple linear model of how mutual dynamic influences or perturbations can spread within the underlying structural brain network, a notion that was advocated previously [18,41,42]. We had previously reported that the brain network Laplacian can be decomposed into its constituent "eigenmodes", which play an important role in both healthy brain function [18,30,43–45] and pathophysiology of disease [43,46–48]. We showed that such a spectral graph theoretic model is capable of reproducing spatial patterns of functional connectivity in the BOLD fMRI regime [30], whereby FC and SC share eigenvectors and their eigenvalues are exponentially related. Just a few Laplacian eigenmodes could reconstruct entire FC matrices, and significantly outperform the far more complex and time-consuming generative neural mass model of [3,35] on the same empirical data.

Unlike prior work however, we show here that a graph-spectral decomposition is possible at much higher frequencies as well, ignoring all non-linearities that are operating at the local (node) level. Like previous NMMs, we lump neural populations at each brain region into neural masses, but unlike them we use a linearized (but frequency-rich) local model – see **Figure 1A**. The macroscopic connectome imposes a linear and deterministic modulation of these local signals, which can be captured by a *network transfer function*. The sequestration of local oscillatory dynamics from the macroscopic network in this way enables the characterization of whole brain dynamics deterministically in closed form in Fourier domain, via the eigen-basis expansion of the network Laplacian.

We applied this model to real connectivity data and validated it against measured source-reconstructed MEG recordings in healthy subjects. The model closely matches empirical spatial and spectral MEG patterns. In particular, the model displays a prominent alpha peak, and, intriguingly, the eigenmodes corresponding to the alpha oscillations have the same posterior-dominant spatial distribution that is repeatedly seen in eyes-closed alpha power distributions. To our knowledge, this is the first network-based explanation of how the alpha power is distributed over the brain, and indeed of the entire spatio-spectral patterning of brain rhythms measurable on MEG.

Results

Closed form solution of steady state spectra

The steady state spectral response induced by the connectome at angular frequency ω, can be expressed as a summation over the eigenmodes $\boldsymbol{u_i}(\omega)$ and eigenvalues $\lambda_i(\omega)$ of the graph Laplacian $\boldsymbol{\mathcal{L}}(\omega)$:

$$\boldsymbol{X}(\omega) = \sum_i \frac{\boldsymbol{u_i}(\omega)\boldsymbol{u_i^H}(\omega)}{j\omega + \frac{1}{\tau_G}\lambda_i(\omega)F_e(\omega)} H_{local}(\omega)\boldsymbol{P}(\omega)$$

τ_G is a time constant, $F_e(\omega)$ is a gamma-shaped neural response function, and $H_{local}(\omega)$ is linearized lumped local spectral response (derivation can be found in Methods). The spectral-domain output $\boldsymbol{X}(\omega)$ and input $\boldsymbol{P}(\omega)$ are vector-valued variables. This steady state model of brain activity includes only 7 global model parameters.

Graph Laplacian eigenmodes mediate a diversity of frequency responses

First we demonstrate the spectra produced by graph eigenmodes as per our theory. Each eigenmode produces a frequency response based on its frequency-dependent eigenvalue (**Figure 1D,E**). Small eigenvalues undergo a larger shift due to frequency, while the large ones stay more stable and tightly clustered around the nominal eigenvalue (i.e. at $\omega = 0$) (see **Figure 1C**). **Figure 1D** shows the transit in the complex plane of a single eigenmode's frequency response, starting at low frequencies in the bottom right quadrant, and moving to the upper left quadrant at high frequencies. The magnitude, given by distance from origin, suggests that most eigenmodes have two prominent lobes, roughly corresponding to alpha and beta rhythms, respectively. In contrast, the lowest few eigenmodes start off far from the origin, indicative of a low-pass response. The magnitude of these complex-valued curves shown in figure **1E** reinforces these impressions, with clear alpha and beta peaks, as well as slower rhythms of the lowest eigenmodes.

The spatial patterns of the first 4 eigenmodes of $\boldsymbol{\mathcal{L}}(\omega)$, evaluated at the alpha peak of 8 Hz, are shown in **Figure 1F**. The first eigenmode u_1 is essentially a sensorimotor network which, due to its lack of an alpha peak (see panel **E**), does not contribute to the spatial patterning of theoretical alpha power. The eigenmodes $u_2 - u_4$ give strong alpha frequency responses, and in turn are strongly distributed spatially in posterior areas but also include other regions and prominently resemble many elements of the **default mode network** and the so-called **structural core** of the human connectome (See also **Figure 5)**. While $u_1 - u_4$ are highly consistent and reproducible, higher modes are increasingly sensitive to axonal velocity and frequency (not shown here).

Since the model relies on connectome topology, we demonstrate in **Figure 2** that different connectivity matrices produce different frequency responses: A) the individual's structural connectivity matrix, B) HCP average template connectivity matrix, C) uniform connectivity matrix of ones, and D) a randomly generated matrix. All modeled power spectra show a broad alpha peak at around 10 Hz and a narrower beta peak at around 20 Hz. This is expected, since these general spectral properties are governed by the local linearized neural mass model. The alpha peak is predominantly contained in the low eigenmodes, up to eigen-index 10 or so. Although the alpha and beta peaks are innately present under default parameters in Figure 1, once we optimize parameters, the peaks become stronger. However, it is important to note that different eigenmodes accommodate a diversity of frequency responses; for instance the lowest eigenmodes show a low-frequency response with no alpha peak whatsoever. In the frequency responses from biologically realistic individual and HCP template connectomes, there is a diversity of spectral responses amongst eigenmodes that is lacking in the response produced by the unrealistic uniform and randomized matrices. Since graph topology appears so critical to the power spectrum it induces, we explored whether and how sparsity of random graphs mediates spectral power (**Figure 2D-F**). At incrementally increasing sparsity levels, the diversity of spectral responses of different eigenmodes increases and approaches that of realistic connectomes. Therefore, graph eigenmodes induce unique and diverse frequency responses that depend strikingly on the topology of the graph.

Network eigenmodes exhibit strong spatial patterning in their frequency responses, even with identical local oscillations (**Figure 3**). We evaluated the model spectral response using the subject-specific $C^{individual}$ matrices of 4 representative subjects (**Figure 3A-D**). The model power spectra strikingly resembles empirical MEG spectra, correctly displaying both the alpha and beta peaks on average, and similar regional variability as in real data.

Regional averages of empirical and modeled power spectra of the entire group after full parameter optimization over individual subjects using the proposed simulated annealing algorithm are shown in figure 3E. The model closely replicates the observed power spectrum equally well with both $C^{individual}$ (green) and $C^{template}$ (blue). Thus, in most cases we can safely replace the subject-specific connectome with the template connectome. In contrast, when non-optimized default parameters were used (black), it resulted in a bad fit, especially at high frequencies, suggesting that individualized parameter optimization is essential to produce realistic spectra. The Wilson-Cowan neural mass model [34,49] using our in-house MATLAB implementation [36], was generally able to produce characteristic alpha and beta frequency peaks (yellow) but does not resemble empirical wideband spectra. Note that no regionally-varying NMM parameters were used in order to achieve a proper comparison with our model, but both models were optimized with the same simulated annealing algorithm. Nevertheless, these data confirm our intuition that the average spectral power signal can be produced by almost any neural model, whereas its regional variations around the canonical spectrum are presently being modeled via the connectome. Finally, no model is capable of reproducing higher frequencies in the higher beta and gamma range seen in MEG, since by design and by biophysical intuition these frequencies arise from local neural assemblies rather than from modulation by macroscopic networks.

Parameter optimization via simulated annealing

The simulated annealing optimization algorithm provided a set of optimized parameters $\{\tau_e, \tau_i,\ \tau_c,\ g_{ei},\ g_{ii},\ \alpha,\ \upsilon\}$: (see supplementary **Table 1)**. **Figure 4A** shows violin plots of the optimized values, indicating that the range is adequate for parameter exploration. The time constants τ_e, τ_i showed tight clustering but the rest of the parameters showed high variability across subjects. The optimal parameters are in a biologically plausible range, similar to values reported in numerous neural mass models. The annealing algorithm aimed to maximize the Pearson's correlation between MEG and modeled spectra ("Spectral correlation"). The convergence plots shown in **Figure 4B**, one curve for each subject, indicates substantial improvement from default choice as optimization proceeds. The distribution of optimized spectral correlations is shown in 4**C-F**. Panel **C** shows the correlation distribution after grid search, while **D** shows the distribution after annealing, which clearly gives better R values. Panel **E** shows annealing using not individual connectomes but the template HCP one. As stated earlier, there is no discernible difference in performance due to this replacement. Panel **F** shows the performance using identical parameters for all subjects – as expected it performs less well compared to the individually optimized model, which gives a unimodal distribution of Pearson's R centered around $\mathrm{R} = 0.83, \mathrm{p} < 10^{-6}$.

Graph spectral model recapitulates the spatial distribution of MEG power

Next we establish that the model can correctly reproduce region-specific spectra, even though it uses identical local oscillations. We integrated the spectral area in the range 8-12 Hz for alpha and 13-25 Hz for beta, of each brain region separately. We define "*spatial correlation*" (as compared to spectral correlation above) as Pearson's R between the *regional distribution* of empirical MEG and model-predicted power within a given frequency band. In this paper we focus exclusively on the alpha and beta bands.

Specific eigenmodes capture spatial distributions of alpha and beta band activity. We plotted the spatial correlation achieved by each eigenmode against empirical MEG regional alpha and beta power,

averaged over all subjects in **Figure 5A**. In **Supplementary Figure 1** we show these spatial correlation curves for all 36 subjects. Only a small number of eigen-modes are tuned to each power band; alpha is generally better captured by low eigenmodes while beta by middle eigenmodes. A scatter plot of all eigenmodes' alpha and beta power spatial correlation is shown in panel **B**, suggesting that when an eigenmode is correlated to alpha power, it is roughly anti-correlated to beta power. This correlation-of-correlations is highly significantly negative ($r=-0.255$, $p<0.0001$). While on average, individual eigenmodes are not highly predictive of alpha or beta, in individual subjects they have much higher R values up to 0.5. **Figure 5C,D** show the spatial pattern of the most spatially correlated eigenmode for alpha (#3) and beta (#33) respectively. These selected eigenmodes have the expected posterior distribution for alpha and are widespread for beta. Panel **E** shows a histogram of the correlation between the alpha band and beta band spatial correlations. It can be seen that across all subjects' alpha and beta band spatial correlation curves for the eigenmodes are in turn anti-correlated. Panel **F** shows histograms of the spatial correlations across subjects of the eigenmode with maximal spatial correlation with empirical alpha (green) and the same eigenmode's spatial correlation with empirical beta (blue). Again, we can see a clear anti-correlation. Together, these results confirm that graph eigenmodes might be tuned to specific frequencies, and their spatial patterns might govern the spatial presentation of different brain rhythms.

Figure 6 depicts the spatial distribution of alpha band power (8-12 Hz) over the entire brain, and **Figure 7** shows spatial distribution of beta power (13 – 25 Hz), for a representative subject. Regions are color coded by regional power scaled by mean power over all regions. A different "glass brain" rendering is shown in **Supplementary Figures 2,3**. **Alpha power distribution**. The alpha power was best modeled by a combination of the 10 best-matching eigenmodes ($R = 0.53$). The posterior and occipital dominance of alpha power is clearly observed predicted alpha distribution, with strong effect size in temporal, occipital and medial posterior areas. However, the model predicts some lateral frontal involvement that is not observed in real data. The two strongest eigenmodes that contribute to the summed model (bottom two rows) also show very similar spatial organization. **Beta power distribution**. Empirical beta power (**Figure 7** top) is spread throughout the cortex, especially frontal and premotor cortex. A combination of five best matching eigenmodes produced the best model match to the source localized pattern ($R = 0.57$). The two best-matching eigenmodes have Pearson's correlation coefficients of 0.42 and 0.41.

Alternate non-linear model. The Wilson-Cowan neural mass model did not succeed in correctly predicting the spatial patterns of alpha or beta power – see **Supplementary Figure 4**. This could be because in our implementation we enforced uniform local parameters with no regional variability. However, this is the appropriate comparison, since our proposed model also does not require regionally-varying parameters.

Peak model performance over sorted, selected eigenmodes. Since only a few eigenmodes appear to contribute substantially, we hypothesized that spatial correlations could be improved by selecting a subset of eigenmodes. Therefore we developed a sorting strategy whereby we first rank the eigenmodes in descending order of spatial correlation for a given subject and given frequency band. Then we perform summation over only these eigenmodes according to Eq (10), each time incrementally adding a new eigenmode to the sum. The spatial correlation of these "sorted-summed" eigenmodes against empirical MEG data are plotted in **Figure 8A,B** as a function of increasing number of eigenmodes. Figure **8A** gives the spatial correlation curves for alpha band and **8B** for beta band; one curve for each subject. The thick solid curves represent the average over all subjects. The spatial correlation initially increases as we add more well-fitting eigenmodes, but peaks around 10 for alpha and 5 eigenmodes for beta power, and begins declining thereafter. Addition of the remaining eigenmodes only serves to reduce the spatial correlation. This behavior is observed in almost all subjects we studied.

The distribution of peak spatial correlations, using optimized parameters and individual connectomes of all subjects is plotted in **panel C**, as well as two alternatives: a) NMM and b) graph model with 1000 instances of 80% sparse randomly generated connectomes, and optimized parameters. The proposed model gives very strong spatial correlation in alpha band (r distribution centered at 0.6), and NMM gives very poor correlation (r centered at 0). Interestingly, the random connectomes also appear to have some ability to capture these spatial patterns (r centered at 0.35), perhaps due to the implicit search within the random eigenmodes of the best-matching ones, which on average will give at least a few eigenmodes that look like MEG power purely by chance. **Panel D** shows analogous results for beta band spatial power correlations. Again our model does the best (r distribution centered at 0.5), but its comparative performance against alternate approaches is not as striking as in alpha.

