Graph neural fields: a framework for spatiotemporal dynamical models on the human connectome

Marco Aqil[1¤*], Selen Atasoy[2,3], Morten L. Kringelbach[2,3], Rikkert Hindriks[1],

1 Department of Mathematics, Vrije Universiteit, Amsterdam, The Netherlands
2 Centre for Eudaimonia and Human Flourishing, University of Oxford, Oxford, UK
3 Center for Music in the Brain, University of Aarhus, Aarhus, Denmark

¤Current Address: Spinoza Centre for Neuroimaging, Amsterdam, The Netherlands
* m.aqil@spinozacentre.nl

Abstract

Tools from the field of graph signal processing, in particular the graph Laplacian operator, have recently been successfully applied to the investigation of structure-function relationships in the human brain. The eigenvectors of the human connectome graph Laplacian, dubbed "connectome harmonics", have been shown to relate to the functionally relevant resting-state networks. Whole-brain modelling of brain activity combines structural connectivity with local dynamical models to provide insight into the large-scale functional organization of the human brain. In this study, we employ the graph Laplacian and its properties to define and implement a large class of neural activity models directly on the human connectome. These models, consisting of systems of stochastic integrodifferential equations on graphs, are dubbed *graph neural fields*, in analogy with the well-established continuous neural fields. We obtain analytic predictions for harmonic and temporal power spectra, as well as functional connectivity and coherence matrices, of graph neural fields, with a technique dubbed CHAOSS (shorthand for *Connectome-Harmonic Analysis Of Spatiotemporal Spectra*). Combining graph neural fields with appropriate observation models allows for estimating model parameters from experimental data as obtained from electroencephalography (EEG), magnetoencephalography (MEG), or functional magnetic resonance imaging (fMRI); as an example application, we study a stochastic Wilson-Cowan graph neural field model on a high-resolution connectome, and show that the model equilibrium fluctuations can reproduce the empirically observed harmonic power spectrum of BOLD fMRI data. Graph neural fields natively allow the inclusion of important features of cortical anatomy and fast computations of observable quantities for comparison with multimodal empirical data. They thus appear particularly suitable for modelling whole-brain activity at mesoscopic scales, and opening new potential avenues for connectome-graph-based investigations of structure-function relationships.

Author summary

The human brain can be seen as an interconnected network of many thousands neuronal "populations"; in turn, each population contains thousands of neurons, and each is connected both to its neighbors on the cortex, and crucially also to distant populations thanks to long-range white matter fibers. This extremely complex network, unique to each of us, is known as the "human connectome graph". In this work, we develop a novel approach to investigate how the neural activity that is necessary for our life and experience of the world arises from an individual human connectome graph. For the first time, we implement a mathematical model of neuronal activity directly on a high-resolution connectome graph, and show that it can reproduce the spatial patterns of activity observed in the real brain with magnetic resonance imaging. This new kind of model, made of equations implemented directly on connectome graphs, could help us better understand how brain function is shaped by computational principles and anatomy, but also how it is affected by pathology and lesions.

Introduction

The spatiotemporal dynamics of human resting-state brain activity is organized in functionally relevant ways, with perhaps the best-known example being the "resting-state networks" [1]. How the repertoire of resting-state brain activity arises from the underlying anatomical structure, i.e. the connectome, is a highly non-trivial question: it has been shown that structural connections imply functional ones, but that the converse is not necessarily true [2]; furthermore, specific discordant attributes of structural and functional connectivity have been found by network analyses [3]. Research on structure-function questions can be broadly divided into data-driven (analysis), theory-driven (modelling), and combinations thereof. In this work, we combine techniques from graph signal processing (analysis) and neural field equations (modelling) to outline a promising new approach for the investigation of whole-brain structure-function relationships.

A recent trend of particular interest in neuroimaging data analysis is the application of methods from the field of graph signal processing [4–8]. In these applications, anatomical information obtained from diffusion tensor imaging (DTI) and structural MRI is used to construct the *connectome graph* [9], and combined with functional imaging data such as BOLD-fMRI or EEG/MEG to investigate structure-function relationships in the human brain (see [10, 11] for reviews). The workhorse of graph signal processing analysis is the *graph Laplacian operator*, or simply graph Laplacian. Originally formulated as the graph-equivalent of the Laplace-Beltrami operator for Riemannian manifolds [12, 13], the graph Laplacian is now established as a valuable tool in its own right [10]. The eigenvectors of the graph Laplacian provide a generalization of the Fourier transform to graphs, and therefore also a complete orthogonal basis for functions on the graph. In the context of the human connectome graph, the eigenvectors of the graph Laplacian are referred to as *connectome harmonics* by analogy with the harmonic eigenfunctions of the Laplace-Beltrami operator. Of relevance to the current work, several connectome harmonics have been shown to be related to specific resting-state networks [9]. More recent studies have provided additional evidence for this claim [14], and others used a similar approach to explain how distinct electrophysiological resting-state networks emerge from the structural connectome graph [15]. Furthermore in [9], for the first time, to the best of our knowledge, a model of neural activity making use of the graph Laplacian was implemented, and used to suggest the role of excitatory-inhibitory dynamics as possible underlying mechanism for the self-organization of resting-state activity patterns. In other very recent work [16] graph-Laplacian-based techniques were employed to model MEG oscillations. Considering these developments, the combination of neural activity modelling and graph signal processing techniques appears as a promising direction for further inquiry.

Whole-brain models are models of neural activity that are defined on the entire cortex and possibly on subcortical structures. This is generally achieved either by parcellating the cortex into a network of a few dozens of macroscopic, coupled regions of interest (ROIs), or by approximating the cortex as a bidimensional manifold, and studying continuous integrodifferential equations in a flat or spherical geometry (See [17] for a review). In this study, relying on graph signal processing methods such as the graph Laplacian and graph filtering [6, 8], we show how to define and implement a large class of whole-brain models of neural activity on arbitrary metric graphs[1], and in particular on a non-parcellated mesoscopic human connectome. These models consist of systems of integrodifferential equations, and are dubbed *graph neural fields* by analogy with their continuous counterparts. We obtain analytic expressions for harmonic and temporal

[1] that is, graphs equipped with a suitable distance metric

power spectra, as well as functional connectivity and coherence matrices, of graph neural fields, with a technique dubbed CHAOSS (shorthand for *Connectome-Harmonic Analysis Of Spatiotemporal Spectra*). When combined with appropriate observation models, graph neural fields can be fitted to and compared with functional data obtained from different imaging modalities such as EEG/MEG, fMRI, and positron emission tomography (PET). As an example application, we study a Wilson-Cowan stochastic graph neural field model, implemented on a single-subject unparcellated connectome. We show that the model can accurately reproduce the empirical harmonic spectrum of resting-state BOLD signal fluctuations. In sum, graph neural fields provide a computationally efficient and versatile modelling framework that is tailored for connectome-graph-based structure-function investigations, and particularly suitable for modelling whole-brain activity on mesoscopic scales. Graph neural fields present immediate application in the investigation of the relationship between individual anatomy, pathology, and lesions with functional activity; and furthermore provide a model-based approach to test novel graph signal processing neuroimaging hypotheses and analyses.

In Section *Results* we start by providing an analytic solution and a numerical implementation of the damped-wave equation on the human connectome graph, since this equation is of interest in the context of modelling neural activity propagation. Next, we show how to implement the Wilson-Cowan stochastic graph neural field model on arbitrary metric graphs. We obtain results of linear stability analysis, CHAOSS, and numerical simulations, first on a one-dimensional graph with 1000 vertices, and then on a single-subject connectome consisting of approximately 18000 cortical vertices and 10000 white matter tracts. The simplified context of a 1-dimensional graph is useful to study the effect of simple graph properties, such as inter-node distance and long-range connections, on model dynamics; moving to a real-world application, we fit the full-connectome model parameters to the experimentally observed harmonic power spectrum of resting-state fMRI data of a single subject, showing excellent agreement between the analytically predicted, numerically simulated, and empirical power spectra. In Section *Methods* we provide a more general and detailed description of the framework of graph neural fields. We define spatiotemporal convolution on graphs through the weighted graph Laplacian, and obtain the graph-equivalents of several connectivity kernels and reaction-diffusion systems that are of interest for neural activity modelling. Finally, we show how to derive analytic expressions for harmonic and temporal power spectra, as well as coherence and functional connectivity matrices, of graph neural fields (CHAOSS).

Results

Damped wave equation on the human connectome graph

The damped-wave equation describes the dynamics of simultaneous diffusion and wave propagation, and is thus of interest in the context of modelling activity propagation in neural tissue [18]. Here, we solve the graph equivalent of the damped-wave equation and implement it on the human connectome graph. In one-dimensional continuous space, the damped-wave equation is

$$a\frac{\partial^2 u(x,t)}{\partial t^2} + b\frac{\partial u(x,t)}{\partial t} = \frac{\partial^2 u(x,t)}{\partial x^2}, \tag{1}$$

where a and b are scalar parameters. Its graph-equivalent is given by

$$D_t u(t) = \Delta u(t), \tag{2}$$

where $u(t)$ is a function on the graph, Δ is the distance-weighted graph Laplacian, and $D_t = ad^2/dt^2 + bd/dt$. Since the graph Laplacian is a constant matrix, we can straightforwardly obtain the exact solution at time t:

$$\hat{u}(t) = K(t)\hat{u}(0) + \tilde{K}(t)\dot{\hat{u}}(0), \tag{3}$$

where $(\hat{u}(0), \dot{\hat{u}}(0))$ are the initial conditions, the superscript $\hat{u}$ indicates the graph Fourier transform (see *Methods*), and

$$K(t) = \frac{r_1 e^{r_2 t} - r_2 e^{r_1 t}}{r_1 - r_2}, \qquad \tilde{K}(t) = \frac{e^{r_1 t} - e^{r_2 t}}{r_1 - r_2}, \qquad r_{1,2} = \frac{-b \pm \sqrt{b^2 + 4a\Lambda}}{2a}, \tag{4}$$

Λ is the diagonal matrix of graph Laplacian eigenvalues. Having obtained an exact solution, we can efficiently simulate the time-evolution of the damped-wave equation on arbitrary metric graphs, for example with the following numerical scheme:

$$\hat{u}(t+\delta t) = K(\delta t)\hat{u}(t) + \tilde{K}(\delta t)\frac{\hat{u}(t) - \hat{u}(t-\delta t)}{\delta t}, \tag{5}$$

We also note that the Telegrapher's Equation, which is of interest in the context of modelling action potentials [19]

$$a\frac{\partial^2 u(x,t)}{\partial t^2} + b\frac{\partial u(x,t)}{\partial t} + cu(x,t) = \frac{\partial^2 u(x,t)}{\partial x^2}, \tag{6}$$

can also be implemented on metric graphs simply by substituting Λ with $(\Lambda - \mathrm{Diag}(c))$ in Eq (4).

Fig 1 shows an example implementation of Eq (5) on the human connectome graph. The initial condition is a Gaussian centered in the occipital cortex of the left hemisphere,

itself obtained by applying the graph Gaussian convolution filter (Table 2) to an impulse function. Note how the initially localized activity propagates with a characteristic speed and wavelength. Damping of the wave, caused by the diffusion term, can also be observed. Fig 2 shows an implementation with different parameters, giving rise to faster wavefront propagation and less damping. Fig 3 shows the same implementation as Fig 2, from a different point of view and with a narrower color-scale. This is to emphasize that, aside from the dominant surface-based wavefront, activity also propagates non-locally along white-matter fibers.