Collectively, we conclude that the graph model is able to fit both the spectral and spatial features of empirical source localized MEG data, and that the optimal fits performed on individual subjects occurs at widely varying subject-specific parameter choices.

Discussion

The proposed linear graph spectral model of neural oscillatory activity is a step towards understanding the fundamental relationship between network topology and the macroscopic whole-brain dynamics. The objective is not just to model brain activity phenomenologically, but to analytically derive the mesoscopic laws that drive macroscopic dynamics. This novel model of the structure-function relationship has the following key distinguishing features: *1) Hierarchical*: the model's complexity depends on the level of hierarchy being modeled: complex, non-linear and chaotic dynamics can be accommodated at the local level, but linear graph model is sufficient at the macro-scale. *2) Graph-based*: Macroscopic dynamics is mainly governed by the connectome, hence linear approximations allow the steady-state frequency response to be specified by the graph Laplacian eigen-decomposition, borrowing heavily from **spectral graph theory** [37–40]. *3) Analytic*: The model is available in closed form, without the need for numerical simulations. *4) Low-dimensional*: Simple, global and universal rules specified with very few parameters, all global and apply at every node, are able to achieve sufficiently complex dynamics. The model is incredibly easy to evaluate, taking no more than a few seconds per brain. We proposed a simulated annealing algorithm to infer model parameters directly from a subject's MEG data. The optimized model matches observed MEG data quite well. No time-consuming simulations of coupled neural masses or chaotic oscillators were needed; indeed, the latter greatly underperformed our model. We report several novel findings with potentially important implications, discussed below.

Recapitulating regional power spectra at all frequencies

Our main result is the robust demonstration of the model on 36 subjects' MEG data. The representative examples shown in Figures 3,6-8 indicate that the graph model recapitulates the observed source localized MEG power spectra for the 68 parcellated brain regions, correctly reproducing the prominent alpha and beta peaks. For each region, the model is also correctly able to predict the full bandwidth power spectra, including the $1/\omega$ fall-off over the entire frequency range of interest.

Revealing sources of heterogeneity in brain activity patterns

The match between model and data is strongest when the model uses empirical macroscopic connectomes obtained from healthy subjects' diffusion weighted MRI scans, followed by tractography. The use of "null" connectomes - uniform connectivity of ones and randomized connectivity matrix, respectively, did far

worse than actual human connectomes (Figure 8), supporting the fact that the latter is the key mediator of real brain activity. The match was not significantly different when using a template HCP connectome versus the individual subject's own connectomes (Figures 3E, 4C,D), suggesting that, for the purpose of capturing the gross topography of brain activity, it is sufficient to use a template connectome, and disregard individual variability.

However, this does not mean that the model is incapable of capturing individual variability: indeed, we designed a comprehensive parameter fitting algorithm using simulated annealing minimization on individual subjects' MEG data of a suitably defined cost function based on Pearson R statistic as a way to capture all relevant spectral features. Using this fitting procedure, we were able to obtain the range of optimally-fitted parameters across the entire study cohort. As shown in Figure 4A, the range is broad in most cases, implying that there is significant inter-subject variability of model parameters, even if a template connectome is used for all. We tested the possibility that a group-averaged parameter set might also succeed in matching real data on individuals. But as shown in Figure 3E, this was found to be a poor choice, supporting the key role of individual variability of model parameters (but not variability in the connectome).

Macroscopic brain rhythms are governed by the connectome

A predominant view assumes that different brain rhythms are produced by groups of neurons with similar characteristic frequencies, which might synchronize and act as "pacemakers." How could this view explain why alpha and beta power are spatially stereotyped across subjects, and why the alpha signal is especially prominent in posterior areas? Although practically any computer model of cortical activity can be tuned, with suitable parameter choice, to oscillate at alpha frequency, e.g. [5,16,20,22,35,50,51], none of them are able to parsimoniously recapitulate the posterior origin of alpha. Thus the prominence of posterior alpha might be explained by the hypothesized existence of alpha generators in posterior areas. Indeed, most oscillator models of local dynamics are capable of producing these rhythms at any desired frequency [5,35,52–54], and therefore it is common to tweak their parameters to reproduce alpha rhythm. Local networks of simulated multicompartmental neurons can produce oscillations in the range 8–20 Hz [5], and, in a non-linear continuum theory, peaks at various frequencies in the range 2–16Hz were obtained depending on the parameters [53]. Specifically, the role of thalamus as pacemaker has motivated thalamocortical models [11,16] that are capable of resonances in various ranges. Neural field models of the thalamocortical loop [16] can also predict slow-wave and spindle oscillations in sleep, and alpha, beta, and higher-frequency oscillations in the waking state. In these thalamocortical models, the posterior alpha can arise by postulating a differential effect in weights of the posterior versus anterior thalamic projections, e.g. [51]. Ultimately, hypotheses requiring local rhythm generators suffer from lack of parsimony and specificity: a separate pacemaker must be postulated for each spectral peak at just the right location [55].

An alternative view emerges from our results that macroscopic brain rhythms are governed by the structural connectome. Even with global model parameters, using the exact same local cortical dynamics captured by the local transfer function $H_{local}(\omega)$, driven by identically distributed random noise $\boldsymbol{P}(\omega)$, we are capable of predicting prominent spectral (Figure 3) and spatial (Figures 6,7) patterning that is quite realistic. This is especially true in the lower frequency range: indeed the model correctly predicts not just the frequency spectra in alpha and beta ranges, but also their spatial patterns – i.e. posterior alpha and distributed but roughly frontal beta. Although this is not definitive proof, it raises the intriguing possibility that the macroscopic spatial distribution of the spectra of brain signals *does not require spatial heterogeneity of local signal sources, nor regionally variable parameters*. Rather, it implies that the most prominent *patterning of brain activity (especially alpha) may be governed by the topology of the macroscopic network* rather than by local, regionally-varying drivers. Nevertheless, a deeper exploration is required of

the topography of the dominant eigenmodes of our linear model, in order to understand the spatial gradients postulated previously [16,51].

Emergence of linearity from chaotic brain dynamics

The non-linear and chaotic dynamics of brain signals may at first appear to preclude deterministic or analytic modeling of any kind. Yet, vast swathes of neuroscientific terrain are surprisingly deterministic, reproducible and conserved across individuals and even species. Brain rhythms generally fall within identical frequency bands and spatial maps [4,16,32]. Based on the hypothesis that the emergent behavior of long-range interactions can be independent of detailed local dynamics of individual neurons [13–18], and may be largely governed by long-range connectivity [19–22], we have reported here a minimal linear model of how the brain connectome serves as a spatial-spectral filter that modulates the underlying non-linear signals emanating from local circuits. Nevertheless, we recognize the limitations of a linear model and its inability to capture inherent non-linearities across all levels in the system.

Other limitations and extensions

The model currently examines resting-state activity, but future extensions will include prediction of functional connectivity, task-induced modulations of neural oscillations and causal modeling of external stimuli, e.g. transcranial magnetic and direct current stimulation. The current implementation does not incorporate complex local dynamics, but future work will explore using non-white internal noise and chaotic dynamics for local assemblies. This may allow us to examine higher gamma frequencies. Although our model incorporates latency information derived from path distances, we plan to explore path-specific propagation velocities derived from white matter microstructural metrics such as axon diameter distributions and myelin thickness. Future work will also examine the specific topographic features of the structural connectome that may best describe canonical neural activity spectra. Finally, we plan to examine the ability of the model to predict time-varying structure-function relationships.

Potential applications

Mathematical encapsulation of the structure-function relationship can potentiate novel approaches for mapping and monitoring brain diseases such as autism, schizophrenia, epilepsy and dementia, since early functional changes are more readily and sensitively measured using fMRI and MEG, compared to structural changes. Because of the complementary sensitivity, temporal and spatial resolutions of diffusion MRI, MEG, EEG and fMRI, combining these modalities may be able to reveal fine spatiotemporal structures of neuronal activity that would otherwise remain undetected if using only one modality. Current efforts at fusing multimodalities are interpretive, phenomenological or statistical, with limited cognizance of underlying neuronal processes. Thus, the ability of the presented model to quantitatively and parsimoniously capture the structure-function relationship may be key to achieving true multi-modality integration.

Online Methods

Spectral graph model development

Notation. In our notation, vectors and matrices are represented by boldface, and scalars by normal font. We denote frequency of a signal, in Hertz, by symbol f, and the corresponding angular frequency as $\omega = 2\pi f$. The connectivity matrix is denoted by $\boldsymbol{C} = \{c_{jk}\}$, consisting of connectivity strength c_{ij} between any two pair of regions j, k.

Canonical rate model over a graph. We use a canonical rate model to describe neural activity across two hierarchical levels – local cortical levels and long-range mesoscopic levels. At each level of the hierarchy of brain circuits, we hypothesize a simple linear rate model of recurrent reverberatory activity given by

$$\frac{dx_{e/i}(t)}{dt} = -\frac{1}{\tau_{e/i}} f_{e/i}(t) * x_{e/i}(t) + \frac{1}{\tau_{e/i}} f_{e/i}(t) * \sum_{j,k} c_{jk} x_{e/i}\left(t - \tau_{jk}^{v}\right) + p_{e/i}(t) \qquad (1)$$

where $x_{e/i}(t)$ is the mean signal of the excitatory/inhibitory populations and $p_{e/i}(t)$ is internal noise source reflecting local cortical column computations or input. The transit of signals, from pre-synaptic membranes, through dendritic arbors and axonal projections, is sought to be captured into ensemble average neural impulse response functions $f_e(t) = \frac{t}{\tau_e} \exp\left(-\frac{t}{\tau_e}\right)$ and $f_i(t) = \frac{t}{\tau_i} \exp\left(-\frac{t}{\tau_i}\right)$ respectively. We disregard the non-linearity of the neural response, hence the output at the terminal to a presynaptic input $u(t)$ is the simple convolution $x_e(t) = f_e(t) * u(t)$. The neural responses $f_{e/i}(t)$ are Gamma-shaped responses (**Figure 1B**) parameterized by time constants $\tau_{e/i}$ that here represent the end result of both synaptic membrane capacitance and the distribution of dendritic/axonal delays introduced by the arborization. NMMs typically use a single classical exponential decay term for membrane capacitance only, since NMMs model highly local cell assemblies where multisynaptic connections are infrequent and axonal and dendritic transport delays are usually incorporated explicitly via connectivity weights and delays. Since our lumped model was designed for relatively large cortical regions, we employ the Gamma-shaped $f_{e/i}$ to correctly capture not just classical membrane capacitance but also the expected diversity of dendritic transport delays. The dynamics of the entire assembly modeled via a self-decaying term $\tau_{e/i} \frac{d\boldsymbol{x}}{dt} \propto -f_{e/i}(t) * \boldsymbol{x}(t)$, typically used in most rate or NMM models, but the difference here is that we chose to apply convolution with neural response $f_{e/i}(t)$ within the decay process. We believe this is necessary to ensure that the dynamics of the population cannot participate in the internal recurrent dynamics of the region until the signal has passed through one instance of the neuronal response. Since this neural response is meant to capture a distribution of local circuit delays, its time constants $\tau_{e/i}$ are purposefully far longer (up to 20ms) than expected from membrane capacitance alone. Studies of cortical lag times using paired electrode recordings between primary and higher cortices demonstrate this. A short visual stimulus causes a neural response in the ferret V1 within 20ms post-stimulus, in the primary barrel field within 16-36ms, and the entire visual cortex becomes engaged 48-70ms after stimulus [6]. Brief deflection of a single barrel whisker in the mouse evokes a somatotopically mapped cortical depolarization that remains localized to its C2 barrel column only for a few milliseconds, thence rapidly spreading to a large part of sensorimotor cortex within tens of milliseconds, a mechanism considered essential for the integration of sensory information [56,57]. Interestingly, the evoked response curve in S1 from the [56] study had a prominent Gamma shape. Of note, the duration of S1 response (~50ms) was considerably longer than the time to first sensory response in C2 (7.2ms) [56]. Interestingly, feedback projections from higher to lower areas take ~50ms, hence have a much slower apparent propagation velocity (0.15-0.25m/s) than what would be predicted by axonal conduction alone (1-3m/s) [6].

Individual neural elements are connected to each other via connection strengths c_{jk}. Let the cortico-cortical fiber conduction speed be v, which here is assumed to be a global constant independent of the pathway under question. For a given pathway connecting regions j and k of length d_{jk}, the conduction delay of a signal propagating from region j to region k will be given by $\tau^{v}{}_{jk} = d_{jk}/v$. Hence signals from neighboring elements also participate in the same recurrent dynamics, giving the 2nd term of Eq (1). Equation (1) will serve as our canonical rate model, and will be reproduced at all levels of the hierarchy, and only the connectivity strengths will vary depending on the level of hierarchy we are modeling, as explained below.

Local neural assemblies. The local connectivities c_{jk}^{local} are assumed to be all-to-all, giving a complete graph. Further, the axonal delays τ_{jk}^{v} associated with purely local connections were already incorporated in the lumped impulse responses $f_{e/i}(t)$. Hence we assert:

$$c_{jk}^{local} = c_{e/i}, \qquad \tau_{jk}^{v} = 0, \forall\, j, k$$

From spectral graph theory, a complete graph has all equal eigenvalues which allows the local network to be lumped into gain constants, and the summation removed. Indeed, rewriting $x_{e/i}(t)$ as the mean signal of all the excitatory/inhibitory cells and setting the gains $g_{ee} = 1 - c_e N_e$ and $g_{ii} = 1 - c_i N_i$ we get

$$\frac{dx_{e/i}(t)}{dt} = -\frac{g_{ee/ii}}{\tau_{e/i}} f_{e/i}(t) * x_{e/i}(t) + p_{e/i}(t). \qquad (2)$$

Given the Fourier Transform pairs $\frac{d}{dt} \leftrightarrow j\omega$, $f_{e/i}(t) \leftrightarrow F_{e/i}(\omega) = \frac{1/\tau_{e/i}^2}{(j\omega + 1/\tau_{e/i})^2}$, we take the Fourier transform of Eq(1) and obtain the local assembly's frequency spectrum:

$$X_{e/i}(\omega) = \left(j\omega + \frac{g_{ee/ii}}{\tau_{e/i}} F_{e/i}(\omega)\right)^{-1} P_{e/i}(\omega) \qquad (3)$$

Writing this in terms of transfer functions $X_e(\omega) = H_e(\omega)P_e(\omega), X_i(\omega) = H_i(\omega)P_i(\omega)$ we get the lumped local system illustrated in **Figure 1A**. Finally, we must also account for signals that alternate between the two populations, which is given by the transfer function

$$H_{ei}(\omega) = H_e(\omega)H_i(\omega)/(1 + g_{ei}H_e(\omega)H_i(\omega))$$

We fix $g_{ee} = 1$ without loss of generality, and let the other terms g_{ii}, g_{ei} be model parameters to be fitted. Finally, the total cortical transfer function is the sum

$$H_{local}(\omega) = H_e(\omega) + H_i(\omega) + H_{ei}(\omega) \qquad (4)$$

and $X_{local}(\omega) = H_{local}(\omega)P(\omega)$ represents all neural activity in this region, whether from excitatory or inhibitory cells. The canonical local activity is therefore defined by the Fourier transform pair: $x_{local}(t) \leftrightarrow X_{local}(\omega)$.

Macroscopic scale: signal evolution on the entire graph

We use the same canonical network dynamics as Eq (1), but now the inter-regional connectivity c_{jk} is non-zero and given by the structural connectome. Similarly, axonal conductance delays are determined by fiber length and conductance speed $\tau_{jk}^{v} = v/d_{jk}$. Further, the external driving signals at each node is the local neural activity $x_{local}(t)$ defined above rather than a noise process $p(t)$. In the interest of parsimony we set each node of the macroscopic graph to have the same internal power spectrum $X_{local}(\omega)$ - i.e. all regions are experiencing the same transfer function, driven by identically distributed (but of course not identical) noise. At this scale, activity measured at graph nodes is no longer excitatory or inhibitory, but mixed, and the corticocortical connections are all between long, pyramidal excitatory-only cells. Thus for the k-th node

$$\frac{dx_k(t)}{dt} = -\frac{1}{\tau_G} f_e(t) * x_k(t) + \frac{\alpha}{\tau_G} f_e(t) * \sum_j c_{jk} x_j\left(t - \tau_{jk}^{v}\right) + x_{local,k}(t)$$

Here we have introduced a global coupling constant α, similar to most connectivity-coupled neural mass models, that seeks to control the relative weight given to long-range afferents compared to local signals. We have also introduced a new time constant, τ_G, which is an excitatory time constant and it may be the same as the previously used constant τ_e. However, we allow the possibility that the long-range projection neurons might display a different capacitance and morphology compared to local circuits, hence we have introduced τ_G (subscript G is for "graph" or "global").

Stacking all equations from all nodes and using vector valued signals $\boldsymbol{x}(t) = \{x_k(t)\}$, we can write

$$\frac{d\boldsymbol{x}(t)}{dt} = -\frac{1}{\tau_G} f_e(t) * \boldsymbol{x}(t) + \frac{\alpha}{\tau_G} f_e(t) * C\{\boldsymbol{x}(t - \tau_{jk}^{v})\} + \boldsymbol{x}_{\boldsymbol{local}}(t) \qquad (5)$$

where the braces $\{\cdot\}$ represent all elements of a matrix indexed by j, k.

We wish to evaluate the frequency spectrum of the above. In Fourier space, delays become phases; hence we use the transform pairs $\frac{d\boldsymbol{x}}{dt} \leftrightarrow j\omega \boldsymbol{X}(\omega)$ and $\boldsymbol{x}(t-\tau) \leftrightarrow e^{-j\tau\omega}\boldsymbol{X}(\omega)$. Therefore, define a "complex connectivity matrix" at any given angular frequency ω as $\boldsymbol{C}^*(\omega) = \{c_{jk} \exp(-j\omega\, \tau^{v}{}_{jk})\}$. Then define a normalized complex connectivity matrix at frequency ω as

$$\boldsymbol{\mathcal{C}}(\omega) = diag\left(\frac{1}{\sqrt{\boldsymbol{deg}}}\right) \boldsymbol{C}^*(\omega)\, diag\left(\frac{1}{\sqrt{\boldsymbol{deg}}}\right) \qquad (6)$$

Taking the Fourier transform of Eq (5), we get

$$\left(j\omega \boldsymbol{X}(\omega) + \frac{1}{\tau_G} F_e(\omega)\, \big(I - \alpha \boldsymbol{\mathcal{C}}(\omega)\big) \boldsymbol{X}(\omega)\right) = H_{local}(\omega) \boldsymbol{P}(\omega) \qquad (7)$$

where we assumed identically distributed noise signals driving both the excitatory and inhibitory local populations at each node, such that $P_{e,k}(\omega) = P_{i,k}(\omega) = P_k(\omega)$ at the k-th node. We then collected all nodes' driving inputs in the vector $\boldsymbol{P}(\omega) = \{P_k(\omega), \forall k\}$.

Defining the complex Laplacian matrix

$$\boldsymbol{\mathcal{L}}(\omega) = \boldsymbol{I} - \alpha \boldsymbol{\mathcal{C}}(\omega)$$

where $\boldsymbol{I}$ is the identity matrix of size $N \times N$, and the degree vector $\boldsymbol{deg}$ is defined as $deg_k = \sum_j c_{jk}$. This complex Laplacian will be evaluated via the eigen-decomposition

$$\boldsymbol{\mathcal{L}}(\omega) = \boldsymbol{U}(\omega) \boldsymbol{\Lambda}(\omega) \boldsymbol{U}(\omega)^H \qquad (8)$$

where $\boldsymbol{\Lambda}(\omega) = diag([\lambda_1(\omega), \dots, \lambda_N(\omega)])$ is a diagonal matrix consisting of the eigenvalues of the complex Laplacian matrix of the connectivity graph $\boldsymbol{\mathcal{C}}(\omega)$, at the angular frequency ω.
Hence

$$\boldsymbol{X}(\omega) = \left(j\omega \boldsymbol{I} + \frac{1}{\tau_G} F_e(\omega) \boldsymbol{\mathcal{L}}(\omega)\right)^{-1} H_{local}(\omega) \boldsymbol{P}(\omega) \qquad (9)$$

In order to solve this we invoke the eigen-decomposition of $\boldsymbol{\mathcal{L}}(\omega)$, and that $\boldsymbol{U}(\omega)\boldsymbol{U}(\omega)^{\boldsymbol{H}} = \boldsymbol{I}$. It can then be shown easily that

$$\boldsymbol{X}(\omega) = \sum_i \frac{\boldsymbol{u_i}(\omega) \boldsymbol{u_i^H}(\omega)}{j\omega + \frac{1}{\tau_G} \lambda_i(\omega) F_e(\omega)} H_{local}(\omega) \boldsymbol{P}(\omega) \qquad (10)$$

This then is the steady state frequency response of the whole brain dynamics. In steady state, we assume that each cortical region is driven by internal noise that spans all frequencies, i.e. white noise. Hence we assume that the driving function $\boldsymbol{p}(t)$ is an uncorrelared Gaussian noise process, such that $\boldsymbol{P}(\omega) = \mathbb{1}$, where $\mathbb{1}$ is a vector of ones. This asserts identical cortical responses at each brain region.

Equation (10) encapsulates the entire model, and it is deterministic and admits a closed form solution, once the graph Laplacian eigenspectrum is known. There are very few model parameters, seven in total: a, $\tau_e, \tau_i, \tau_G, v, g_{ii}, g_{ei}$, which are all global and apply at every node. Note that the entire model is based on a single equation of graph dynamics, Eq (1), which is repeatedly applied to each level of the hierarchy. Here we used two levels: a mesoscopic level where connectivity is all-to-all, and a macroscopic level, where

connectivity is measured from fiber architecture. In theory, this template could be refined into finer levels, where neural responses become increasingly non-linear, and connectivity becomes sparser and structured.

Experimental Procedures and Model Comparisons

Alternative benchmark model for comparison. In order to put the proposed model in context, we also implemented for comparison a Wilson-Cowan neural mass model [17,34,36,49] with similar dimensionality. Although NMMs like this can and have been implemented with regionally varying local parameters, here we enforced uniform, regionally non-varying local parameters, meaning all parcellated brain regions shared the same local and global parameters. This is a fair comparison since the proposed model is also regionally non-varying. The purpose of this exercise is to ascertain whether a non-regional NMM can also predict spatial power variations purely as a consequence of network transmission, like the proposed model, using the same model optimization procedure (see below). This NMM incorporates a transmission velocity parameter that introduces a delay based on fiber tract lengths extracted from diffusion MRI, but, unlike our model, does not seek to explicitly evaluate a frequency response based on these delays.

Study cohort. We acquired MEG, anatomical MRI, and diffusion MRI for 36 healthy adult subjects (23 Male, 13 female; 26 left-handed, 10 right-handed; mean age 21.75 years (range: 7-51 years). All study procedures were approved by the institutional review board at the University of California at San Francisco (UCSF) and are in accordance with the ethics standards of the Helsinki Declaration of 1975 as revised in 2008.

MRI. A 3 Tesla TIM Trio MR scanner (Siemens, Erlangen, Germany) was used to perform MRI using a 32-channel phased-array radiofrequency head coil. High-resolution MRI of each subject's brain was collected using an axial 3D magnetization prepared rapid-acquisition gradient-echo (MPRAGE) T1-weighted sequence (echo time [TE] = 1.64 ms, repetition time [TR] = 2530 ms, TI = 1200 ms, flip angle of 7 degrees) with a 256-mm field of view (FOV), and 160 1.0-mm contiguous partitions at a 256×256 matrix. Whole-brain diffusion weighted images were collected at b = $1000 s/mm^2$ with 30 directions using 2-mm voxel resolution in-plane and through-plane.

Magneto-encephalography (MEG) data. MEG recordings were acquired at UCSF using a 275-channel CTF Omega 2000 whole-head MEG system from VSM MedTech (Coquitlam, BC, Canada). All subjects were instructed to keep their eyes closed for five minutes while their MEGs were recorded at a sampling frequency of 1200 Hz.

Data Processing

Region parcellation. The T1-weighted images were parcellated into 68 cortical regions and 18 subcortical regions using the using the Desikan-Killiany atlas available in the FreeSurfer software [58]. To do this, the subject specific T1-weighted images were back-projected to the atlas using affine registration, as described in our previous studies [18,59].

Structural Connectivity Networks. We constructed different structural connectivity networks with the same Desikan-Killiany parcellations to access the capabilities of our proposed model. Firstly, we obtained openly available diffusion MRI data from the MGH-USC Human Connectome Project to create an average template connectome. As in our previous studies [18,59], subject specific structural connectivity was computed on diffusion MRI data: *Bedpostx* was used to determine the orientation of brain fibers in conjunction with *FLIRT*, as implemented in the *FSL* software [60]. In order to determine the elements of the adjacency matrix, we performed tractography using *probtrackx2*. We initiated 4000 streamlines from each seed voxel corresponding to a cortical or subcortical gray matter structure and tracked how many of these streamlines reached a target gray matter structure. The weighted connection between the two structures $c_{i,j}$, was

defined as the number of streamlines initiated by voxels in region i that reach any voxel within region j, normalized by the sum of the source and target region volumes ($c_{i,j} = \frac{streamlines}{v_i + v_j}$). This normalization prevents large brain regions from having high connectivity simply due to having initiated or received many streamlines. Afterwards, connection strengths are averaged between both directions ($c_{i,j}$ and $c_{j,i}$) to form undirected edges. It is common in neuroimaging literature to threshold connectivity to remove weakly connected edges, as this can greatly influence the implied topology of the graph. In our work, we chose not to apply further thresholding, as unlike conventional graph theoretic metrics, linear models of spread and consequently network eigenmodes are relatively insensitive to implied topology induced by presence (or lack) of weak nonzero connections. However, to determine the geographic location of an edge, the top 95% of non-zero voxels by streamline count were computed for both edge directions. The consensus edge was defined as the union between both post-threshold sets.

MEG processing and source reconstruction. MEG recordings were down-sampled from 1200 Hz to 600 Hz, then digitally filtered to remove DC offset and any other noisy artifact outside of the 1 to 160 Hz bandpass range. Since MEG data are in sensor space, meaning they represent the signal observable from electrodes placed on the scalp, this data needs to be "inverted" in order to infer the neuronal activity that has generated the observed signal by solving the so-called inverse problem. Several effective methods exist for performing *source localization* [61–63]. Here we eschew the common technique of solving for a small number of discrete dipole sources which is not fully appropriate in the context of inferring resting state activity, since the latter is neither spatially sparse not localized. Instead, we used adaptive spatial filtering algorithms from the NUTMEG software tool written in house [64] in MATLAB (The MathWorks, Inc., Natick, Massachusetts, United States). To prepare for source localization, all MEG sensor locations were co-registered to each subject's anatomical MRI scans. The lead field (forward model) for each subject was calculated in NUTMEG using a multiple local-spheres head model (three-orientation lead field) and an 8 mm voxel grid which generated more than 5000 dipole sources, all sources were normalized to have a norm of 1. Finally, the MEG recordings were projected into source space using a beamformer spatial filter. Source estimates tend to have a bias towards superficial currents and the estimates are more error-prone when we approach subcortical regions, therefore, only the sources belonging to the 68 cortical regions were selected to be averaged around the centroid. Specifically, all dipole sources were labeled based on the Desikan-Killiany parcellations, then sources within a 20 mm radial distance to the centroid of each brain region were extracted, the average time course of each region's extracted sources served as empirical resting-state data for our proposed model.