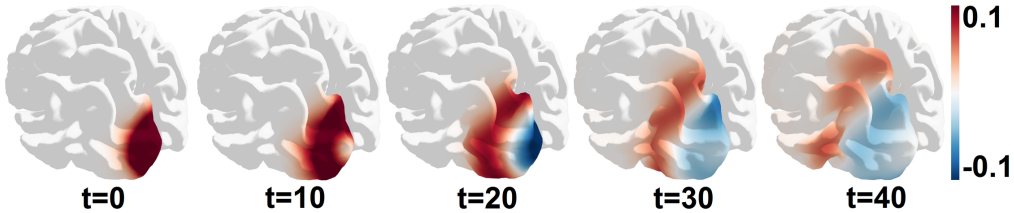

Fig 1. The damped-wave equation on the human connectome gives rise to propagation with characteristic speed and wavelength. Shown are snapshots of simulated cortical activity that is governed by the damped wave equation with time-step $\delta t = 1$ and parameters $a = 3 \cdot 10^5$, $b = 5 \cdot 10^3$.

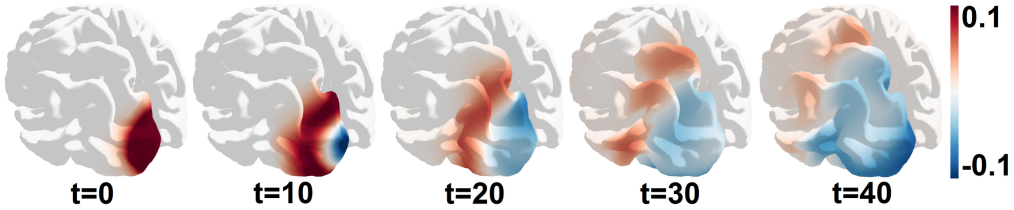

Fig 2. Varying the parameters of the damped-wave equation alters the dynamics of propagation on the human connectome. Shown are snapshots of simulated cortical activity that is governed by the damped wave equation with time-step $\delta t = 1$ and parameters $a = 1.5 \cdot 10^5$, $b = 2.5 \cdot 10^3$.

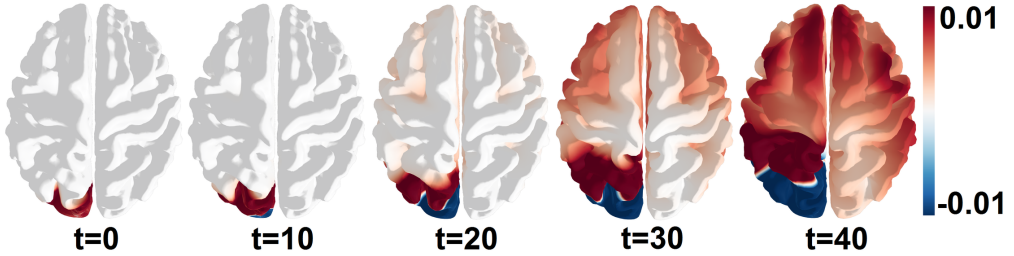

Fig 3. Dynamics of the damped-wave equation on the human connectome include non-local propagation along white-matter fibers. Shown are snapshots of simulated cortical activity that is governed by the damped wave equation with time-step $\delta t = 1$ and parameters $a = 1.5 \cdot 10^5$, $b = 2.5 \cdot 10^3$.

Stochastic Wilson-Cowan equations on graphs

The Wilson-Cowan model [20] is a widely used and successful model of cortical dynamics. In this section we show how to use the framework of graph neural fields (*Methods*) to implement the stochastic Wilson-Cowan equations on an arbitrary graphs equipped with a suitable distance metric, and how to compute spatiotemporal observables (CHAOSS). We then illustrate the effects of weighted and non-local graph edges on the dynamics in the simplified context of a one-dimensional graph, before moving on to a real-world application with fMRI data.

General formulation

In continuous space, the stochastic Wilson-Cowan model is described by the following system of integrodifferential equations:

$$\tau_E \frac{\partial E}{\partial t} = -d_E E + \mathrm{S}\left[\alpha_{EE}(K_{EE} \otimes E) - \alpha_{IE}(K_{IE} \otimes I) + P\right] + \sigma \xi_E, \tag{7}$$

$$\tau_I \frac{\partial I}{\partial t} = -d_I I + \mathrm{S}\left[\alpha_{EI}(K_{EI} \otimes E) - \alpha_{II}(K_{II} \otimes I) + Q\right] + \sigma \xi_I, \tag{8}$$

where the symbol $\otimes$ denotes a convolution integral, and we have omitted for brevity the spatiotemporal dependency of $E(x,t)$, $I(x,t)$, $\xi_E(x,t)$ and $\xi_I(x,t)$. In other words, this model posits the existence of two neuronal populations (Excitatory and Inhibitory) at each location in space. The fraction of active neurons in each population ($E(x,t)$, $I(x,t)$) evolve according to a spontaneous decay with rate d_E and d_I, a sigmoid-mediated activation term containing the four combinations of population interactions (*E-E, I-E, E-I, I-I*) as well as the external input terms P and Q, stochastic noise realizations $\xi_E(x,t)$ and $\xi_I(x,t)$ of intensity σ, and with the temporal scaling parameters τ_E and τ_I. The propagation of activity and interaction among neuronal populations is modeled by spatial convolution integrals with four, potentially different, kernels ($K_{EE}, K_{IE}, K_{EI}, K_{II}$). For arbitrary symmetric spatial kernels, convolutions can be formulated as linear matrix-vector products on graphs (*Methods* Eq (45)). Therefore, the stochastic Wilson-Cowan equations on graphs reduce to

$$\tau_E \dot{E}(t) = -d_E E(t) + \mathrm{S}\big[\alpha_{EE} K_{EE} E(t) - \alpha_{IE} K_{IE} I(t) + P\big] + \sigma \xi_E(t) \tag{9}$$

$$\tau_I \dot{I}(t) = -d_I I(t) + \mathrm{S}\big[\alpha_{EI} K_{EI} E(t) - \alpha_{II} K_{II} I(t) + Q\big] + \sigma \xi_I(t) \tag{10}$$

$E(t)$, $I(t)$, $\xi_E(t)$ and $\xi_I(t)$ are functions on the graph, i.e. vectors of size n, where n is the number of vertices in the graph; the convolutions are implemented via the graph-filters K_{**}, which are matrices of size (n,n). In particular, for the case of Gaussian convolutions, the filters are given by (Table 2)

$$K_{EE} = U e^{\sigma_{EE}^2 \Lambda/2} U^T, \quad K_{IE} = U e^{\sigma_{IE}^2 \Lambda/2} U^T, \tag{11}$$

$$K_{EI} = U e^{\sigma_{EI}^2 \Lambda/2} U^T, \quad K_{II} = U e^{\sigma_{II}^2 \Lambda/2} U^T. \tag{12}$$

With the distance-weighted graph Laplacian $\Delta = U \Lambda U^T$, so that Λ is the diagonal matrix of graph Laplacian eigenvalues, U and U^T the corresponding eigenvectors.

This model formulation also allows for non-Gaussian convolution kernels; and the inclusion of a stochastic term allows for characterization of resting-state activity as noise-induced fluctuations about a stable steady-state (E^*, I^*) [38].

Symbol	Meaning
E	Fraction of active excitatory subpopulations.
I	Fraction of active inhibitory subpopulations.
τ_E, τ_I	Time constants.
d_E, d_I	Activity decay rates.
$\mathrm{S}[x]$	$1/(1+e^{-x})$.
$\alpha_{EE}, \alpha_{IE}, \alpha_{EI}, \alpha_{II}$	Strength of connectivity between populations.
$K_{EE}, K_{IE}, K_{EI}, K_{II}$	Convolution kernels, graph filters.
$\sigma_{EE}, \sigma_{IE}, \sigma_{EI}, \sigma_{II}$	Standard deviation of Gaussian kernels.
P	Constant input to E populations.
Q	Constant input to I populations.
σ	Noise amplitude.
ξ_E, ξ_I	Noise realizations.

Table 1. Meaning of symbols in the continuous Wilson-Cowan equations and their graph equivalents.

Linear stability analysis

In order to compute meaningful spatiotemporal observables with CHAOSS, it is first necessary to obtain solutions the steady-state equations and their linear stability. For generality and compactness of notation, let us define a new column vector $\mathbf{u}$(t) as the concatenation of $E(t)$ and $I(t)$; We express $\mathbf{d}$ for the diagonal matrix containing the damping parameters d_E and d_I and $\boldsymbol{\tau}$ for the matrix containing the timescale parameters τ_E and τ_I. The matrix K contains the four (arbitrary) graph-filters, $\mathbf{X}$ is the concatenated vector encoding subcortical inputs P and Q. We can now write the original system of Eq (9-10) with a single equation

$$\boldsymbol{\tau}\dot{\mathbf{u}}(t) = -\mathbf{d}\mathbf{u}(t) + \mathrm{S}\big[K\mathbf{u}(t) + \mathbf{X}\big] + \sigma\boldsymbol{\xi} \tag{13}$$

Note that this expression potentially allows for space-dependent model parameters. The steady-state(s) $\mathbf{u}^*$ can be obtained by setting the time-derivative and noise amplitude σ to zero and solving the resulting steady state (matrix) equation:

$$\mathrm{S}\big[K\mathbf{u}^* + \mathbf{X}\big] = \mathbf{d}\mathbf{u}^* \tag{14}$$

This equation does not have an analytic, exact, closed form solution (in fact, it doesn't even necessarily have a solution. The sigmoid function is bound between -1 and 1, but $\mathbf{d}\mathbf{u}^*$ is not). Furthermore, in the context of a whole-brain model on a mesoscopic connectome, the equation is very high dimensional (twice the number of vertices n, with $n \sim 18000$ in our case). This makes a brute-force numerical approach to the determination of steady states computationally inefficient, especially because it would have to be repeated for each parameter set under examination.

Solutions to the steady state equations If we restrict our analysis to *spatially homogeneous* steady states and *space-independent* model parameters, the steady state

equation simplifies to the following 2-dimensional system, rather than the original $2n$-dimensional Eq (14):

$$\frac{1}{1+exp(\alpha_{IE}I^* - \alpha_{EE}E^* - P)} = d_E E^* \tag{15}$$

$$\frac{1}{1+exp(\alpha_{II}I^* - \alpha_{EI}E^* - Q)} = d_I I^* \tag{16}$$

Solutions to this 2-dimensional system can be rapidly obtained numerically for any set of parameters. The biologically valid steady state(s) of the model are given by the solutions (E^*, I^*) with $E^*, I^* \in [0,1]$, since E and I here represent the fraction of active neurons within the respective population. Once a valid steady state is obtained, its stability can be determined through the Jacobian eigenspectrum, and verified with numerical simulations.

General Jacobian To obtain the Jacobian eigenspectrum and determine the stability of a steady state, we have to linearize Eq (13) as

$$\dot{\mathbf{u}}(t) = -\boldsymbol{\tau}^{-1}\mathbf{d}\mathbf{u}(t) + \boldsymbol{\tau}^{-1}\mathrm{S}\big[K\mathbf{u}(t) + \mathbf{X}\big] = F(\mathbf{u}) \approx F(\mathbf{u}^*) + \left.\frac{\partial F(\mathbf{u})}{\partial \mathbf{u}}\right|_{\mathbf{u}^*} \delta\mathbf{u} \tag{17}$$

With $F(\mathbf{u}^*) = 0$ by definition since $\mathbf{u}^*$ is a steady state, and defining a small perturbation abut the steady state $\delta\mathbf{u} = (\mathbf{u} - \mathbf{u}^*)$. The Jacobian is then:

$$J(\mathbf{u}^*) = \left.\frac{\partial F(\mathbf{u})}{\partial \mathbf{u}}\right|_{\mathbf{u}^*} = -\boldsymbol{\tau}^{-1}\mathbf{d} + \mathrm{Diag}\left(\boldsymbol{\tau}^{-1}\mathrm{S}'[K\mathbf{u}^* + \mathbf{X}]\right) K \tag{18}$$

To simplify this expression and allow further analytic progress, we apply the property of the sigmoid derivative $\mathrm{S}'(x) = \mathrm{S}(x)(1 - \mathrm{S}(x))$; the steady state equation $\mathrm{S}\big[K\mathbf{u}^* + \mathbf{X}\big] = \mathbf{d}\mathbf{u}^*$; and finally denoting with $\circ$ the Hadamard (element-wise) product we obtain

$$J(\mathbf{u}^*) = -\boldsymbol{\tau}^{-1}\mathbf{d} + \mathrm{Diag}\left(\boldsymbol{\tau}^{-1}\mathbf{d}\mathbf{u}^* \circ (\mathbf{1} - \mathbf{d}\mathbf{u}^*)\right) K \tag{19}$$

We have thus obtained a general expression for the Jacobian of the Wilson-Cowan model on graphs, which holds also for non-homogeneous steady states and/or space-dependent parameters. In order to evaluate the linear stability of any steady state, it is sufficient plug in the steady state $\mathbf{u}^* = (E^*, I^*)$ solution to Eq (14) in Eq (19), and study the eigenspectrum of the resulting Jacobian.