Model Optimization using Simulated Annealing

We defined goodness of fit (GOF) as the Pearson correlation between empirical source localized MEG power spectra and simulated model power spectra, averaged over all 68 regions of a subject's brain. The proposed model has very few (seven) global parameters compared to neural mass models, and is available in closed-form. However GOF is still non-convex, with multiple possible local minima. To improve the odds that we capture the global minimum, we chose to implement the probabilistic approach of simulated annealing [65]. The algorithm samples a set of parameters within a set of boundaries by generating an initial trial solution and choosing the next solution from the current point by a probability distribution with a scale depending on the current "temperature" parameter. While the algorithm always accepts new trial points that map to cost-function values lower than the previous cost-function evaluations, it will also accept solutions that have cost-function evaluations greater than the previous one to move out of local minima. The acceptance probability function is $1/(1 + \frac{\Delta}{e^{\max(T)}})$, where T is the current temperature and Δ is the difference of the new minus old cost-function evaluations.

The initial parameter values and boundary constraints for each parameter are given in Table 1:

	Initial Value	**Lower/Upper Boundary**
Time constants $\{\tau_e, \tau_i, \tau_G\}$	12, 3, and 6 ms, respectively	[5ms, 20ms]
Gains $\{g_{ei}, g_{ii}\}$	4 and 1 respectively	[0.5, 5]
Transmission velocity v	5 m/s	[5 m/s, 20 m/s]
Alpha α	1	[0.1, 1]

All simulated annealing runs were allowed to iterate over the parameter space for a maximum of $N_p \times 3000$ iterations, where N_p is the number of parameters in the model.

As a comparison, we performed the same optimization procedure to a regionally non-varying Wilson-Cowan neural mass model [34,49]. We have recently reported a similar simulated annealing optimization procedure on this model [36].

Acknowledgements

This work was supported by NIH grants R01EB022717, R01DC013979, R01NS100440, R01DC017696, and UCOP-MRP-17-454755. The template HCP connectome used in the preparation of this work were obtained from the MGH-USC Human Connectome Project (HCP) database (https://ida.loni.usc.edu/login.jsp). The HCP project is supported by the National Institute of Dental and Craniofacial Research (NIDCR), the National Institute of Mental Health (NIMH) and the National Institute of Neurological Disorders and Stroke (NINDS). Collectively, the HCP is the result of efforts of co-investigators from the University of Southern California, Martinos Center at Massachusetts General Hospital (MGH), Washington University, and the University of Minnesota.

Author contributions:
Conceived the study: AR, PM, SN
Wrote model algorithm: AR
Experimental data acquisition and analysis: CC, XX, EP, JO
Software implementation: AR, CC, XX
Wrote manuscript: AR, XX, SN
Edited / improved manuscript: AR, CC, XX, EP, JO, PM, SN

Competing Interests statement: All authors declare that they do not have a competing interest as it pertains to this study.

REFERENCES

1. Hagmann, P. *et al.* Mapping the structural core of human cerebral cortex. *PLoS Biol.* **6,** e159 (2008).
2. Iturria-Medina, Y. Anatomical brain networks on the prediction of abnormal brain States. *Brain Connect.* **3,** 1–21 (2013).
3. Honey, C. J. *et al.* Predicting human resting-state functional connectivity from structural connectivity. *Proc. Natl. Acad. Sci. U. S. A.* **106,** 2035–40 (2009).
4. Freeman, W. J. & Zhai, J. Simulated power spectral density (PSD) of background electrocorticogram (ECoG). *Cogn. Neurodyn.* **3,** 97–103 (2008).
5. Liley, D. T., Alexander, D. M., Wright, J. J. & Aldous, M. D. Alpha rhythm emerges from large-scale networks of realistically coupled multicompartmental model cortical neurons. *Network* **10,** 79–92

(1999).
6. Roland, P. E., Hilgetag, C. C. & Deco, G. Tracing evolution of spatio-temporal dynamics of the cerebral cortex: cortico-cortical communication dynamics. *Front. Syst. Neurosci.* **8,** 76 (2014).
7. Markounikau, V., Igel, C., Grinvald, A. & Jancke, D. A dynamic neural field model of mesoscopic cortical activity captured with voltage-sensitive dye imaging. *PLoS Comput. Biol.* **6,** (2010).
8. Schaffer, E. S., Ostojic, S. & Abbott, L. F. A complex-valued firing-rate model that approximates the dynamics of spiking networks. *PLoS Comput. Biol.* **9,** e1003301 (2013).
9. Markram, H. The blue brain project. *Nat. Rev. Neurosci.* **7,** 153–160
10. Markram, H. *et al.* Reconstruction and Simulation of Neocortical Microcircuitry. *Cell* **163,** 456–492 (2015).
11. Izhikevich, E. M. & Edelman, G. M. Large-scale model of mammalian thalamocortical systems. *Proc. Natl. Acad. Sci.* **105,** 3593–3598 (2008).
12. Potjans, T. C. & Diesmann, M. The Cell-Type Specific Cortical Microcircuit: Relating Structure and Activity in a Full-Scale Spiking Network Model. *Cereb. Cortex* **24,** 785–806 (2014).
13. Shimizu, H. & Haken, H. Co-operative dynamics in organelles. *J. Theor. Biol.* **104,** 261–73 (1983).
14. Mišić, B., Sporns, O. & McIntosh, A. R. Communication efficiency and congestion of signal traffic in large-scale brain networks. *PLoS Comput. Biol.* **10,** e1003427 (2014).
15. Mišić, B. *et al.* Cooperative and Competitive Spreading Dynamics on the Human Connectome. *Neuron* **86,** 1518–1529 (2015).
16. Robinson, P. A., Rennie, C. J., Rowe, D. L., O'Connor, S. C. & Gordon, E. Multiscale brain modelling. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **360,** 1043–50 (2005).
17. Destexhe, A. & Sejnowski, T. J. The Wilson-Cowan model, 36 years later. *Biol. Cybern.* **101,** 1–2 (2009).
18. Abdelnour, F., Voss, H. U. & Raj, A. Network diffusion accurately models the relationship between structural and functional brain connectivity networks. *Neuroimage* **90,** 335–47 (2014).
19. Abdelnour F, Raj A, Dayan M, Devinsky O, T. T. Estimating function from structure in epileptics using graph diffusion model. in *Proceedings of the IEEE International Symposium on Biomedical Imaging* Paper ID 528 (2015).
20. Nakagawa, T. T. *et al.* How delays matter in an oscillatory whole-brain spiking-neuron network model for MEG alpha-rhythms at rest. *Neuroimage* **87,** 383–94 (2014).
21. Jirsa, V. K., Jantzen, K. J., Fuchs, A. & Kelso, J. A. S. Spatiotemporal forward solution of the EEG and MEG using network modeling. *IEEE Trans. Med. Imaging* **21,** 493–504 (2002).
22. Deco, G., Senden, M. & Jirsa, V. How anatomy shapes dynamics: a semi-analytical study of the brain at rest by a simple spin model. *Front. Comput. Neurosci.* **6,** 68 (2012).
23. Achard, S., Salvador, R., Whitcher, B., Suckling, J. & Bullmore, E. A resilient, low-frequency, small-world human brain functional network with highly connected association cortical hubs. *J. Neurosci.* **26,** 63–72 (2006).
24. Bullmore, E., Bullmore, E., Sporns, O. & Sporns, O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* **10,** 186–198 (2009).
25. Strogatz, S. H. Exploring complex networks. *Nature* **410,** 268–76 (2001).
26. van den Heuvel, M. P., Mandl, R. C. W., Kahn, R. S. & Hulshoff Pol, H. E. Functionally linked resting-state networks reflect the underlying structural connectivity architecture of the human brain. *Hum. Brain Mapp.* **30,** 3127–3141 (2009).
27. Hermundstad, A. M. *et al.* Structural foundations of resting-state and task-based functional connectivity in the human brain. *Proc. Natl. Acad. Sci.* **110,** 6169–6174 (2013).
28. Rubinov, M., Sporns, O., van Leeuwen, C. & Breakspear, M. Symbiotic relationship between brain structure and dynamics. *BMC Neurosci.* **10,** 55 (2009).
29. Ghosh, A., Rho, Y., McIntosh, A. R., Kötter, R. & Jirsa, V. K. Cortical network dynamics with time delays reveals functional connectivity in the resting brain. *Cogn. Neurodyn.* **2,** 115–120 (2008).
30. Abdelnour, F., Dayan, M., Devinsky, O., Thesen, T. & Raj, A. Functional brain connectivity is

predictable from anatomic network's Laplacian eigen-structure. *Neuroimage* **172,** 728–739 (2018).
31. Park, H.-J. & Friston, K. Structural and Functional Brain Networks: From Connections to Cognition. *Science (80-.).* **342,** 1238411–1238411 (2013).
32. He, B. J., Zempel, J. M., Snyder, A. Z. & Raichle, M. E. The temporal structures and functional significance of scale-free brain activity. *Neuron* **66,** 353–69 (2010).
33. El Boustani, S. & Destexhe, A. A master equation formalism for macroscopic modeling of asynchronous irregular activity states. *Neural Comput.* **21,** 46–100 (2009).
34. Wilson, H. R. & Cowan, J. D. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik* **13,** 55–80 (1973).
35. David, O. & Friston, K. J. A neural mass model for MEG/EEG: coupling and neuronal dynamics. *Neuroimage* **20,** 1743–55 (2003).
36. Xie, X. *et al.* Identifiability in connectome based neural mass models. *bioRxiv* 480012 (2018). doi:10.1101/480012
37. Larsen, R., Nielsen, M., Sporring, J., Zhang, F. & Hancock, E. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*. **4191,** (Springer Berlin Heidelberg, 2006).
38. Kondor & Lafferty. Diffusion kernels on graphs and other discrete structures. (2002).
39. Auffarth, B. Spectral graph clustering. *Univ. Barcelona, course Rep.* 1–12 (2007).
40. A.Ng, M. Jordan, Y. W. On Spectral Clustering: Analysis and an algorithm. *Adv. Neural Inf. Process. Syst.* 849–856 (2002).
41. Galán, R. F. On How Network Architecture Determines the Dominant Patterns of Spontaneous Neural Activity. *PLoS One* **3,** e2148 (2008).
42. Goni, J. *et al.* Resting-brain functional connectivity predicted by analytic measures of network communication. *Proc. Natl. Acad. Sci.* **111,** 833–838 (2014).
43. Wang, M. B., Owen, J. P., Mukherjee, P. & Raj, A. Brain network eigenmodes provide a robust and compact representation of the structural connectome in health and disease. *PLOS Comput. Biol.* **13,** e1005550 (2017).
44. Abdelnour, F., Dayan, M., Devinsky, O., Thesen, T. & Raj, A. Estimating brain's functional graph from the structural graph's Laplacian. in *SPIE Optical Engineering + Applications* (eds. Papadakis, M., Goyal, V. K. & Van De Ville, D.) 95970N (International Society for Optics and Photonics, 2015).
45. Atasoy, S., Donnelly, I. & Pearson, J. Human brain networks function in connectome-specific harmonic waves. *Nat. Commun.* **7,** 10340 (2016).
46. Abdelnour, F., Mueller, S. & Raj, A. Relating Cortical Atrophy in Temporal Lobe Epilepsy with Graph Diffusion-Based Network Models. *PLoS Comput. Biol.* **11,** (2015).
47. Abdelnour, F., Raj, A., Devinsky, O. & Thesen, T. Network Analysis on Predicting Mean Diffusivity Change at Group Level in Temporal Lobe Epilepsy. *Brain Connect.* **6,** 607–620 (2016).
48. Raj, A., Kuceyeski, A. & Weiner, M. A network diffusion model of disease progression in dementia. *Neuron* **73,** 1204–15 (2012).
49. Muldoon, S. F. *et al.* Stimulation-Based Control of Dynamic Brain Networks. *PLOS Comput. Biol.* **12,** e1005076 (2016).
50. Nunez, P. L. & Srinivasan, R. A theoretical basis for standing and traveling brain waves measured with human EEG with implications for an integrated consciousness. *Clin. Neurophysiol.* **117,** 2424–35 (2006).
51. Vijayan, S., Ching, S., Purdon, P. L., Brown, E. N. & Kopell, N. J. Thalamocortical mechanisms for the anteriorization of α rhythms during propofol-induced unconsciousness. *J. Neurosci.* **33,** 11070–5 (2013).
52. van Rotterdam, A., Lopes da Silva, F. H., van den Ende, J., Viergever, M. A. & Hermans, A. J. A model of the spatial-temporal characteristics of the alpha rhythm. *Bull. Math. Biol.* **44,** 283–305 (1982).
53. Liley, D. T. J., Cadusch, P. J. & Dafilis, M. P. A spatially continuous mean field theory of electrocortical activity. *Netw. Comput. Neural Syst.* **13,** 67–113 (2002).
54. Spiegler, A. & Jirsa, V. Systematic approximations of neural fields through networks of neural masses

in the virtual brain. *Neuroimage* **83,** 704–25 (2013).
55. Nunez, P. L. A study of origins of the time dependencies of scalp EEG: i--theoretical basis. *IEEE Trans. Biomed. Eng.* **28,** 271–80 (1981).
56. Ferezou, I. *et al.* Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice. *Neuron* **56,** 907–23 (2007).
57. Polack, P.-O. & Contreras, D. Long-range parallel processing and local recurrent activity in the visual cortex of the mouse. *J. Neurosci.* **32,** 11120–31 (2012).
58. Fischl, B. *et al.* Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* **33,** 341–55 (2002).
59. Owen, J. P. *et al.* The structural connectome of the human brain in agenesis of the corpus callosum. *Neuroimage* **70,** 340–355 (2013).
60. Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W. & Smith, S. M. FSL. *Neuroimage* **62,** 782–90 (2012).
61. Wipf, D. P., Owen, J. P., Attias, H. T., Sekihara, K. & Nagarajan, S. S. Robust Bayesian estimation of the location, orientation, and time course of multiple correlated neural sources using MEG. *Neuroimage* **49,** 641–55 (2010).
62. Zumer, J. M., Attias, H. T., Sekihara, K. & Nagarajan, S. S. Probabilistic algorithms for MEG/EEG source reconstruction using temporal basis functions learned from data. *Neuroimage* **41,** 924–940 (2008).
63. Jerbi, K. *et al.* Localization of realistic cortical activity in MEG using current multipoles. *Neuroimage* **22,** 779–93 (2004).
64. Dalal, S. S. *et al.* NUTMEG: a neuromagnetic source reconstruction toolbox. *Neurol. Clin. Neurophysiol.* **2004,** 52 (2004).
65. Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. Optimization by simulated annealing. *Science* **220,** 671–80 (1983).

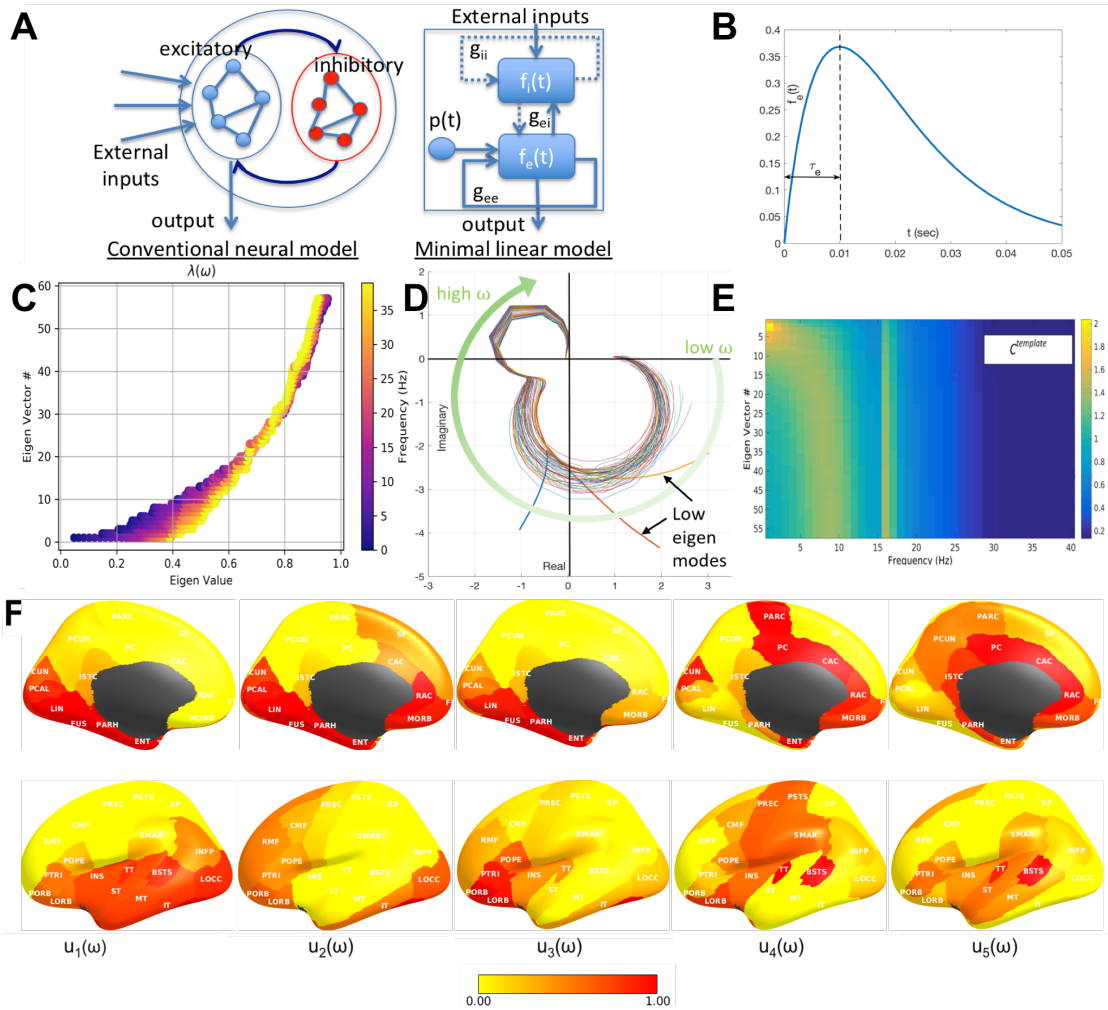

Figure 1: **Approximating a non-linear simulation model of local neural assemblies with a linear model**. **A:** Conventional neural models typically instantiate a large assembly of excitatory and inhibitory neurons, which are modeled as fully connected internally. External inputs and outputs are gated through the excitatory neurons only, and inhibitory neurons are considered strictly local. The proposed linear model condenses these local assemblies into lumped linear systems $f_e(t)$ and $f_i(t)$, Gamma-shaped functions having time constants τ_e and τ_i – see panel **B**. The recurrent architecture of the two pools within a local area is captured by the gain terms g_{ee}, g_{ii}, g_{ei}, indicating the loops created by recurrents within excitatory, inhibitory and cross-populations. **C:** The frequency-dependent behavior of the eigenvalue-spectrum of the complex Laplacian $\boldsymbol{\mathcal{L}}(\omega)$. Each dot represents the absolute value of each eigenvalue $\lambda(\omega)$, plotted against the eigenvector index; its color represents the frequency ω - low (blue) to high (yellow). Small eigenvalues show a larger shift due to frequency compared to large ones. ***D***: Frequency response of each eigenmode plotted on the complex plane with default model parameters. Each curve represents the transit in the complex plane of a single eigenmode's frequency response, starting at low frequencies in the bottom right quadrant, and moving characteristically to the upper left quadrant at high frequencies. The magnitude of the response, given by the distance from the origin, suggests that most eigenmodes have two prominent lobes, roughly corresponding to alpha and beta rhythms, respectively. In contrast, the lowest few eigenmodes start off far from the origin, indicative of a low-pass response. **E:** Magnitude of the frequency response of each eigenmode reinforces these impressions more clearly. **F**: The spatial patterns of the top 5 eigenmodes of $\boldsymbol{\mathcal{L}}(\omega)$, evaluated at the alpha frequency, 10 Hz. The first 4 eigenmodes $\mathbf{u}_1 - \mathbf{u}_4$, give strong alpha frequency responses, and in turn are strongly distributed spatially in posterior areas. Also see **Figure 5**. But they also include other regions and prominently resemble many elements of the **default mode network** and the **structural core** of the human connectome, especially $\mathbf{u}_5$. $\mathbf{u}_4$ resembles the **sensorimotor network**. These low eigenmodes are highly consistent and reproducible, but higher ones increasingly depend on the axonal velocity and frequency, and are not shown here.

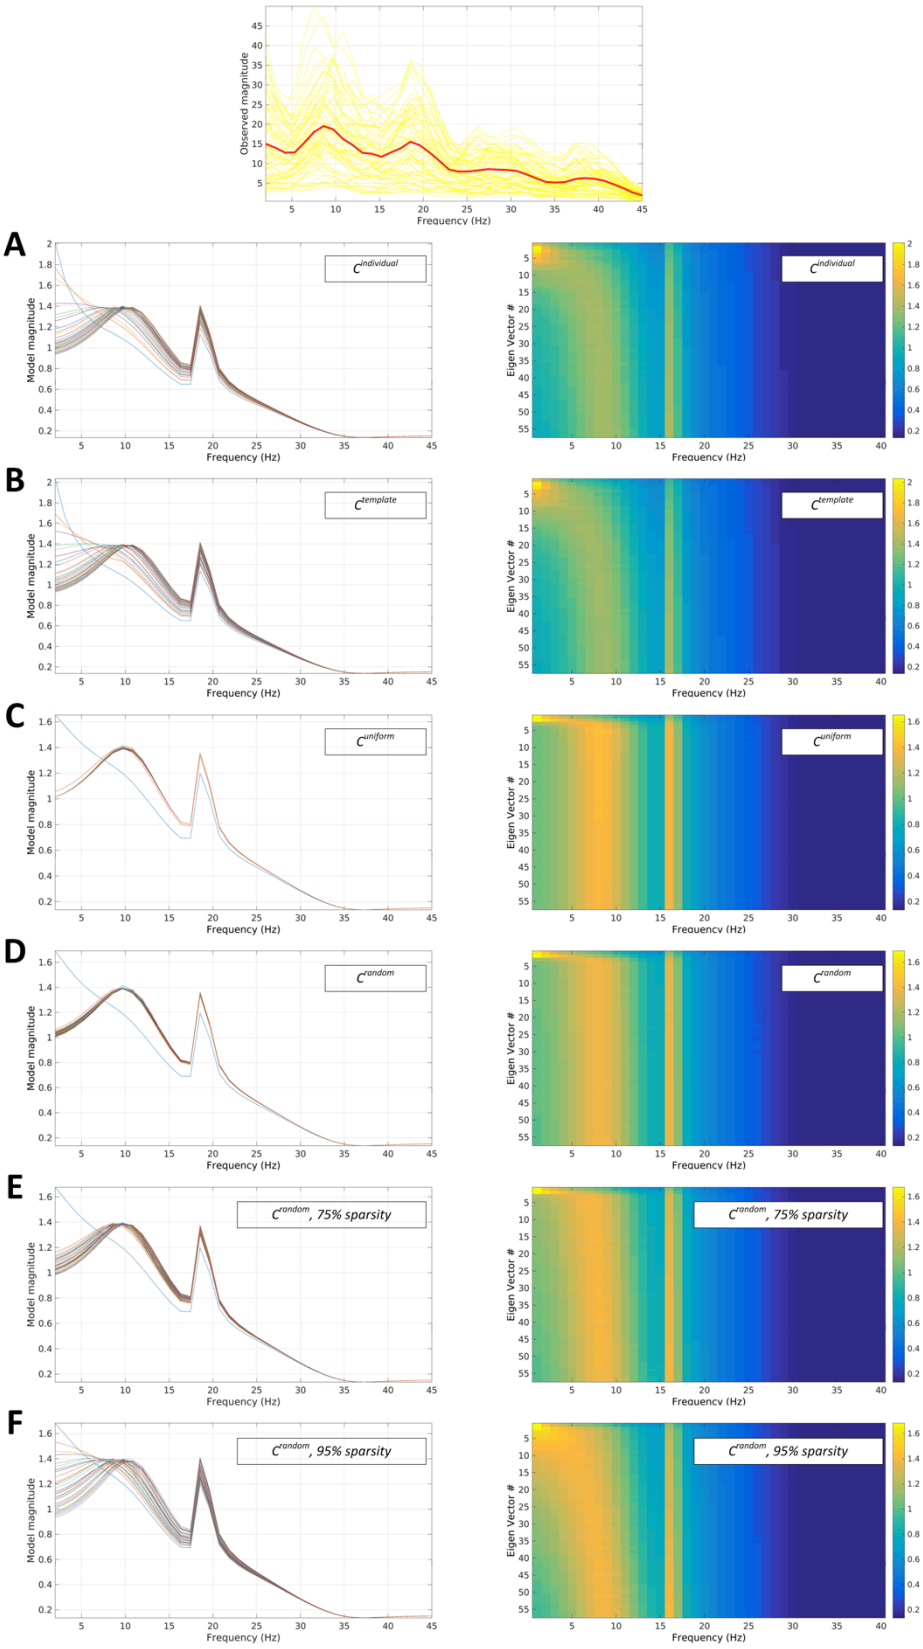

Figure 2: Spectral graph model predictions of MEG spectra for one representative subject. Top – Observed MEG power spectrum for each of the 68 parcellated brain regions. Average spectra for each brain region are shown in yellow, and the average spectrum across all brain regions is shown in red. The subsequent rows show each eigenmode's spectral magnitude response with model parameters optimized to match the observed spectrum ($\tau_e = 0.0073$, $\tau_i = 0.0085$, $\tau_G = 0.0061$, $g_{ei} = 2.9469$ $g_{ii} = 4.4865$, $\nu = 18.3071$ and $\alpha = 0.4639$). **A**: Model using

subject's individual structural connectivity matrix. **B**: Model using a template structural connectivity matrix obtained by averaging structural connectivity from 80 HCP subjects. **C**: Model using uniform connectivity matrix of ones. **D**: Model using randomized connectivity matrix with no sparsity. **E**: Model using randomized connectivity matrix with 75% sparsity. **F**: Model using randomized connectivity matrix with 95% sparsity. In all cases the connectome modulates the spectral response in delta-beta range, leaving the higher gamma frequencies unchanged. In general, mainly the low eigenmodes $(\mathbf{u_1} - \mathbf{u_{20}})$ appear to modulate the frequency response in any significant manner, and may be considered responsible for the diversity of spectra observed in the model.