Analytic form of the Jacobian eigenspectrum On a mesoscopic human connectome graph such as the one we use here, the general Jacobian of Eq (19) is a dense matrix with more than 10^8 elements. Its eigenspectrum can be calculated numerically, but such a computation is not particularly fast, and has to be repeated for each steady state of each parameter set under examination. However, we restrict the problem to *homogeneous* steady states and *space-independent* model parameters, it is possible to use the properties of the graph Laplacian to obtain an analytic expression for the Jacobian

eigenspectrum, which can be computed extremely quickly[2]. These assumptions can be relaxed to allow for eigenmode-dependent[3] model parameters. Define the scalar, steady-state-dependent parameters:

$$a = d_E E^*(1 - d_E E^*), \qquad b = d_I I^*(1 - d_I I^*) \tag{20}$$

We can use the general Jacobian of Eq (19), with Gaussian kernels, to write explicitly the *linearized* Wilson-Cowan equations for the time-evolution of a perturbation about a homogeneous steady state:

$$\tau_E \dot{E}(t) = -d_E E(t) + a\alpha_{EE} U e^{\sigma_{EE}^2 \Lambda/2} U^T E(t) - a\alpha_{IE} U e^{\sigma_{IE}^2 \Lambda/2} U^T I(t) + \sigma \xi_E(t) \tag{21}$$

$$\tau_I \dot{I}(t) = -d_I I(t) + b\alpha_{EI} U e^{\sigma_{EI}^2 \Lambda/2} U^T E(t) - b\alpha_{II} U e^{\sigma_{II}^2 \Lambda/2} U^T I(t) + \sigma \xi_I(t) \tag{22}$$

Applying the graph Fourier transform U^T, the equations are diagonalized. Each eigenmode of the graph Laplacian therefore behaves independently as a 2-dimensional linear system, with the Jacobian for the kth eigenmode being

$$J_k = \begin{bmatrix} -\frac{d_E}{\tau_E} + \frac{a}{\tau_E}\alpha_{EE} e^{\sigma_{EE}^2 \lambda_k/2} & -\frac{a}{\tau_E}\alpha_{IE} e^{\sigma_{IE}^2 \lambda_k/2} \\ \frac{b}{\tau_I}\alpha_{EI} e^{\sigma_{EI}^2 \lambda_k/2} & -\frac{b}{\tau_I}\alpha_{II} e^{\sigma_{II}^2 \lambda/2} - \frac{d_I}{\tau_I} \end{bmatrix} \tag{23}$$

The Jacobian eigenvalues can then be computed directly as

$$\lambda_{1,2}^{J_k} = \frac{\mathrm{Tr}\,(J_k) \pm \sqrt{\mathrm{Tr}^2\,(J_k) - 4\mathrm{Det}\,(J_k)}}{2} \tag{24}$$

From the Jacobian eigenspectrum thus obtained, we can directly infer the stability character of a steady state. Note that in order to obtain meaningful predictions for spatiotemporal observables, the steady state under examination has to be stable, that is, J_k must have no eigenvalues with positive real parts for all λ_k. We have used Gaussian kernels in the derivation, but the result can be straightforwardly generalized to all other kernels. The decoupling of the linearized Wilson-Cowan equations in the graph-Fourier domain has another important consequence: since the dimensionality of the system reduces from n^2 to n (where n is the number of vertices in the graph), very efficient numerical simulations of the linearized equations can be carried out directly in the graph Fourier domain.

[2]These assumptions are required because generally space-dependent parameters would be expressed by a non-constant diagonal matrix, that would not commute with the graph Fourier transform U^T.

[3]Eigenmode-dependent parameters would be expressed by a diagonal matrix in the graph Fourier domain, therefore by definition commuting with U^T.

Wilson-Cowan model CHAOSS

After having obtained steady state solutions, and determined stability character of the steady state, we can apply the *Connectome-Harmonic Analysis Of Spatiotemporal Spectra* (*Methods*) to characterize the spatiotemporal dynamics of resting-state brain activity. CHAOSS predictions, combined with a suitable observation model, can then be compared with empirical neuroimaging data, for example EEG, MEG, or fMRI. Consider linearized stochastic graph neural field equations in the graph Fourier domain:

$$D_t \hat{u}_k(t) = J_k \hat{u}_k(t) + \sqrt{B}\hat{\xi}_k(t), \tag{25}$$

Where $k = 1, ..., n$ indexes the graph Laplacian eigenmodes. For Wilson-Cowan model (Eq (21-22)), we have: $D_t = d/dt$, J_k is given in Eq (23) for the case with Gaussian kernels, and B is

$$B = \begin{bmatrix} \sigma^2/\tau_E^2 & 0 \\ 0 & \sigma^2/\tau_I^2 \end{bmatrix}, \tag{26}$$

In terms of the elements of the matrices J_k and B, the theoretical prediction for the harmonic-temporal power spectrum of the excitatory neural population activity becomes (*Methods* Eq (72)):

$$[S_k(\omega)]_{00} = \frac{[B]_{00}\left([J_k]_{11}^2 + \omega^2\right) + [J_k]_{01}^2[B]_{11}}{\left([J_k]_{00}[J_k]_{11} - [J_k]_{01}[J_k]_{10} - \omega^2\right)^2 + \omega^2\left([J_k]_{00} + [J_k]_{11}\right)^2} \tag{27}$$

The double-digits numerical subscripts in Eq (27) refer to the row-column element of the respective matrix. Eq (27) can be used to compute the separate harmonic and temporal power spectra of the model, as well as the correlation and functional connectivity matrices (*Methods* Eq (77)). Equivalent formulas for the inhibitory population can also be similarly derived.

By integrating $[S_k(\omega)]_{00}$ over all temporal frequencies, an explicit expression for the harmonic power spectrum of excitatory activity can be obtained:

$$H_E(k) = \frac{[B]_{00}([J_k]_{00}[J_k]_{11} - [J_k]_{01}[J_k]_{10}) + [J_k]_{11}^2[B]_{00} + [J_k]_{01}^2[B]_{11}}{2([J_k]_{01}[J_k]_{10} - [J_k]_{00}[J_k]_{11})([J_k]_{00} + [J_k]_{11})} \tag{28}$$

Eq (27) and (28) represent a general result that does not only apply to the Wilson-Cowan model. In fact, these equations describe the harmonic-temporal and harmonic power spectra of stochastic equilibrium fluctuations for the first population of any graph neural field model with two interacting populations and a first-order temporal differential operator.

Effects of distance-weighting and long-range connectivity

We characterize the effects internode spacing and the presence of long-range connectivity on model dynamics by implementing the Wilson-Cowan model on a one-dimensional graph with 1000 nodes. Numerical simulations were carried out with a time-step $\delta t = 5 \cdot 10^{-5} s$ and an observation time of 10^5 time-steps, which corresponds to $5s$ of

simulated activity.

To show the effects of distance-weighting in graph neural fields, we note how, for the parameter set of S1 Table, increasing the distance between nodes leads to the emergence of an oscillatory resonance that eventually destabilizes and gives way to limit-cycle activity. Keeping the number of nodes constant, increasing the internode spacing h alters the stability of the steady state from broadband activity ($h = 10^{-5}$m), to oscillatory resonance ($h = 10^{-4}$m), to oscillatory instability ($h = 2 \cdot 10^{-4}$m). The spatial and temporal power spectra for the case $h = 10^{-4}$m are shown Figure 4A and B, respectively. These results demonstrate that by using the weighted graph Laplacian, the dynamics of graph neural fields becomes dependent on how metric properties of the graph i.e. on how the graph is embedded in three-dimensional space. If the combinatorial graph Laplacian is used, as done in [9], the dynamics only depends on the topology of the cortical mesh and in this sense does not take into account the physical properties of the cortex.

The presence of *fast*, long-range connectivity can impact the power spectrum and functional interactions of equilibrium fluctuations, as well as the stability of steady-states. To illustrate this, we add a single non-local edge between nodes 250 and 750 to the one-dimensional graph with $h = 10^{-4}$m. The Euclidean distance between these two nodes is $500 \cdot h = 5 \cdot 10^{-2}$m= 5cm. In the healthy brain, myelination of white-matter fibers allows long-range (cortico-cortical) activity propagation to take place at speeds ~ 200 times greater than local (intra-cortical) propagation [21]. To model myelination, we set the length of the non-local edge to be the Euclidean distance between the nodes, divided by a factor of 200 (similarly to the construction of the human connectome graph Laplacian, where the length of cortico-cortical edges is set to be their path-length distance along DTI fibers, divided by a factor of 200). Therefore, the effective length of the non-local edge is $2.5 \cdot 10^{-4}$m. Figure 5 shows the effects of the presence of the non-local edge on the spatial and temporal power spectra of the equilibrium fluctuations. The most pronounced effect is damping of the oscillatory resonance in the temporal power spectrum, thus rendering the fluctuations more stable. Furthermore, the edge leads to a faint but discernible alteration in the functional connectivity (see Panels C and D).

Interestingly, when the model operates in the pathological i.e. non-stable regime ($h = 2 \cdot 10^{-4}$m), addition of a single non-local edge stabilizes the steady state, thus leading to healthy equilibrium fluctuations (see Figure 6). The non-local edge also creates a visible increase in the functional connectivity between the nodes involved, and a change in the pattern in neighboring nodes (Panels C and D). As noted above, these significant effects of long-range connectivity are observed if the effective length of the non-local edge is small enough for non-local activity propagation to interact with local activity propagation. For these one-dimensional simulations, this happens if the speed factor is larger than $\sim$50.

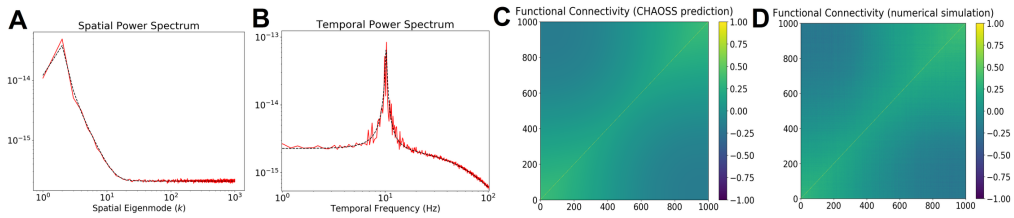

Fig 4. Effects of internode spacing on graph neural field dynamics. Panels A and B show, respectively, the spatial and temporal power spectrum of equilibrium fluctuations in the one-dimensional graph for internode spacing $h = 10^{-4}m$. The dotted black correspond to the theoretical prediction and the red lines are obtained through numerical simulations. Panels C and D show the model functional connectivity as obtained, respectively, by analytic predictions and numerical simulations.

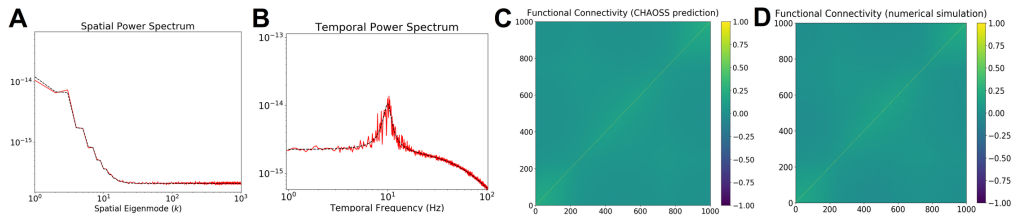

Fig 5. Suppression of oscillatory resonance by long-range connectivity. Panels A and B show, respectively, the spatial and temporal power spectrum of equilibrium fluctuations in the one-dimensional graph for internode spacing $h = 10^{-4}m$ after the addition of a non-local edge between nodes 250 and 750. Compare with Figure 4 to note the visible suppression of oscillatory resonance, and the slight change in functional connectivity engendered by a single non-local edge. The dotted black correspond to the theoretical prediction and the red lines are obtained through numerical simulations. Panels C and D show the model functional connectivity as obtained, respectively, by analytic predictions and numerical simulations.

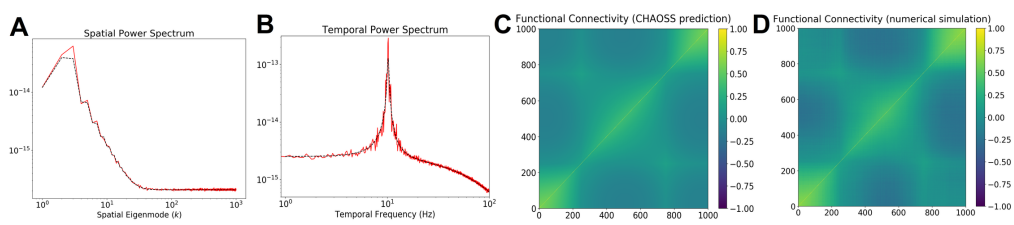

Fig 6. Abortion of pathological oscillations by long-range connectivity. Panels A and B show, respectively, the spatial and temporal power spectrum of equilibrium fluctuations in the one-dimensional graph for internode spacing $h = 2 \cdot 10^{-4}m$ after the addition of a non-local edge between nodes 250 and 750. Before the addition, the dynamics was placed in an unstable limit-cycle regime. The dotted black correspond to the theoretical prediction and the red lines are obtained through numerical simulations.Panels C and D show the model functional connectivity as obtained, respectively, by analytic predictions and numerical simulations.

Application to resting-state functional MRI

To illustrate a real-world application of graph neural fields, we implement the stochastic Wilson-Cowan model on the connectome graph of a single subject, and use it reproduce the harmonic power spectrum observed in resting-state fMRI data. The harmonic power spectrum of resting-state fMRI was computed, according to its definition, as the temporal mean of the squared graph-Fourier transform of the fMRI timecourses. To regularize the empirical fMRI harmonic power spectrum, we computed its log-log binned median with 300 bins, following [22]. Eigenmodes above $k \sim 15000$ contained artifacts due to reaching the limits of fMRI spatial resolution, and were thus removed. The model power spectra were binned in the same way. The parameters of the Wilson-Cowan model were optimized with a basinhopping procedure, aiming to minimize the residual difference between observed and predicted spectra. In the fitting, we allowed for a linear rescaling of the theoretical spectrum so that

$$H_{\mathrm{fMRI}}(k) = \beta H_E(k), \tag{29}$$

where $H_{\mathrm{fMRI}}(k)$ is the harmonic power spectrum of the fMRI data, $H_E(k)$ is the harmonic power spectrum of the excitatory population in the Wilson-Cowan model, and β is a rescaling parameter. Fig 7 shows the best fitting theoretical harmonic power spectra, together with the empirical fMRI spectrum. The model is able to reproduce the experimental spectrum. The parameters values obtained from the fitting are reported in S2 Table and appear consistent.

To verify the accuracy of predictions obtained with the connectome-harmonic analysis of spatiotemporal spectra, we implement numerical simulations of the linearized Wilson-Cowan equations on the human connectome graph. Simulations were carried out with a time-step value $\delta t = 10^{-4}s$ and an observation time of 10^6 time-steps, corresponding to $100s$ of simulated brain activity. Figure 8 shows snapshots of the simulated resting-state activity at several different times. Note that the two hemispheric meshes appear physically separate, but inter-hemispheric propagation is allowed through white-matter callosal tracts (not shown in the figure).

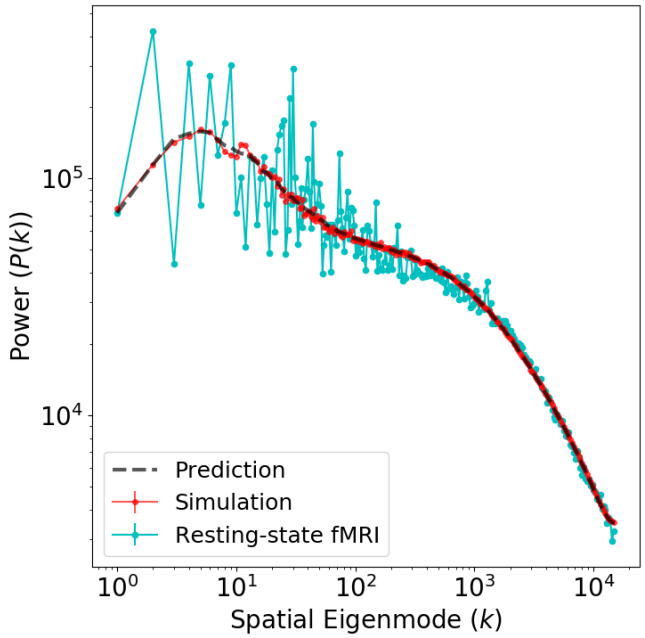

Fig 7. Stochastic Wilson-Cowan graph neural field model predicts the experimental harmonic power spectrum. Shown is the empirically observed BOLD-fMRI harmonic spectrum (cyan line), together with the theoretical (dotted black line) and numerical (red line) predictions from the stochastic Wilson-Cowan equations. The numerical prediction was obtained by taking the median of three independent simulations.

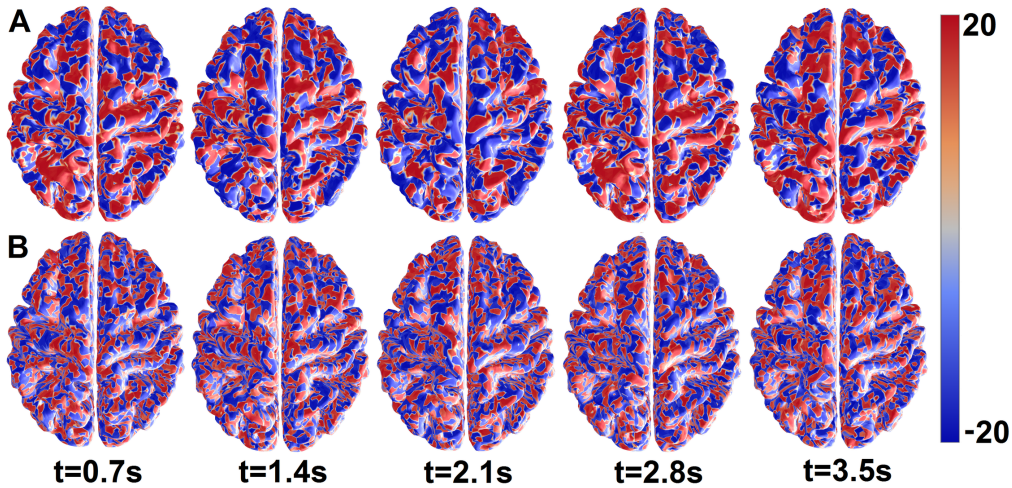

Fig 8. Measured and simulated resting-state brain activity. Panel A shows resting-state brain activity, as fluctuations of the BOLD fMRI signal about the mean at each node. Panel B shows snapshots of activity from the stochastic Wilson-Cowan equations simulated using the parameters of Table in S2 Table. The model activity was temporally downsampled to match the TR of fMRI data, and rescaled by $\sqrt{\beta}$ to match the scale of the BOLD signal. No spatial or temporal smoothings were applied.

Methods

Laplacian operators on graphs

In this section we provide a derivation of the weighted graph Laplacian operator in terms of graph differential operators. The weighted graph Laplacian is distinguished from the combinatorial graph Laplacian often used in analysis studies [9], as it allows physical properties of the cortex to be taken into account, which is necessary to implement "realistic" graph neural field models.

The combinatorial Laplacian. Consider an undirected graph with n vertices. The graph's binary *adjacency matrix* $\tilde{A}$ is defined by

$$\tilde{A}_{ij} = \begin{cases} 1 & \text{if} \quad i \sim j \\ 0 & \text{otherwise} \end{cases} \tag{30}$$

where $i \sim j$ means that vertices i and j are connected by an edge. The graph's *degree matrix* $\tilde{D}$ is a diagonal matrix whose diagonal entries are defined by

$$\tilde{D}_{ii} = \sum_{j=1}^{n} \tilde{A}_{ij}. \tag{31}$$

It hence counts the number of edges for each vertex i. The binary or *combinatorial graph Laplacian*, denoted by $\tilde{\Delta}$, is defined as

$$\tilde{\Delta} = \tilde{A} - \tilde{D}. \tag{32}$$

The combinatorial graph Laplacian and its normalized version do not carry information about the distances between cortical vertices and therefore are invariant under topological but non-isometric deformations of graph. Neural activity modeled in terms of the combinatorial graph Laplacian therefore is a topological graph invariant, whereas real neural activity does depend on the metric properties of the graph. The combinatorial graph Laplacian, however, can be adjusted so as to take into account the metric properties of the graph, yielding the *weighted graph Laplacian*. Below, we provide a derivation of the weighted graph Laplacian in terms of the graph directional derivaties of a graph function.

The weighted Laplacian. Let f be a *graph function* i.e. a function defined on the vertices of a graph and let M be the graph's distance matrix. Thus, the $(i,j)^{th}$ entry M_{ij} of M equals the distance between nodes i and j in a particular metric. We note that for this derivation, it is irrelevant how M is obtained. In the context of connectomes, the elements of M can be defined in terms of suitably scaled Euclidean distances, geodesic distances over the cortical manifold, or as the lengths of the white matter fibers connecting the vertices. The *first-order graph directional derivative* $\partial_j f_i$ of f at vertex i in the direction of vertex j is defined as

$$\partial_j f_i = \frac{\tilde{A}_{ij}}{M_{ij}}(f_j - f_i), \tag{33}$$

Note that according to this definition, $\partial_j f_i = 0$ if vertex j is not connected to vertex i,

and that $\partial_i f_i = 0$. Also note that ∂_j is a linear operator on the vector space of graph signals. Furthermore, since $\tilde{A}_{ij}^2 = \tilde{A}_{ij}$, the *second-order graph directional derivative* $\partial_j^2 f_i$ of f at vertex i in the direction of vertex j is defined as

$$\partial_j(\partial_j f_i) = \partial_j^2 f_i = \frac{\tilde{A}_{ij}}{M_{ij}^2}(f_i - f_j). \tag{34}$$

Following the definition of the Laplacian operator in Euclidean space as the sum of second-order partial derivatives, the *weighted graph Laplacian*, denoted as Δ is defined as [4]

$$\Delta f_i = -\sum_{j=1}^{n} \partial_j^2 f_i. \tag{35}$$

To see the relation with the combinatorical graph Laplacian, we note that Δ can be written in matrix form as

$$\Delta = A - D, \tag{36}$$

where A and D are the *weighted adjacency matrix* and *weighted degree matrix*, respectively, which are defined as $A_{ij} = \tilde{A}_{ij}/M_{ij}^2$ and $D_{ii} = \sum_{j=1}^{n} A_{ij}$, respectively. Thus, the weighted graph Laplacian can be obtained by using the weighted versions of the adjacency and degree matrices in the definition of the combinatorical graph Laplacian.