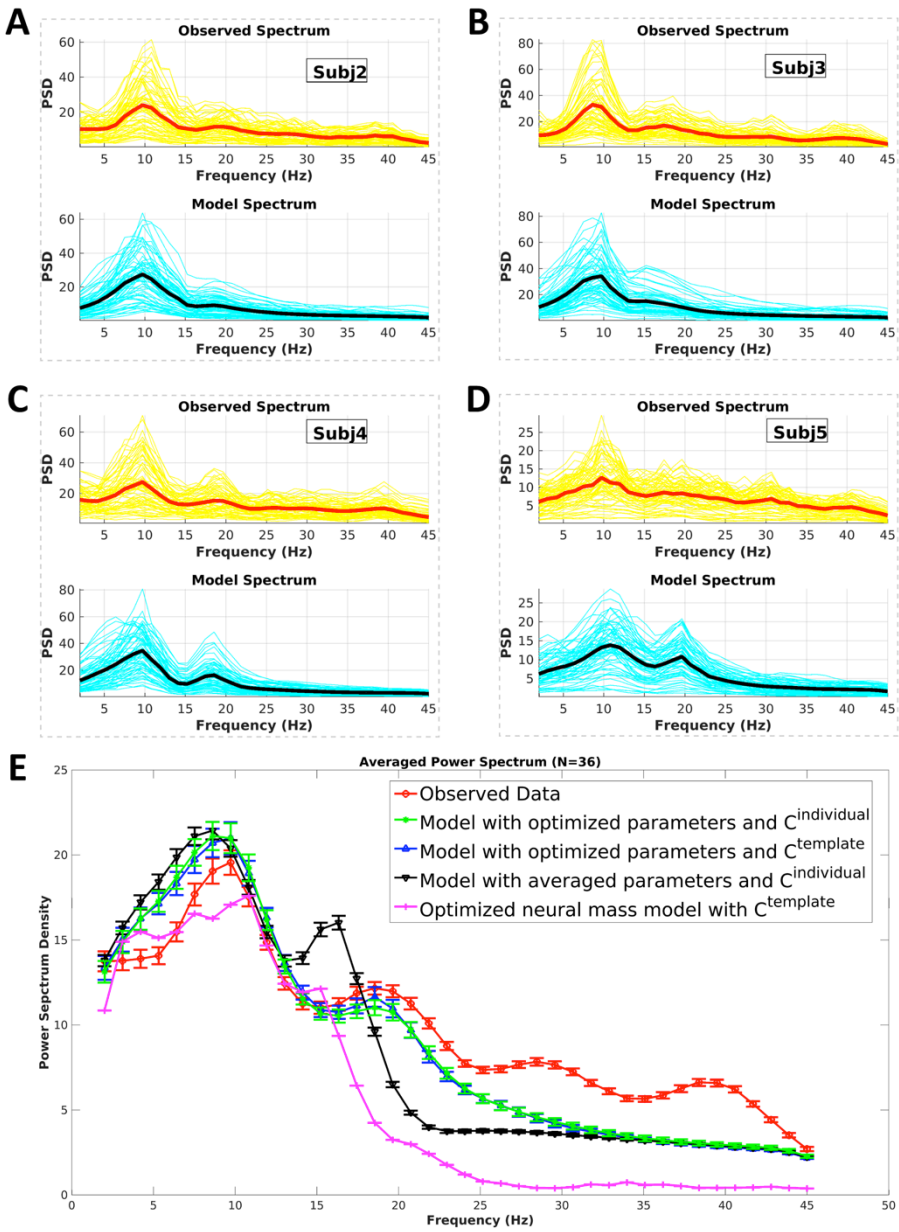

Figure 3: Spectral graph model correctly captures MEG spectra across subjects. A-D: The observed spectra and spectral graph model's simulated spectra for four representative subjects. Red and yellow curves illustrate source localized average spectra and region-wise spectra respectively, while black and cyan curves illustrate modeled average spectra and region-wise spectra respectively. **E:** The average observed spectrum for all 36 subjects is shown in red. The average simulated spectral graph model's power spectra produced by (green) optimized parameters with each individual subject's connectome, (blue) optimized parameters with

template connectome from HCP dataset, and (black) average optimized parameters applied to each individual subject's connectome are shown as comparisons. Additionally, the neural mass model's average simulated power spectra with each subject's optimized global parameters and HCP template connectome is shown in pink.

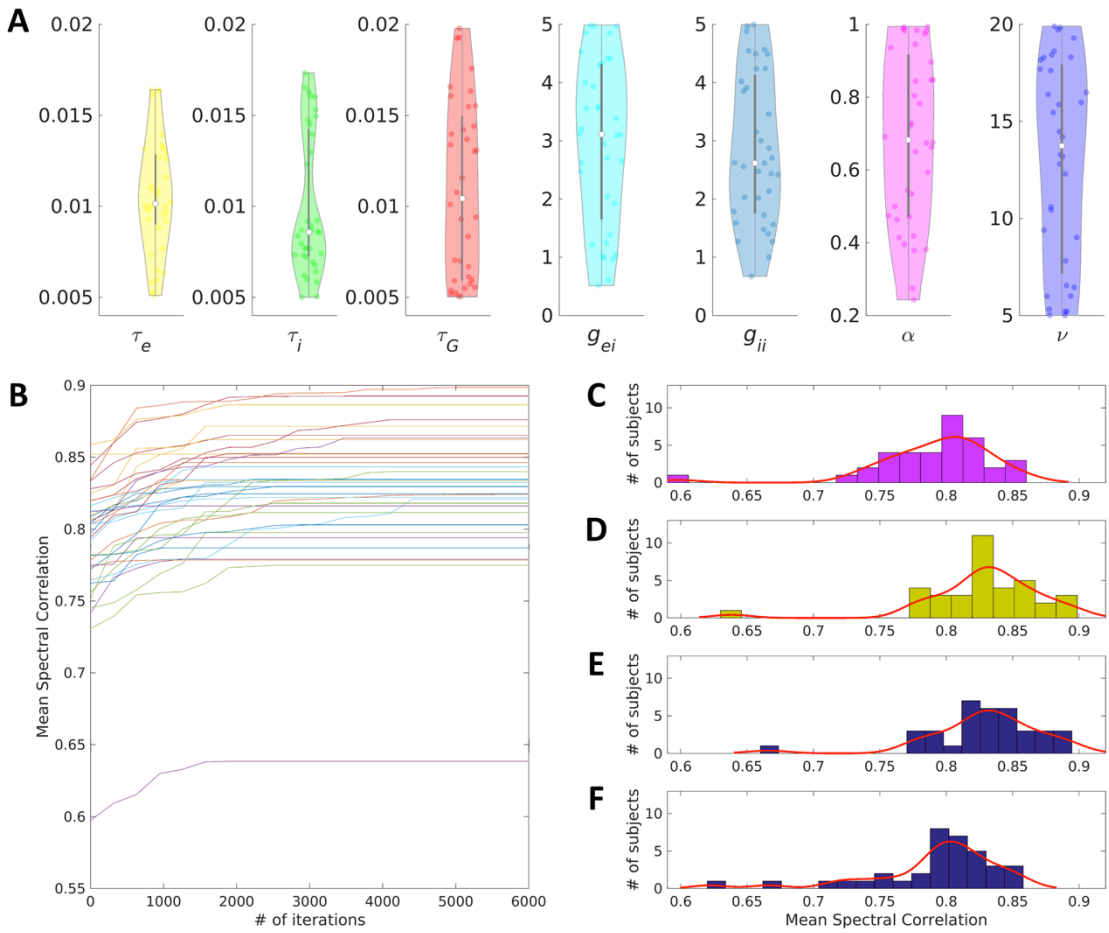

Figure 4: Model optimization via simulated annealing. **A**: **Distribution of optimized model parameter values for all 36 subjects.** The simulated annealing algorithm performed model optimization for the set of parameters $\{\tau_e,\ \tau_i,\ \tau_c,\ g_{ei},\ g_{ii},\ \alpha,\ \upsilon\}$ on all subjects. The optimized values for each parameter are shown in violin plots with each dot representing one subject. **Performance of optimization procedure**. **B**: Spectral Pearson correlation between model and source localized MEG spectra at each iteration of simulated annealing. Each curve shows the spectral correlation achieved by the model optimized for a single subject, averaged over all regions. Each accepted iteration increased the mean correlation values until the algorithm converged to a set of optimized parameters. Individually-optimized mean spectral correlation values for all 36 subjects with individual connectomes using grid search (**C**) gives inferior performance compared with simulated annealing (**D**). Performance of **D** is similar to that achieved using the template HCP connectome (**E**), but higher than using identical parameters for all subjects (**F**) – here the mean of all subjects' optimized parameters was used. In almost all subjects, the optimal spectral fit shown in **D** is highly significant (average $R = 0.83, p < 10^{-6}$).

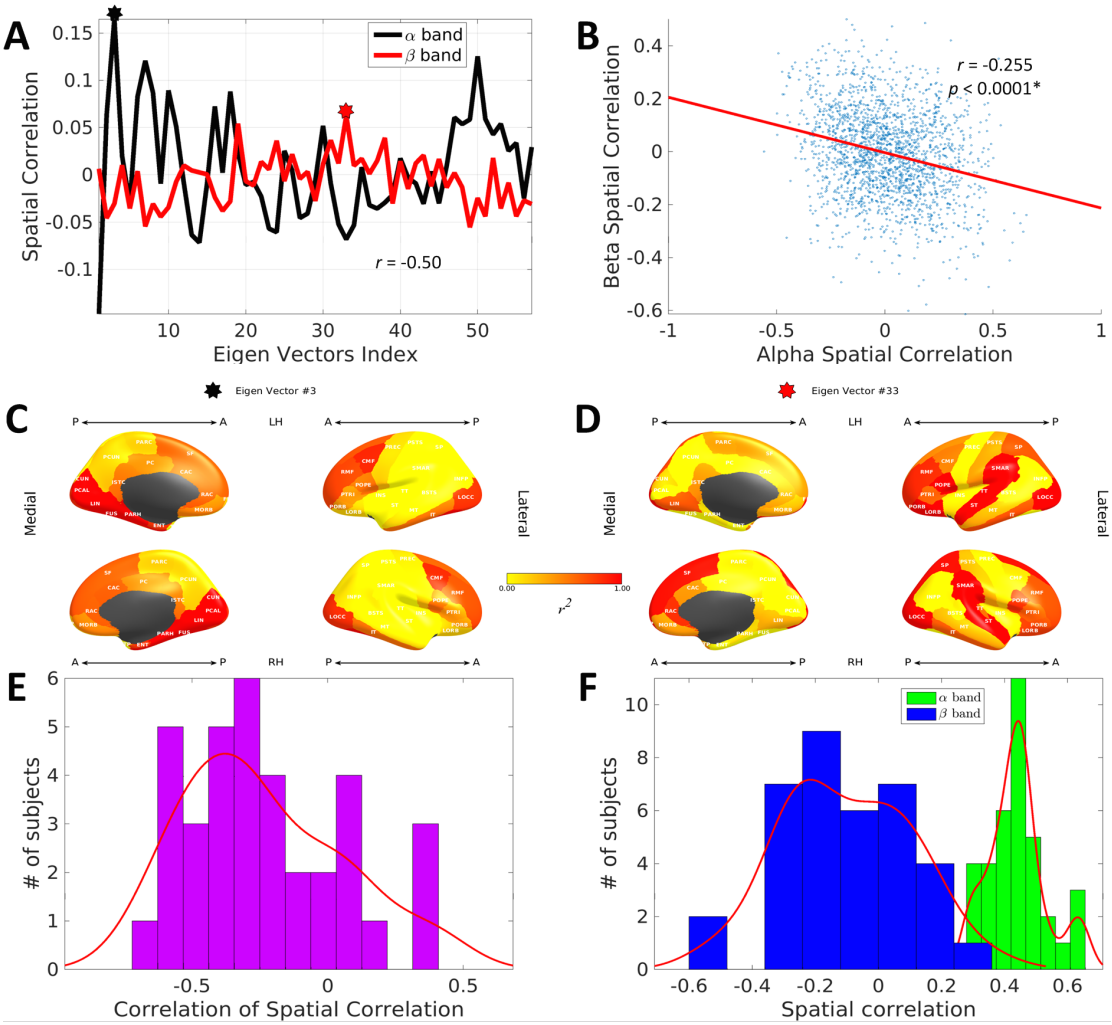

Figure 5: Specific eigenmodes capture spatial distributions of alpha and beta band activity. Here "spatial correlation" is defined as Pearson's R statistic of the correlation between the regional distribution of empirical MEG and model-predicted power within a given frequency band. **A**: Spatial correlation was computed for each eigenmode for all subjects in the alpha and beta bands, and the average spatial correlation for the eigenmodes are shown. **B**: When a given eigenmode is correlated to alpha power, it is roughly anti-correlated to beta power. To demonstrate this further, a scatter plot of all eigenmodes' alpha and beta power spatial correlation is shown. The r-statistic (r = -0.255, p<0.0001) of this correlation-of-correlations is highly significantly negative. **C** and **D** show the cortical surface renderings of the spatial pattern of the most spatially correlated eigenmode for alpha (#3) and beta (#33) respectively (indicated by * in panel **A**). **E:** Per subject, the histogram of the correlation between the alpha band and beta band spatial correlation curves. It can be seen that for most subjects alpha and beta band spatial correlation curves for the eigenmodes are in turn anti-correlated. **F**: Green histogram shows the distribution of spatial correlation achieved by the maximally correlated eigenmode with each subject's observed data in the alpha band. The blue histogram shows the spatial correlation achieved by the same eigenmodes with data in the beta band. This suggests that the eigenmode that is maximally correlated with alpha spatial pattern is in turn uncorrelated with the beta spatial pattern.