Convolution of graph signals

In order to define graph neural fields, we need to have a graph-theoretical analog of the continuous spatiotemporal convolution $(K \otimes u)(x,t)$

$$(K \otimes u)(x,t) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} K(x-x', t-t')u(x',t')dx'dt', \tag{37}$$

To obtain this, we use the convolution theorem to represent the convolution in the spatiotemporal Fourier domain as $\hat{K}(k,\omega)\hat{u}(k,\omega)$, where k denotes wavenumber amd ω denotes temporal frequency. When the kernel is real-valued and spatially symmetric, its Fourier transform is real-valued and even in k, so that $\hat{K}(k,\omega)$ can be viewed as a function of $-k^2$. To generalize the convolution operator to graphs, we use the fact that $-k^2$ is the eigenvalue of the spatial Fourier basis function e^{ikx} under the Laplace operator d^2/dx^2. Since the Laplace operator on a graph is given by its Laplacian matrix Δ, the graph filter corresponding to the spatiotemporal convolution $(K \otimes u)(x,t)$ can be defined by substituting the eigenvalues of Δ for the values $-k^2$ in $\hat{K}(-k^2,\omega)$.

The graph Fourier transform. We consider a graph with n vertices and Laplacian matrix Δ, together with an eigendecomposition $\Delta = U\Lambda U^T$, where U is an $n \times n$ orthogonal matrix containing an eigenbasis of Δ in its columns and Λ is the $n \times n$ diagonal matrix that contains the corresponding eigenvalues $\lambda_1 \geq \lambda_2 \geq, \cdots, \geq \lambda_n \geq 0$ on its diagonal. The *graph Fourier transform* of a function $u(t)$ on the graph is now defined by

$$\hat{u}(t) = U^T u(t), \tag{38}$$

[4] One could also write $\nabla = [\partial_1, ..., \partial_N]$ for the "graph-gradient", and then define $\Delta = -\nabla \cdot \nabla$. The minus sign arises to maintain the analogy with the continuous Laplacian.

where the transformation U^T expresses $u(t)$ in the eigenbasis of Δ. The vertex-domain signal $u(t)$ can be recovered again by applying the *inverse graph Fourier transform* $U^{-T} = U$ to $\hat{u}(t)$. For clarity, note that the graph Fourier transform is not related to the temporal Fourier transform and that $u(t)$ does not have to depend on time to apply it. In particular, for grid graphs (i.e. graphs whose drawing, embedded in some Euclidean space, forms a regular tiling), the graph Fourier transform is equivalent, in the continuum limit, to the spatial Fourier transform in Euclidean space. However, the graph Fourier transform as defined here is a more general concept, since it can be applied to arbitrary nondirected graphs, possibly with non-local edges, such as the human connectome.

Filtering of graph signals. Let $K(x,t)$ be a kernel with Fourier transform $\hat{K}(-k^2,\omega)$ and let $u(t)$ be a graph signal with spatial Fourier transform $\hat{u}(t)$. We denote the temporal Fourier transforms of $u(t)$ and $\hat{u}(t)$ by $u(\omega)$ and $\hat{u}(\omega)$, respectively. We define the graph kernel $\hat{K}_g$ associated with the continuous kernel K

$$\hat{K}_g = \mathrm{diag}(\hat{K}(\lambda_1,\omega),\cdots,\hat{K}(\lambda_n,\omega)). \tag{39}$$

In the graph Fourier and temporal frequency domains, the filtered signal is hence per definition given by

$$\hat{u}^{\mathrm{filt}}(\omega) = \hat{K}_g \hat{u}(\omega) \tag{40}$$

Applying the inverse graph Fourier transform U, we obtain the filtered signal in the graph domain

$$u^{\mathrm{filt}}(\omega) = U\hat{K}_g\hat{u}(\omega) = U\hat{K}_g U^T u(\omega) = K_g u(\omega), \tag{41}$$

where we have defined $K_g = U\hat{K}_g U^T$, the graph domain representation of the filter. In other words, K_g is the graph filtering operator corresponding to the continuous kernel K.

The above representations of the kernel and the filtered signal are in the temporal frequency domain. To obtain the corresponding time domain representations, we note that the ith entry of $u^{\mathrm{filt}}(\omega)$ is given by

$$u_i^{\mathrm{filt}}(\omega) = \sum_{j=1}^{n} K_g^{i,j}(\omega) u_j(\omega), \tag{42}$$

where $K_g^{i,j}(\omega)$ denotes the $(i,j)^{th}$ entry of $K_g(\omega)$. Using the convolution theorem, the time domain representation of the filtered signal is

$$u_i^{\mathrm{filt}}(t) = \sum_{j=1}^{n} (K_g(i,j) \otimes u_j)(t). \tag{43}$$

Collecting the terms for all n entries in a column vector yields the filtered signal in the graph and temporal domain and we write it as $(K_g \otimes_g u)(t)$:

$$u^{\mathrm{filt}}(t) = (K_g \otimes_g u)(t). \tag{44}$$

In case of a purely spatial kernel $K(x,t) = K(x)$, with $K(x)$ symmetric and with spatial

Fourier transform $\hat{K}(-k^2)$, the graph kernel $\hat{K}_g$ is independent of frequency so that the filtering operator reduces to the frequency-independent linear transformation

$$(K_g \otimes_g u)(t) = K_g u(t), \tag{45}$$

In case of a purely temporal kernel $K(x,t) = g_\Theta(t)$, with $g_\Theta(t) = g(t)\Theta(t)$, where $g(t)$ is the temporal kernel, and $\Theta(t)$ the Heaviside step function, which ensures that integration is temporally causal. Graph filtering reduces to a temporal convolution:

$$(K_g \otimes_g u)(t) = (g_\Theta \otimes u)(t). \tag{46}$$

Lastly, in case of a separable kernel $K(x,t) = w(x)g_\Theta(t)$, the graph kernel decomposes as

$$\hat{K}_g = g_\Theta(\omega)\mathrm{diag}(\hat{w}(\lambda_1), \cdots, \hat{w}(\lambda_n)), \tag{47}$$

so that the filtered signal is given by

$$(K_g \otimes_g u)(t) = (g_\Theta \otimes (K_g u))(t), \tag{48}$$

In this case, the filtering occurs in two stages: the signal of first spatially filtered by K_g and subsequently temporally filtered by the convolution with $g_\Theta(t)$.

Examples of graph kernels. We have defined spatiotemporal graph kernels for maximal generality, but in practice, spatial kernels are particularly efficient and simple to use, since they reduce to simple linear products in the graph domain. Table 2 lists several commonly used continuous spatial kernels and their equivalent on graphs. The graph Laplacian converges to the Laplacian operator in the continuum limit; therefore on grid graphs (i.e. graphs whose drawing, embedded in some Euclidean space, forms a regular tiling), the graph filters act simply like discretized versions of their continuous counterparts. For example, applying the graph Gaussian kernel to a function defined on a 2-dimensional square-grid graph is equivalent to blurring/smoothing a 2-dimensional image with a spatial Gaussian kernel of the same size, and closed boundary conditions (open boundaries can be implemented by extending the graph beyond the image size, and periodic boundaries by adding edges connecting nodes on opposite sides). More relevantly to the current work, this approach generalizes to arbitrary non-grid graphs, potentially with non-local edges, such as the human connectome, and is therefore more general than grid-based discretizations of continuous convolution kernels.

Kernel	**Euclidean domain**	**Fourier domain**	$\hat{\boldsymbol{K}}_{\boldsymbol{g}}$
Gaussian	$e^{-x^2/2\sigma^2}$	$e^{-\sigma^2 k^2/2}$	$e^{\sigma^2\lambda_k/2}$
Exponential	$e^{-\alpha\lvert x\rvert}$	$\frac{1}{\alpha^2+k^2}$	$\frac{1}{\alpha^2-\lambda_k}$
Mexican hat	$(1-(x/\sigma)^2)e^{-x^2/2\sigma^2}$	$k^2 e^{-\sigma^2 k^2/2}$	$-\lambda_k e^{\sigma^2\lambda_k/2}$
Rectangular	$\mathrm{rect}(ax)$	$\mathrm{sinc}\left(\frac{k}{2\pi a}\right)$	$\mathrm{sinc}\left(\frac{\sqrt{-\lambda_k}}{2\pi a}\right)$
Triangular	$\mathrm{tri}(ax)$	$\mathrm{sinc}^2\left(\frac{k}{2\pi a}\right)$	$\mathrm{sinc}^2\left(\frac{\sqrt{-\lambda_k}}{2\pi a}\right)$

Table 2. Spatiotemporal convolution kernels in the Euclidean, Fourier, and graph domains. Normalization factors are omitted. The sampling function sinc is defined as $\mathrm{sinc}(x) = \sin(\pi x)/(\pi x)$.

Reaction-diffusion neural activity models

In this section we show how graph filters can also be used to implement the graph equivalents of neural activity models that can be directly written as partial differential equations [18, 23] and, among others, comprise damped wave and reaction-diffusion equations.

Consider the following scalar reaction-diffusion model

$$D_t u(x,t) = \frac{\partial^2}{\partial x^2} u(x,t) + \sigma\xi(x,t), \tag{49}$$

where D_t denotes a temporal differential operator that describes local reactions, $\partial^2/\partial x^2$ is the diffusion term and $\sigma\xi(x,t)$ is an external forcing function. To obtain the corresponding graph equation we first transform Eq (49) to the spatiotemporal Fourier domain:

$$D(\omega)\hat{u}(k,\omega) = -k^2\hat{u}(k,\omega) + \sigma\hat{\xi}(k,\omega), \tag{50}$$

where $D(\omega)$ denotes the Fourier transform of $D(t)$, and subsequently solve for $\hat{u}(k,\omega)$:

$$\hat{u}(k,\omega) = \sigma\hat{K}(k,\omega)\hat{\xi}(k,\omega), \tag{51}$$

where the spatiotemporal kernel $\hat{K}(k,\omega)$ is given by

$$\hat{K}(k,\omega) = \frac{1}{D(\omega) + k^2}. \tag{52}$$

The corresponding graph kernel is given by

$$\hat{K}_g = (D(\omega) - \Lambda)^{-1}, \tag{53}$$

where Λ denotes the diagonal matrix containing the eigenvalues of the weighted graph Laplacian Δ. Applying $\hat{K}_g$ to the input gives

$$\hat{u}(\omega) = \sigma\hat{K}_g\hat{\xi}(\omega). \tag{54}$$

Transforming back to the graph-temporal domain to obtain $u(t) = \sigma(K_g \otimes \xi)(t)$. To obtain the full system of differential equations, we note that

$$(D(\omega) - \Lambda)\hat{u}(\omega) = \sigma\hat{\xi}(\omega), \tag{55}$$

and transform this equation back to the spatial domain to obtain the following system of ordinary differential equations:

$$D_t u(t) = \Delta u(t) + \sigma\xi(t). \tag{56}$$

This shows that a reaction diffusion equation can directly be defined on a graph by replacing the Laplace operator in continuous space by the weighted graph Laplacian Δ, and solved by computing a suitable graph filter.

Graph neural fields

Continuous neural fields

Neural field models describe the dynamics of cortical activity $u(x,t)$ at time t and cortical location $x \in \Omega$. Here, $\Omega \in \mathbb{R}^3$ denotes the cortical manifold embedded in three-dimensional Euclidean space. Depending on the physical interpretation of the state variable $u(x,t)$, neural fields come in two types, which we will refer in the rest of the text as *Type 1* and *Type 2*. This short description is by no means meant to be exhaustive, and only contains the required background to define graph neural fields; comprehensive treatments of continuous neural fields are provided in [18, 23].

In Type 1 neural fields [18], the state variable $u(x,t)$ describes the average membrane potential at location x and time t. The general form of a neural field model of Type 1 is

$$D_t u(x,t) = \int_\Omega K(d(x,x'))S\left[u(x',t)\right] dx' + \sigma\xi(x,t), \tag{57}$$

where $\sigma\xi(x,t)$ is the external input, $d(x,x')$ is the geodesic distance between cortical locations x and x', K is the spatial kernel of the neural field that describes how the firing-rate $S[u(x',t)]$ at location x' affects the voltage at location x, and S is the firing-rate function that converts voltages to firing-rates. Furthermore, D_t is a temporal differential operator that models the synaptic dynamics. In modelling ongoing cortical activity, $\xi(t,x)$ is usually taken to be a stationary stochastic process. Following earlier studies, we will assume $\xi(t,x)$ to be spatiotemporally white-noise, i.e.

$$\mathbb{E}\left[\xi(x,t)\xi(x',t'\right] = \delta_{xx'}\delta_{tt'}, \tag{58}$$

where $\mathbb{E}$ denotes expectation value and $\delta_{aa'}$ is the Dirac delta function, which is defined by $\delta_{aa'} = 1$ if $a = a'$ and zero otherwise. Note that in principle, colored noise could be used as well. The distance function $d(x,x')$ between cortical locations x and x' as well as the integration over the cortical manifold Ω assume that Ω is equipped with a Riemannian metric. A natural metric is the Euclidean metric induced by the embedding of the cortical manifold in three-dimensional Euclidean space.

In Type 2 neural field models [24, 25], the state variable $u(x,t)$ denotes the fraction of active cells in a local cell population at location x and time t and hence takes values in the interval $[0,1]$. Type 2 neural field models have the form

$$D_t u(x,t) = S\left[\int_\Omega K(d(x,x'))u(x',t)dx'\right] + \sigma\xi(x,t), \tag{59}$$

where S denotes the activation function that maps fractions to fractions and hence takes values in the interval $[0,1]$ and thus has a different interpretation than the firing-rate function in Type 1 neural field models. Note that the only difference between Type 1 and Type 2 neural field models is the placement of the non-linear function S. Most neural field models used in practice are obtained by coupling two or more field models, where each model corresponds to a different neural population. For example, the state variable of the Wilson-Cowan neural field model is two-dimensional and its components correspond to excitatory and inhibitory neural populations. In the rest of the text, we will assume that K is normalized to have unit surface area.

In theoretical studies on neural field models, the cortex is usually assumed to be flat:, i.e. $\Omega = \mathbb{R}^2$ (cortical sheet) or $\Omega = \mathbb{R}^1$ (cortical line) or a closed subset thereof.

The major simplification that occurs in this case is that the cortical metric reduces to the Euclidean metric:

$$d(x, x') = ||x - x'||, \tag{60}$$

and, as a consequence, the integrals in Eq (57) and (59) reduce to convolutions so that Fourier methods can be used in the models' analysis (however, see [27] for a detailed theoretical study of a neural field model on the sphere). Note that in Euclidean space, the kernel K is symmetric, i.e. $K(-x) = K(x)$ for all $x \in \mathbb{R}$. The definition of graph neural fields as given in Section *Graph neural fields* relies on this fact as well. For ease of notation, below we set $\Omega = \mathbb{R}^1$ so that $d(x, x') = |x - x'|$. Note that the double integrals in Eq (57) and (59) can be written as convolutions, and hence it will be possible to implement them on graphs using the previously derived definitions of graph spatiotemporal convolutions.

Graph neural fields

Neural fields given by Eq (57) and (59) can be defined on an arbitrary graph by replacing the spatial filter K by its graph-equivalent K_g as defined in Section *Convolution of graph signals*. Thus, a *graph neural field of Type 1* is a model of the form

$$D_t u(t) = K_g S[u(t)] + \sigma\xi(t), \tag{61}$$

and

$$D_t u(t) = S\left[K_g u(t)\right] + \sigma\xi(t), \tag{62}$$

where $u(t)$ and $\xi(t)$ are graph functions. Thus, the graph-theoretical analogs of the continuous neural fields given by Eq (57) and (59) are ordinary differential equations. When more than one type of neural population is included or when the differential operator D_t is of order higher than one, the continuous neural fields reduce to *systems* of ordinary differential equations.

The continuous neural fields in Eq (57) and (59) are described by partial integro-differential equations in which the integration in done over space. They can also be described by spatiotemporal integral equations by viewing the differential operator as a temporal integral, which leads to a more general class of continuous neural fields. Using the spatiotemporal convolution on graphs, this more general class of neural fields can be formulated on graphs leading to a systems of temporal integral equations. To make this explicit, we use the definition of the spatiotemporal graph filtering operator $K_g\otimes$ to write out the ith component of u_i:

$$u_i(t) = \int_{-\infty}^{\infty} \left(\sum_{j=1}^{n} K_g^{ij}(s) S[u_j(t-s)] \right) ds + \sigma \int_{-\infty}^{t} \xi_i(t') dt'. \tag{63}$$

Thus, the spatiotemporal integrals in continuous neural fields are replaced by temporal integrals in graph neural fields and the spatial structure of the continuous neural field is incorporated into the graph kernels K_g^{ij}. The same applies to neural fields of Type 2. Furthermore, for continuous neural fields of Type 1 with separable kernels, and for special choices of the temporal component of the kernel, the spatiotemporal integral equation can be reduced to a partial integro-differential equation [23, 26]. For graph

neural fields there exists an equivalent subset of models that can be represented by a system of ordinary integro-differential equations.

In case of a purely spatial kernel $K(x,t) = w(x)$ we obtain the following systems of ordinary differential equations:

$$D_t u(t) = K_g S[u(t)] + \sigma\xi(t), \tag{64}$$

and

$$D_t u(t) = S\left[(K_g u_g)(t)\right] + \sigma\xi(t), \tag{65}$$

respectively, and in case of a purely temporal kernel $K(x,t) = g_\Theta(x)$ we obtain the following systems of ordinary differential equations:

$$D_t u(t) = (g_\Theta \otimes_g S[u])(t) + \sigma\xi(t), \tag{66}$$

and

$$D_t u(t) = S\left[(g_\Theta \otimes_g u)(t)\right] + \tag{67}$$

respectively. In case of a separable kernel $K(x,t) = w(x)g_\Theta(t)$ we obtain the following systems of ordinary differential equations:

$$D_t u(t) = (g_\Theta \otimes_g K_g S[u])(t) + \sigma\xi(t), \tag{68}$$

and

$$D_t u(t) = S\left[(g_\Theta \otimes_g K_g u)(t)\right] + \sigma\xi(t), \tag{69}$$

respectively.

Relating graph neural fields to experimental observables

Connectome-harmonic analysis of spatiotemporal spectra (CHAOSS)

To characterize the spatiotemporal statistics of resting-state brain dynamics, we derive analytic predictions for harmonic and temporal spectra and functional connectivity matrices. The derivation relies on the fact that, for space-independent parameters, linearized graph neural field equations decouple in the graph Fourier domain[5]. Each of the n graph Laplacian eigenmodes behaves independently like the following N-dimensional linear system, where N is the number of neuronal population types ($N = 2$ for the Wilson-Cowan model):

$$D_t \hat{u}_k(t) = J_k \hat{u}_k(t) + \sqrt{B}\hat{\xi}_k(t), \tag{70}$$

[5]Alternatively, for the case of space-dependent parameters, observables may be computed from numerical simulations

where $\xi_k(t)$ is an N-dimensional uncorrelated white noise, B is an N-dimensional diagonal matrix containing the noise intensity of the N neural populations, D_t is a temporal differential operator, and J_k is the Jacobian of the kth eigenmode. Taking the temporal Fourier transform we obtain

$$\hat{u}_k(\omega) = [D(\omega) - J_k]^{-1}\sqrt{B}\hat{\xi}_k(\omega), \tag{71}$$

where $D(\omega)$ denotes the temporal Fourier transform of D_t. Abbreviating the graph filter $\hat{K}_g = [D(\omega) - J_k]^{-1}$, the cross-spectral matrix $S_k(\omega)$ of the kth eigenmode is given by:

$$S_k(\omega) = \mathbb{E}[\hat{u}_k(\omega)\hat{u}_k(\omega)^\dagger] = \hat{K}_g\sqrt{B}\mathbb{E}[\hat{\xi}_k(\omega)\hat{\xi}_k(\omega)^\dagger]\sqrt{B}\hat{K}_g^\dagger = \hat{K}_g B\hat{K}_g^\dagger, \tag{72}$$

where, $\dagger$ denotes the conjugate transpose and $\mathbb{E}$ denotes the expected value. Colored noise can be modeled by letting B depend on ω, although this is usually not done in neural field modelling studies. Another possible generalization is to let B depend on the eigenmode k.

Eq (72) gives a closed-form expression for the N-dimensional cross-spectral matrix of the kth eigenmode, where N is the number of neuronal populations in the model. Hence, its diagonal entries $[S_k(\omega)]_{pp}, p = 0, ..., N-1$ describe the power of the pth neuronal population in the kth eigenmode, at temporal frequency ω. The temporal power spectrum $T_p(\omega)$ of the pth neuronal population is obtained by summing over eigenmodes:

$$T_p(\omega) = 2\sum_{k=1}^{n}[S_k(\omega)]_{pp}, \tag{73}$$

Where the factor of 2 arises because on graphs, k ranges only over positive integers between 1 and n (n is the number of eigenmodes). Similarly, the harmonic power spectrum of the pth neuronal population $H_p(k)$ is obtained by integrating over the temporal frequency ω:

$$H_p(k) = \frac{1}{2\pi}\int_{-\infty}^{+\infty}[S_k(\omega)]_{pp}d\omega, \tag{74}$$

where $1/2\pi$ is a normalization constant. When combined with a suitable observation model, these predictions can be compared with or fitted to experimental data from different neuroimaging modalities.

Functional connectivity Furthermore, it is possible to compute the correlation matrix of brain activity for each neuronal population. To construct the covariance matrix of a neuronal population activity Σ_p across all graph vertices, we first construct the covariance matrix $\hat{\Sigma}_p$ in the graph Fourier domain. Because of the independence of eigenmodes, the covariance matrix of the pth population (at lag zero) in the graph Fourier domain $\hat{\Sigma}_p$ is a diagonal matrix with the elements along the diagonal being the values of the spatial power spectrum.

$$\hat{\Sigma}_p = \mathrm{Diag}(H_p(k)) \tag{75}$$

The covariance matrix across all vertices is obtained by transforming back to the vertex domain:

$$\Sigma_p = U\hat{\Sigma}_p U^T \tag{76}$$

The functional connectivity (correlation) matrix F_p, which is often used in fMRI resting-state studies, is obtained by normalizing the covariance matrix

$$F_p = (\Sigma_p^+)^{-1/2}\Sigma_p(\Sigma_p^+)^{-1/2}, \tag{77}$$

where Σ_p^+ denotes Σ_p with all off-diagonal entries set to zero. Seed-based connectivity of the jth node is measured by the jth row (or column) of F_p.

Coherence matrix From the linearized model equations one can also derive the coherence matrix, which measures the strength and latency of linear interactions between pairs of vertices as a function of frequency ω and is often used in EEG and MEG studies [28]. If the noise is assumed to be white, non-linear connectivity measures such as the phase-locking value and amplitude correlations can be analytically computed from the coherence matrix [29]. For simplicity we derive the coherence matrix for the special case of a single population and note that the generalization to multiple populations is straightforward.

The derivation of the coherence matrix is similar to that of the functional connectivity, and starts by expressing the linearized model equations in the vertex domain:

$$D_t u(t) = Ju(t) + \sqrt{B}\hat{\xi}(t), \tag{78}$$

where the n-dimensional matrix J denotes the Jacobian matrix in the vertex-domain. Transforming Eq. (78) to the temporal Fourier domain and taking expectations yields the cross-spectral matrix $S_v(\omega)$ in the vertex domain:

$$S_v(\omega) = \mathbb{E}[u(\omega)u(\omega)^\dagger] = K_g B K_g^\dagger, \tag{79}$$

where $K_g = [D(\omega) - J]^{-1}$. The coherence matrix $C(\omega)$ is obtained by normalization of the cross-spectral matrix in the vertex domain:

$$C(\omega) = (S_v^+(\omega))^{-1/2} S_v(\omega)(S_v^+(\omega))^{-1/2}, \tag{80}$$

where $S_v^+(\omega))$ denotes $S_v(\omega)$ with its off-diagonal entries to zero. The $(i,j)^{th}$ entry of $C(\omega)$ is the coherence between the cortical activity at vertex i and j.