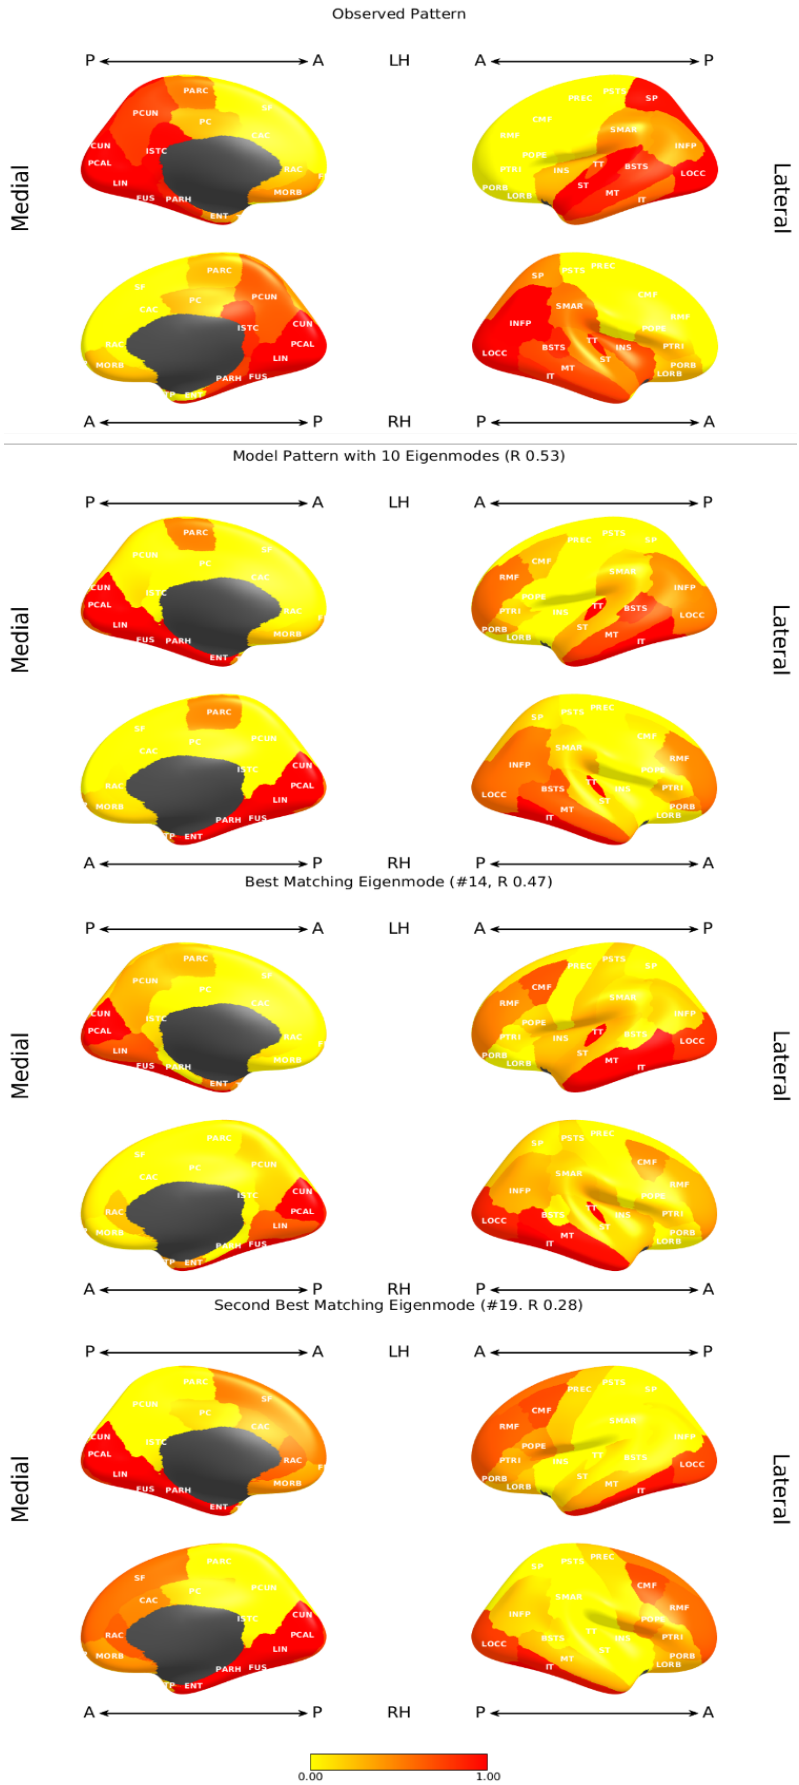

Figure 6: Spatial distribution of the alpha band power across regions. The spatial distribution of alpha band power is shown in cortical surface renderings, from top to bottom: the

observed MEG pattern; the model pattern with 10 eigenmodes; the best matching eigenmode generated by eigenmode 14 only; the second best matching eigenmode generated by eigenmode 19 only.

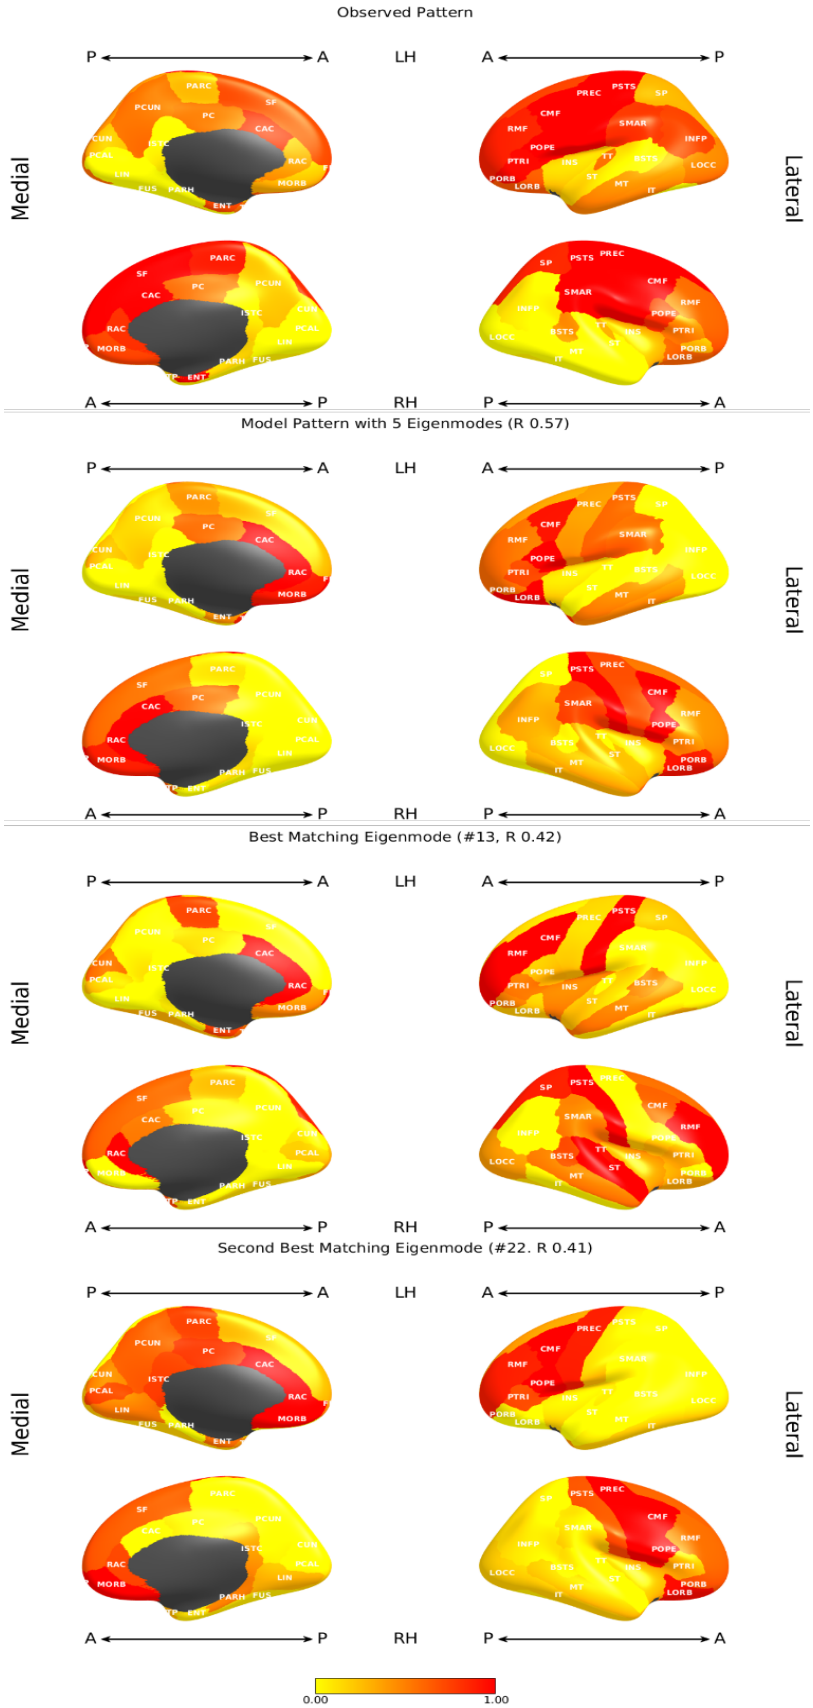

Figure 7: Spatial distribution of beta band power across regions. The spatial distribution of

beta band power are shown in cortical surface renderings, from top to bottom: the observed MEG pattern; the model pattern with 5 eigenmodes; the best matching eigenmode generated by eigenmode 13 only; the second best matching eigenmode generated by eigenmode 22 only.

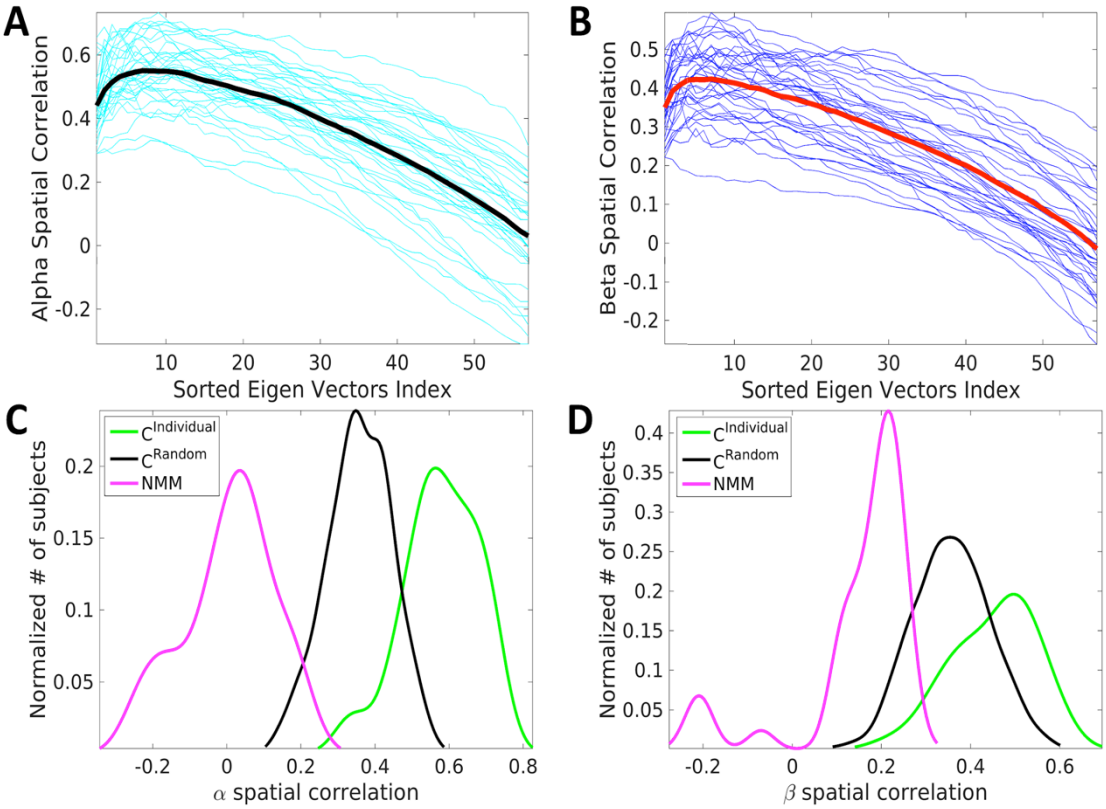

Figure 8: Spatial correlation between model and observed data for all subjects. The eigenmodes decomposed from individual connectomes are sorted by descending eigen values, and the spectral graph model's alpha and beta spatial correlation as these eigenmodes are cumulatively added together are shown in **A** and **B** respectively, the thin cyan and blue lines are subject specific spatial correlations, while the thick black and red line are average of all subjects. The distribution of simulated alpha and beta band maximum spatial correlation for all subjects are shown in **C** and **D**, the green curve is produced by the spectral graph model with individual connectomes and optimized parameters, the black curve is produced by the spectral graph model with 1000 instances of 80% sparse randomly generated connectomes and optimized parameters, and the pink curve is produced by the neural mass model with individual connectomes and optimized global parameters. The NMM gives very poor prediction of the spatial pattern of alpha power (r distribution centered at 0). The random connectomes also appear to have some ability to capture these spatial patterns (r distribution centered at 0.35). This may be understood as a consequence of the implicit search within the random eigenmodes of the best-matching ones, which on average will give at least a few eigenmodes that look like MEG power purely by chance. The model evolved on the human connectome does the best in all cases, but markedly better in alpha compared to beta band.