Anatomical and functional data

We use the same structural MRI and DTI data as in [9], obtained from the Human Connectome Project (https://db.humanconnectome.org/) to construct the anatomical connectome graph of a single subject. In short, MRI data is employed to obtain intra-cortical graph edges based on the surface mesh; DTI data is employed to add long-range cortico-cortical edges to the graph. The main difference with [9] is that instead of constructing the combinatorial (binary) graph Laplacian, here we construct a distance-weighted graph Laplacian. This allows us to take into account physical distance properties of the cortex that are relevant for graph neural fields, and that are otherwise lost. Specifically, intra-cortical edges are weighed by they 3D Euclidean distance; white matter edges are weighed by the distance along the respective DTI fiber path, divided by a factor of 200. This value is chosen to reflect the myelination of white matter fibers, which is known to allow cortico-cortical activity to propagate at speeds ~ 200 times greater in comparison with intra-cortical propagation [21]. Resting-state BOLD fMRI timecourses of the subject were minimally preprocessed (coregistration, motion correction), resampled on the subject connectome graph, and demeaned.

Spatiotemporal observables in numerical simulations

In numerical simulations, the spatial power spectrum is obtained following its standard definition, as the temporal mean of the squared graph Fourier transform of activity fluctuations about a steady state. The temporal power spectrum is estimated with the periodogram method, as implemented in the Scipy Python package. The correlation matrix is obtained by normalizing the covariance matrix of excitatory activity, which was itself estimated with the Numpy Python package.

Code repository

All code used for analysis and simulations is available for use and review at https://github.com/marcoaqil/Graph-Stochastic-Wilson-Cowan-Model

Discussion

In this work, we have presented a general approach to whole-brain neural activity modelling on unparcellated connectomes (*graph neural fields*), by combining tools from graph signal processing and neural field equations. We developed a technique to compute spatiotemporal observables (CHAOSS), and showed that a Wilson-Cowan stochastic graph neural field model can reproduce the empirically observed harmonic spectrum of resting-state BOLD signal fluctuations. Graph neural fields can address some limitations of existing modelling frameworks, and therefore represent a complementary approach resulting particularly suitable for mesoscopic-scale modelling and connectome-graph-based analyses. To discuss advantages and limitations of our approach, it is useful to contextualize it within the existing landscape of whole-brain models.

Existing whole-brain models can be broadly divided into two classes, according to whether they incorporate short-range intracortical connectivity or not. *Region-based models* only take into account long-range connectivity between dozens or few hundreds of macroscopic ROIs, whereas *surface-based models* directly incorporate short-range intra-cortical connectivity as well [26, 30]. It is furthermore possible to distinguish between *discrete* and *continuous* surface-based models. Discrete surface-based models are defined on a (highly-sampled) cortex and are therefore finite-dimensional. In several studies, region-based and discrete surface-based models are collectively referred to as *networks of neural masses* [17, 31, 32]. Continuous surface-based models are better known as *neural field models* and are defined on the entire cortex and are thus infinite-dimensional [23, 26, 33]. Mathematically, discrete surface-based models are finite-dimensional systems of ordinary differential equations, whereas neural field models are partial integro-differential equations.

Region-based models are constructed by parcellating the cortex into a number of regions-of-interest (ROIs), placing a local model in each ROI, and connecting them according to a given connectome (see [2, 17, 32] for reviews). The ROIs are usually obtained from structural or functional cortical atlases and the number of ROIs is in the order of a hundred or less. Connectome-based mass models are characterized by the type of local models and how they are connected i.e. if the connections are weighted or not, excitatory or inhibitory, and if transmission delays are incorporated. A wide variety of local models has been used in the literature, including neural mass models, self-sustained oscillators, chaotic deterministic systems, circuits of spiking neurons, normal-form bifurcation models, rate models, and density models [2, 17, 33]. Region-based models have proven valuable in understanding varies aspects of large-scale cortical dynamics and their roles in cognitive and perceptual processing, but they are limited in one important respect: they do not allow studying the spatiotemporal organization of cortical activity on scales smaller than several squared centimeters and their effects on large-scale pattern formation. This is due to the fact that the dynamics within ROIs are modeled by a single model without spatial extent. This prevents studying the mechanisms underlying a large class of cortical activity patterns that have been observed in experiments, including traveling and spiral waves, sink-source-dynamics as well as their role in shaping macroscopic dynamics. This is a significant limitation, particularly because the role of mesoscopic spatiotemporal dynamics in cognitive and perceptual processing is increasingly being recognized and experimentally studied [34, 35]. Graph neural fields present the advantage of allowing explicit modelling of activity propagation dynamics with spatiotemporal convolutions and graph differential equations on mesoscopic-resolution connectomes, thereby overcoming this limitation.

Whole-brain models that incorporate short-range connectivity are referred to as *surface-*

based because they can are defined either on high-resolution surface-based representations of the cortex [30, 36, 37] or on the entire cortex viewed as a continuous medium. We will refer to these types of models as *discrete* and *continuous* surface-based models, the latter of which are known as *neural field models* [23, 26, 38, 39]. Numerically simulating discrete surface-based models is much more computationally demanding than simulating region-based models as the former typically have dimensions that are one to two orders higher than those of the latter. Numerically simulating neural field models is even more demanding and requires heavy numerical integration in combination with advanced analytical techniques [40]. Moreover, simulating neural field models requires special preparation of cortical meshes to ensure accuracy and numerical stability. [36, 37, 41–43]. In this context, graph neural fields have the advantage of being implementable natively on multimodal structural connectomes obtained from MRI and DTI, thereby minimizing anatomical approximations. Graph neural fields can naturally take into account important physical properties such as cortical folding, hemispheric asymmetries, non-homogeneous structural connectivity, and white matter projections, with a minimal amount of computational power. Furthermore, the cortex in graph neural field need not be a flat or spherical bidimensional manifold. The approach can be straightforwardly extended to include cortical thickness, allowing activity to propagate not only tangentially, but also perpendicularly to the cortical surface. Cortical layers can already be distinguished with high-field and ultra-high field functional fMRI, and are thought to subserve different functions [44]. The ability of graph neural fields to account for cortical thickness in dynamical models of neural activity is therefore a promising property for future development [45].

Our approach presents several limitations. First, some of the analytic results presented here (CHAOSS) rely on the model parameters being space-independent, that is the model parameters are assumed to be the same throughout the cortex. This assumption has the advantage of allowing mathematical analyses that, unlike numerical simulations, are virtually "infinitely" scalable with little computational cost, and was also used in previous studies of continuous neural fields [18]. However, there are more biophysically realistic models that require space-dependent parameters. For example, some recent neural mass network models incorporate neuronal receptors and their densities, which are known to vary across the cortex [46–48]. There are several ways in which it might be possible to overcome this limitation. We remark that it is only our analytic approach that requires space-independent parameters; numerical simulations of graph neural fields could be carried out also with space-dependent parameters (of course, such simulations would be more computationally demanding than their counterparts with space-independent parameters). To preserve analytic tractability while characterizing regional differences, one could attempt to absorb all the relevant space-dependent information into the graph Laplacian. Similarly to the idea of differentially weighing white matter edges to account for myelination, one might weigh differentially graph edges within specific ROIs or specific subsets of vertices. A hybrid approach of space-dependent numerical simulations and space-independent analysis (by averaging values of space-dependent parameters) could be another way to address this issue. Second, we have restricted our approach to spatially symmetric kernels. In some special cases, asymmetric kernels may be practically obtained by introducing suitable asymmetries in the graph edges. For example, consider a grid graph in two dimensions, with additional edges connecting bottom-left and top-right vertices of each square in the grid. Because of the broken lattice symmetry, a Gaussian kernel on this non-grid graph will behave like a spatially elliptic Gaussian, angled at 45 degrees. This is analogous to modelling a spatially asymmetric diffusion process on the graph. Third, another limitation is the use of an undirected and time-independent connectome graph. For maximal generality and

biophysical realism, one might want to study a directed, or even time-dependent (plastic) structural connectome: it is not clear at present if and how this could be implemented in the framework of graph neural fields.

Immediate applications of graph neural fields can be found in the comparison of harmonic spectra, functional connectivity and coherence matrix with single-subject empirical data obtained from different neuroimaging modalities such as fMRI and MEG, as well as different conditions, for example health, pathology, and neuropharmacologically altered states of consciousness [22]. For example, investigating the effects of a reduced myelination speed factor, or pruned white-matter fibers, could be an interesting approach to modelling the effects of pathological or age-related structural alterations on the dynamics of functional activity. Other applications include the implementation of more biophysically realistic models, potentially including space-dependent parameters, and the use of a cortical connectome that includes cortical thickness, accounting for activity propagation across layers perpendicularly to the surface. Aside from whole-brain modelling, graph neural fields may also be used for modelling specific ROIs and stimulus-evoked brain activity. In particular, because of the known retinotopic mapping between visual stimuli and neural activity, the visual cortex presents itself as a very interesting ROI for such developments [49]. Moving beyond neural populations and even the human brain, we note that the graph Laplacian may also be used to implement single-neuron models directly on the full connectome graphs of simple organisms, such as *C. Elegans*, whose connectome has been experimentally mapped at the single-neuron level.

Conclusion

In this study we described a class of whole-brain neural activity models which we refer to as *graph neural fields*, and showed that they can be used to investigate properties of brain activity measured with neuroimaging methods. The formulation of graph neural fields relies on existing concepts from the field of graph signal processing, namely the graph Laplacian operator and graph filtering, and modelling concepts such as neural field equations. This framework allows inclusion of realistic anatomical features, analytic predictions of harmonic-temporal power spectra, correlation, and coherence matrices (CHAOSS), and efficient numerical simulations. We illustrated the practical use of the framework by reproducing the harmonic spectrum of resting-state BOLD fMRI with a stochastic Wilson-Cowan graph neural field model. Future work could build on the methods and results presented here, both from theoretical and applied standpoints.

Supporting information

S1 Table Parameter set for 1D analysis and simulations. This parameter set was obtained by a qualitative comparison of the Wilson-Cowan model's harmonic and temporal spectra with empirical data, and used to illustrate how graph properties affect neural field dynamics in one dimension.

Parameter	Value	Units (S.I)
τ_E	$2.415 \cdot 10^{-1}$	s
τ_I	$5.227 \cdot 10^{-1}$	s
σ_{EE}	$4.577 \cdot 10^{-3}$	m
σ_{IE}	$1.704 \cdot 10^{-3}$	m
σ_{EI}	$9.377 \cdot 10^{-3}$	m
σ_{II}	$2.239 \cdot 10^{-1}$	m
d_E	10^2	-
d_I	9.430	-
α_{EE}	$6.886 \cdot 10^2$	-
α_{IE}	$7.903 \cdot 10^2$	-
α_{EI}	$9.972 \cdot 10^2$	-
α_{II}	$1.223 \cdot 10^2$	-
P	3.824	-
Q	7.120	-
σ	10^{-7}	-

S2 Table Parameter set for connectome-wide analysis and simulations. This parameter set was obtained by quantiatively fitting the Wilson-Cowan model's harmonic power spectrum to that of resting-state fMRI data, and used for all connectome-wide analysis and numerical simulations.