Supplementary figures:

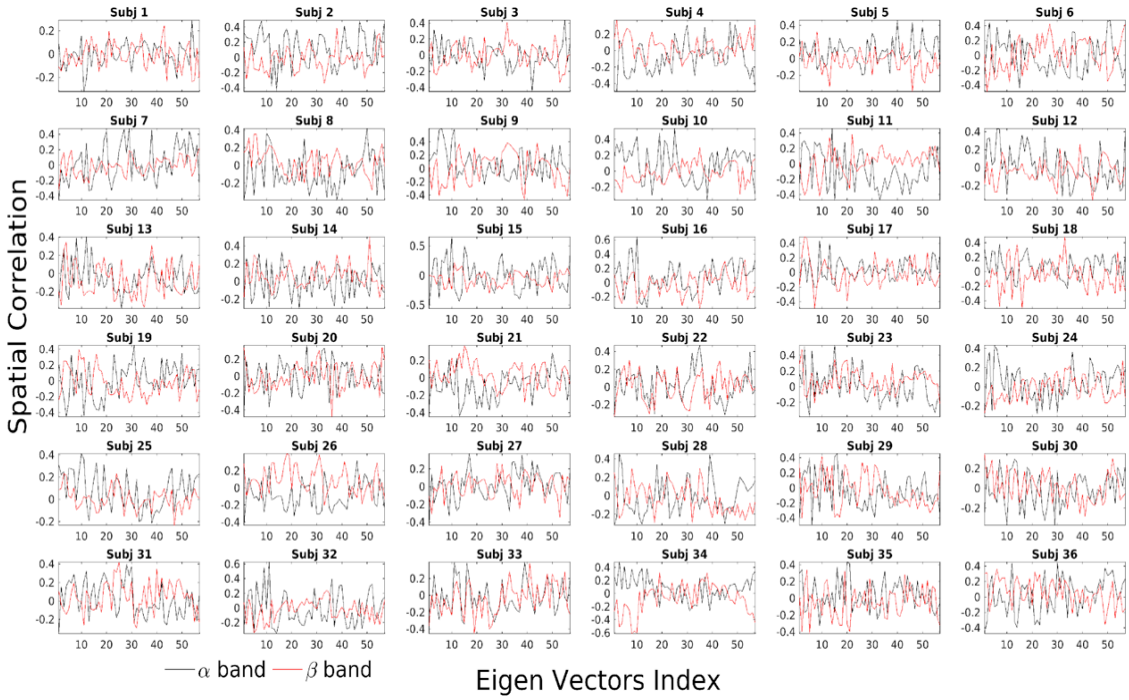

Supplementary Figure 1: Alpha and beta band spatial correlation for 36 subjects. Every subject's individual eigenmode's alpha and beta spatial correlation are shown together before sorting and cumulative summing.

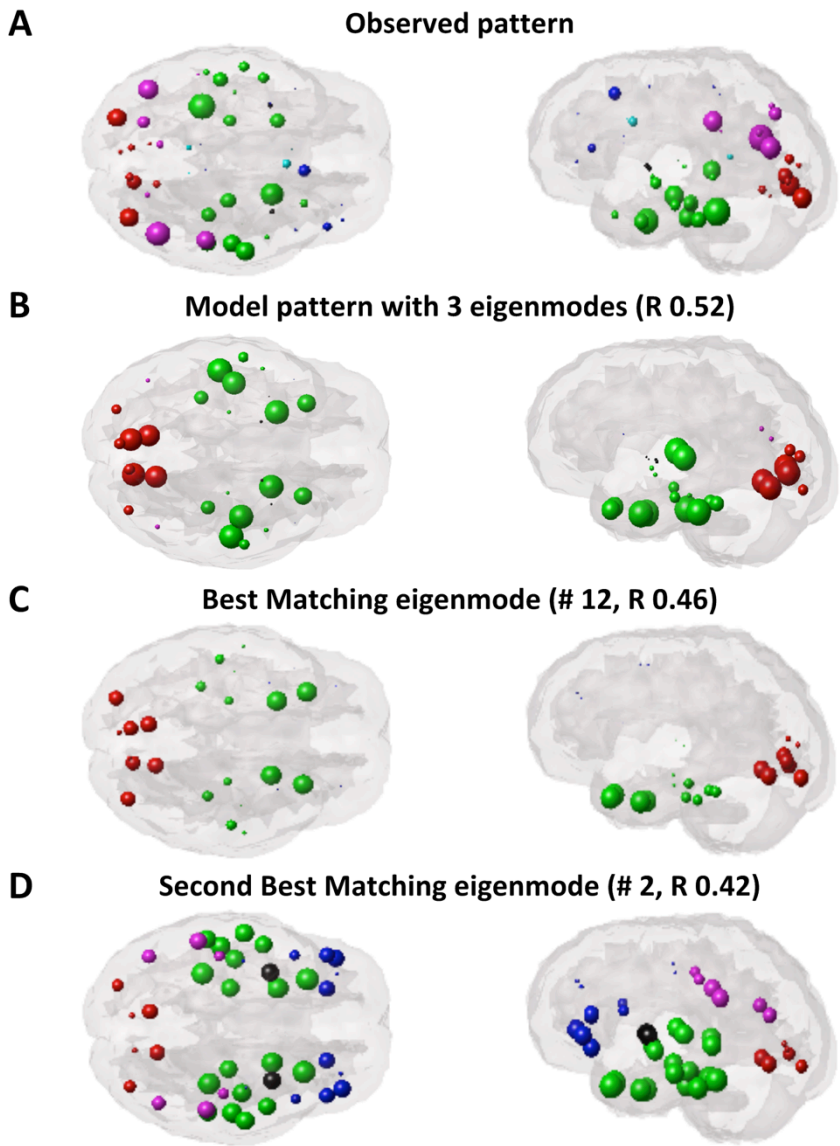

Supplementary Figure 2: Alpha spatial distribution for subject 2. The spatial distribution of alpha band power in a "glass brain" rendering, from top to bottom: **A**: the observed MEG pattern; **B**: the model pattern with 3 eigenmodes; **C**: the best matching eigenmode generated by eigenmode 12 only; **D**: the second best matching eigenmode generated by eigenmode 2 only.

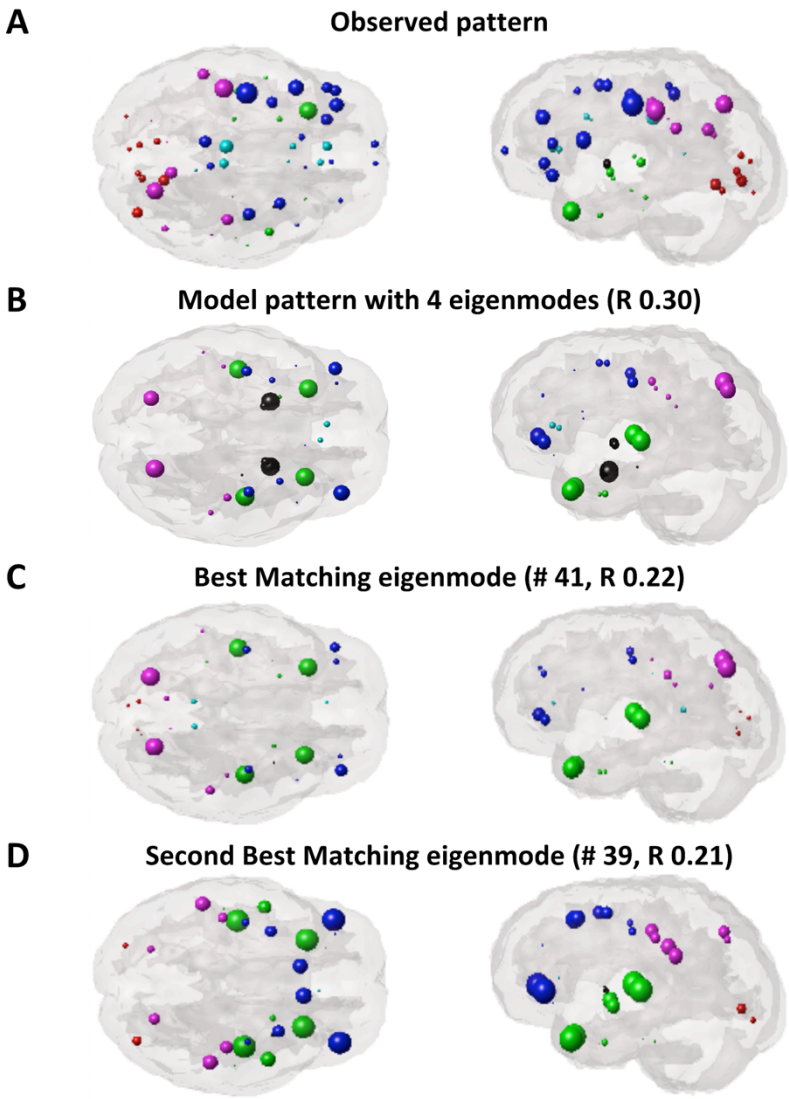

Supplementary Figure 3: Beta spatial distribution for subject 2. The spatial distribution of beta band power in a "glass brain" rendering, from top to bottom: **A**: the observed MEG pattern; **B**: the model pattern with 4 eigenmodes; **C**: the best matching eigenmode generated by eigenmode 41 only; **D**: the second best matching eigenmode generated by eigenmode 39 only.

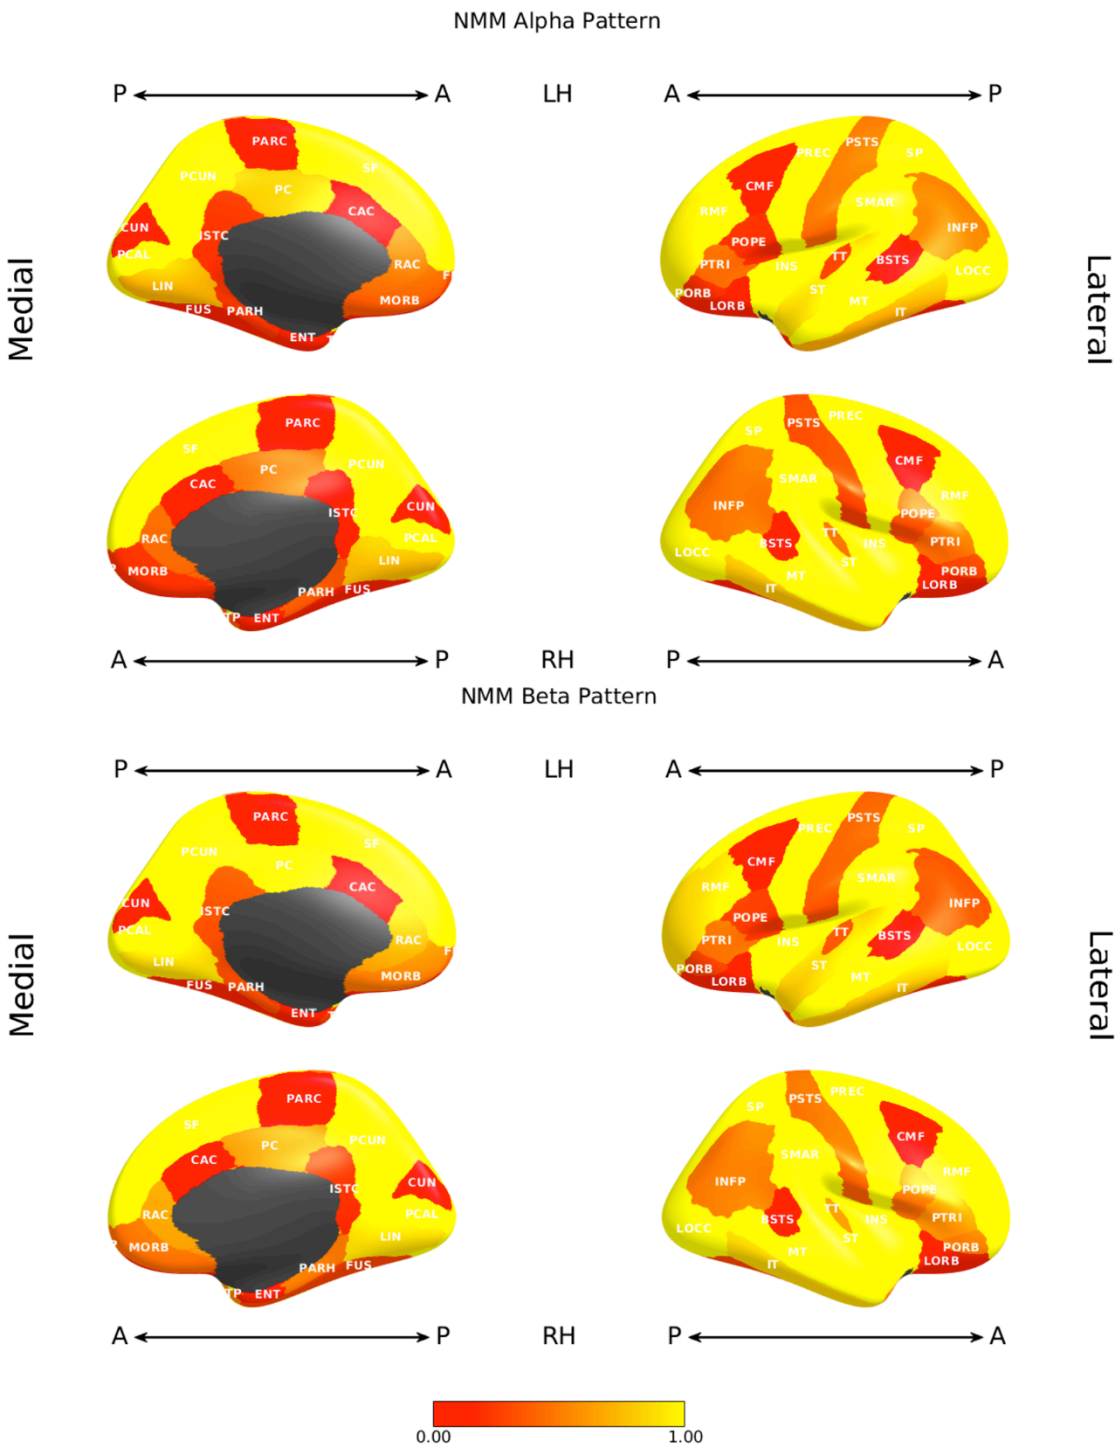

Supplementary Figure 4: Alpha and beta band spatial power distribution for subject 2 predicted by the Wilson-Cowan neural mass model. The spatial distributions are shown in cortical surface renderings: alpha power (top) and beta power (bottom). The empirical MEG spatial patterns of this subject were shown in **Figures 6** and **7** of the main Results section. Clearly, this NMM is unable to capture the spatial patterns of alpha and beta, using identical local model parameters. Better results were reported in numerous prior studies of NMMs, but those allow locally-varying model parameters fine-tuned to match local band power.