Parameter	Value	Units (S.I)
τ_E	$2.024 \cdot 10^{-1}$	s
τ_I	$2.346 \cdot 10^{-1}$	s
σ_{EE}	$1.611 \cdot 10^{-2}$	m
σ_{IE}	$2.022 \cdot 10^{-3}$	m
σ_{EI}	$6.698 \cdot 10^{-2}$	m
σ_{II}	$9.149 \cdot 10^{-2}$	m
d_E	$2.718 \cdot 10^1$	-
d_I	1.240	-
α_{EE}	$1.487 \cdot 10^2$	-
α_{IE}	$2.191 \cdot 10^2$	-
α_{EI}	$2.620 \cdot 10^2$	-
α_{II}	$1.614 \cdot 10^2$	-
P	$2.235 \cdot 10^1$	-
Q	8.450	-
σ	10^{-7}	-

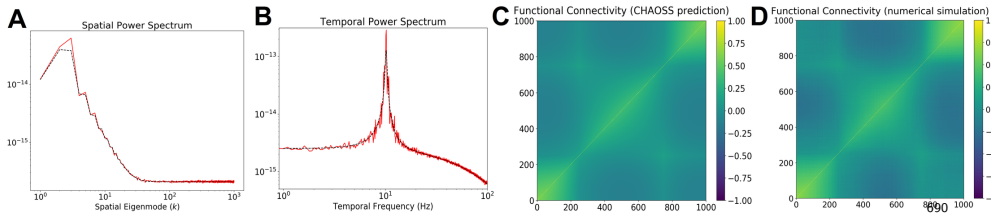

S1 Fig.
Spatial convolutions on 1-dimensional graphs. To illustrate spatial convolution on graphs, we apply different spatial convolution filters from Table 2 to an impulse function centered on the mid-node of a one-dimensional grid-graph with spacing $h = 1$ units. The resulting functions, normalized to have unit amplitude, show the shapes of the graph kernels. Note that the rectangular kernel convolution operator in Panel (E) exhibits the *Gibbs phenomenon* [50], which is a known feature of finite Fourier representations of functions with jump discontinuities. Solutions to this problem have been offered [51], but they are beyond the scope of the current work. Thus, we suggest avoiding spatial kernels with jump discontinuities in the context of graph neural fields. As a side remark, we also note that for large datasets (for example natural images databases), it might be computationally advantageous to apply convolutions with symmetric kernels through graph filters, rather than with standard discrete convolution methods. Spatial convolutions on graphs become linear matrix-vector products, which are highly optimized and easily parallelizable operations; the bulk of the computational cost for graph convolutions consists in the initial computation of the filter itself, which has to be performed only once per kernel.

Acknowledgments

We would like to thank Thomas Yeo and Ruby Kong for providing the mapping between HCP 32k and 10k vertices and Daniele Avitabile for valuable discussions. MLK is supported by the ERC Consolidator Grant: CAREGIVING (n. 615539), Center for Music in the Brain, funded by the Danish National Research Foundation (DNRF117), and Centre for Eudaimonia and Human Flourishing funded by the Pettit and Carlsberg Foundations. RH is supported by the NWO-Wiskundeclusters grant nr. 613.009.105.

References

1. Damoiseaux JS, Rombouts S, Barkhof F, Scheltens P, Stam CJ, Smith SM, et al. Consistent resting-state networks across healthy subjects. Proceedings of the national academy of sciences. 2006;103(37):13848–13853.

2. Deco G, Jirsa VK, McIntosh AR. Emerging concepts for the dynamical organization of resting-state activity in the brain. Nature Reviews Neuroscience. 2011;12(1):43.

3. Lim S, Radicchi F, van den Heuvel MP, Sporns O. Discordant attributes of structural and functional connectivity in a two-layer multiplex network. bioRxiv. 2018; p. 273136.

4. Mohar B, Alavi Y, Chartrand G, Oellermann O. The Laplacian spectrum of graphs. Graph theory, combinatorics, and applications. 1991;2(871-898):12.

5. Sandryhaila A, Moura JM. Discrete signal processing on graphs. IEEE transactions on signal processing. 2013;61(7):1644–1656.

6. Shuman DI, Narang SK, Frossard P, Ortega A, Vandergheynst P. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. IEEE Signal Processing Magazine. 2013;30(3):83–98.

7. Perraudin N, Vandergheynst P. Stationary signal processing on graphs. IEEE Transactions on Signal Processing. 2017;65(13):3462–3477.

8. Ortega A, Frossard P, Kovačević J, Moura JM, Vandergheynst P. Graph signal processing: Overview, challenges, and applications. Proceedings of the IEEE. 2018;106(5):808–828.

9. Atasoy S, Donnelly I, Pearson J. Human brain networks function in connectome-specific harmonic waves. Nature communications. 2016;7:10340.

10. Atasoy S, Deco G, Kringelbach ML, Pearson J. Harmonic brain modes: a unifying framework for linking space and time in brain dynamics. The Neuroscientist. 2018;24(3):277–293.

11. Huang W, Bolton TA, Medaglia JD, Bassett DS, Ribeiro A, Van De Ville D. A Graph Signal Processing Perspective on Functional Brain Imaging. Proceedings of the IEEE. 2018;.

12. Xu G. Discrete Laplace–Beltrami operators and their convergence. Computer aided geometric design. 2004;21(8):767–784.

13. Belkin M, Sun J, Wang Y. Discrete laplace operator on meshed surfaces. In: Proceedings of the twenty-fourth annual symposium on Computational geometry; 2008. p. 278–287.

14. Wang MB, Owen JP, Mukherjee P, Raj A. Brain network eigenmodes provide a robust and compact representation of the structural connectome in health and disease. PLoS computational biology. 2017;13(6):e1005550.

15. Tewarie P, Abeysuriya R, Byrne Á, O'Neill GC, Sotiropoulos SN, Brookes MJ, et al. How do spatially distinct frequency specific MEG networks emerge from one underlying structural connectome? The role of the structural eigenmodes. NeuroImage. 2019;186:211–220.

16. Raj A, Cai C, Xie X, Palacios E, Owen J, Mukherjee P, et al. Spectral graph theory of brain oscillations. Human Brain Mapping. 2020;.

17. Breakspear M. Dynamic models of large-scale brain activity. Nature neuroscience. 2017;20(3):340.

18. Coombes S. Waves, bumps, and patterns in neural field theories. Biological cybernetics. 2005;93(2):91–108.

19. Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. The Journal of physiology. 1952;117(4):500–544.

20. Cowan JD, Neuman J, van Drongelen W. Wilson–Cowan equations for neocortical dynamics. The Journal of Mathematical Neuroscience. 2016;6(1):1.

21. Purves D, Augustine G, Fitzpatrick D, Katz L, LaMantia A, McNamara J, et al. Increased conduction velocity as a result of myelination. Neuroscience. 2001;.

22. Atasoy S, Roseman L, Kaelen M, Kringelbach ML, Deco G, Carhart-Harris RL. Connectome-harmonic decomposition of human brain activity reveals dynamical repertoire re-organization under LSD. Scientific reports. 2017;7(1):17661.

23. Bressloff PC. Spatiotemporal dynamics of continuum neural fields. Journal of Physics A: Mathematical and Theoretical. 2011;45(3):033001.

24. Wilson HR, Cowan JD. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. Kybernetik. 1973;13(2):55–80.

25. Negahbani E, Steyn-Ross DA, Steyn-Ross ML, Wilson MT, Sleigh JW. Noise-induced precursors of state transitions in the stochastic Wilson–Cowan model. The Journal of Mathematical Neuroscience (JMN). 2015;5(1):9.

26. Coombes S. Large-scale neural dynamics: simple and complex. NeuroImage. 2010;52(3):731–739.

27. Visser S, Nicks R, Faugeras O, Coombes S. Standing and travelling waves in a spherical brain model: the Nunez model revisited. Physica D: Nonlinear Phenomena. 2017;349:27–45.

28. Pereda E, Quiroga RQ, Bhattacharya J. Nonlinear multivariate analysis of neurophysiological signals. Progress in neurobiology. 2005;77(1-2):1–37.

29. Nolte G, Galindo-Leon E, Li Z, Liu X, Engel AK. Mathematical relations between measures of brain connectivity estimated from electrophysiological recordings for Gaussian distributed data. bioRxiv. 2019; p. 680678.

30. Proix T, Spiegler A, Schirner M, Rothmeier S, Ritter P, Jirsa VK. How do parcellation size and short-range connectivity affect dynamics in large-scale brain network models? NeuroImage. 2016;142:135–149.

31. Sanz-Leon P, Knock SA, Spiegler A, Jirsa VK. Mathematical framework for large-scale brain network modelling in The Virtual Brain. Neuroimage. 2015;111:385–430.

32. Byrne Á, O'Dea RD, Forrester M, Ross J, Coombes S. Next-generation neural mass and field modelling. Journal of Neurophysiology. 2020;123(2):726–742.

33. Deco G, Jirsa VK, Robinson PA, Breakspear M, Friston K. The dynamic brain: from spiking neurons to neural masses and cortical fields. PLoS Comput Biol. 2008;4(8):e1000092.

34. Roberts JA, Gollo LL, Abeysuriya RG, Roberts G, Mitchell PB, Woolrich MW, et al. Metastable brain waves. Nature communications. 2019;10(1):1056.

35. Muller L, Chavane F, Reynolds J, Sejnowski TJ. Cortical travelling waves: mechanisms and computational principles. Nature Reviews Neuroscience. 2018;19(5):255.

36. Bojak I, Oostendorp TF, Reid AT, Kötter R. Connecting mean field models of neural activity to EEG and fMRI data. Brain topography. 2010;23(2):139–149.

37. Bojak I, Oostendorp TF, Reid AT, Kötter R. Towards a model-based integration of co-registered electroencephalography/functional magnetic resonance imaging data with realistic neural population meshes. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences. 2011;369(1952):3785–3801.

38. Robinson P, Rennie C, Wright J, Bourke P. Steady states and global dynamics of electrical activity in the cerebral cortex. Physical Review E. 1998;58(3):3557.

39. Liley DT, Cadusch PJ, Dafilis MP. A spatially continuous mean field theory of electrocortical activity. Network: Computation in Neural Systems. 2002;13(1):67–113.

40. Martin R. Collocation techniques for solving neural field models on complex cortical geometries. Nottingham Trent University; 2018.

41. Freestone DR, Aram P, Dewar M, Scerri K, Grayden DB, Kadirkamanathan V. A data-driven framework for neural field modelling. NeuroImage. 2011;56(3):1043–1058.

42. Spiegler A, Jirsa V. Systematic approximations of neural fields through networks of neural masses in the virtual brain. Neuroimage. 2013;83:704–725.

43. Sanz Leon P, Knock SA, Woodman MM, Domide L, Mersmann J, McIntosh AR, et al. The Virtual Brain: a simulator of primate brain network dynamics. Frontiers in neuroinformatics. 2013;7:10.

44. Lawrence SJ, Norris DG, de Lange FP. Dissociable laminar profiles of concurrent bottom-up and top-down modulation in the human visual cortex. Elife. 2019;8:e44422.

45. Kuehn E, Sereno MI. Modelling the human cortex in three dimensions. Trends in cognitive sciences. 2018;22(12):1073–1075.

46. Deco G, Cruzat J, Cabral J, Knudsen GM, Carhart-Harris RL, Whybrow PC, et al. Whole-brain multimodal neuroimaging model using serotonin receptor maps explains non-linear functional effects of LSD. Current Biology. 2018;28(19):3065–3074.

47. Kringelbach ML, Cruzat J, Cabral J, Knudsen GM, Carhart-Harris R, Whybrow PC, et al. Dynamic coupling of whole-brain neuronal and neurotransmitter systems. Proceedings of the National Academy of Sciences. 2020;117(17):9566–9576.

48. Herzog R, Mediano PA, Rosas FE, Carhart-Harris R, Sanz Y, Tagliazucchi E, et al. A mechanistic model of the neural entropy increase elicited by psychedelic drugs. bioRxiv. 2020;.

49. Dumoulin SO, Wandell BA. Population receptive field estimates in human visual cortex. Neuroimage. 2008;39(2):647–660.

50. Hewitt E, Hewitt RE. The Gibbs-Wilbraham phenomenon: an episode in Fourier analysis. Archive for history of Exact Sciences. 1979;21(2):129–160.

51. Gottlieb D, Shu CW. On the Gibbs phenomenon and its resolution. SIAM review. 1997;39(4):644–668.