Does Connectome Harmonic Analysis pass the Spin Test?

Raphaël Vock[1,2], Antoine Grigis[1], Benoît Dufumier[1], Edouard Duchesnay[1]

[1]Université Paris–Saclay, CEA, CNRS, UMR9027 Baobab, NeuroSpin, Saclay, France
[2]raphaelvock@gmail.com

Abstract. Connectome harmonic analysis has been proposed as a multimodal approach for studying brain dynamics by decomposing functional MRI signals in a Fourier basis informed by the structural connectome derived from diffusion MRI. In this work we pose the following question: is the propensity of the connectomic Fourier basis to reconstruct resting state fMRI signals truly contingent upon anatomical priors? We present evidence that it is not, by demonstrating that when fewer than $n = 100$ modes are considered the connectomic eigenbasis obtained through state-of-the-art methodology performs similarly to geometrically transformed versions of that same basis. The main theoretical contribution of this paper is the construction of a regular planar embedding of the left hemisphere's cortical surface, which we use to compute a smoothly parametrised family of cortical transformations which form the basis for an improved Spin Test.

Keywords: Connectome harmonic analysis · Diffusion MRI · Functional MRI · Multimodal data integration

1 Introduction

Connectome Harmonic Analysis (CHA) [2] is a multimodal paradigm which seeks to understand brain dynamics via the decomposition of functional MRI (fMRI) signals in a Fourier basis informed by the structural connectome derived from diffusion MRI. By borrowing tools from traditional signal processing, CHA is straightforward to compute, nonparametric, builds upon physical principles for brain dynamics, and demonstrates the ability to accurately compress fMRI signals using a relatively small number of modes (fewer than 1% of the number of dimensions in the original signal). It has been rapidly propelled to the forefront of structure–function coupling models of brain dynamics. Some applications have been given towards the study of consciousness [6] and psychedelic-induced brain state alterations [3, 10].

Pang and colleagues [7] have cast doubts upon the meaningfulness of these connectomic modes by showing that an eigenbasis obtained using the cortical surface's intrinsic geometry does just as well as the connectomic eigenbasis, in terms of either basis's ability to reconstruct fMRI data from a limited number of modes. In the present work we use the well-known Spin Test framework [1] to

provide further evidence for CHA's lack of robustness. The idea of the Spin Test is to stress-test a given fMRI methodology by observing how certain statistics respond to the remapping of the input data via transformations of the cortical surface.

In particular, we seek to answer the following question: does the connectome-informed Fourier eigenbasis induce more efficient representations of resting state fMRI data than that same basis transformed via a continuous remapping of the cortical surface? We demonstrate that the eigenbasis obtained through state-of-the-art methodology [8] does not outperform rotations of that basis. This raises questions as to the robustness of CHA, at least if we are to consider the reconstruction accuracy metric as meaningful. The main methodological contribution of this paper is the development of an improved version of the Spin Test parametrised by a continuous parameter which demonstrably overcomes the shortcomings of the original test.

Previous work has been dedicated to applying the Spin Test to CHA in various other contexts, e.g. [6], but to our knowledge this is the first attempt to study the reconstruction accuracy metric in terms of that framework.

2 Materials and method

2.1 HCP dataset

Our work makes use of the Human Connectome Project (HCP) dataset for both diffusion MRI data, used in the construction of the connectomic eigenbasis (§2.4), and resting state fMRI data used to evaluate bases via the reconstruction accuracy metric (§2.3).

2.2 Functional decomposition bases

From the FreeSurfer [4] population-averaged, high-resolution template for the left hemisphere's cortical surface (fsLR 32k), we remove the interhemispheric barrier in order to obtain a 29,696-vertex triangle mesh corresponding to the idealised geometry of the adult human, left hemisphere, mid-thickness cortical surface. We denote this surface using the symbol $\mathcal{C}$. A resting state fMRI signal on the cortical surface can be represented as a function $f(\mathbf{x}, t)$ of both space and time, where $\mathbf{x}$ denotes a point on the idealised cortical surface of the left hemisphere $\mathcal{C}$, and t is some time point in a fixed interval. The motivation for CHA resides in the need for an anatomically informed linear basis $(\mathbf{e}_1, \ldots, \mathbf{e}_N)$ of $\mathbb{R}^{\mathcal{C}}$, i.e. the space of functional signals at a given time. Note that the basis vectors $\mathbf{e}_i$ are time- and subject-independent, and in particular are computed generically with no prior knowledge of the fMRI signal distribution. A given time-dependent fMRI signal can subsequently be decomposed in that eigenbasis as $f(\mathbf{x}, t) = \sum_{i=1}^{N} c_i(t)\mathbf{e}_i$ where the coefficient $c_i(t)$ depends *only* on time. We may assume without loss of generality that the $\mathbf{e}_i$ are orthonormal, such that the coefficient $c_i(t)$ can be explicitly obtained using the standard inner product over $\mathbb{R}^{\mathcal{C}}$ as $c_i(t) = \langle f(\bullet, t), \mathbf{e}_i \rangle$.

2.3 Evaluation metric

Following [2, 7, 8] we evaluate the quality of a given basis using the reconstruction accuracy metric which is defined as follows. For each resting state fMRI signal $f(\mathbf{x}, t)$ corresponding to one of the 252 subjects used by [7] in the Human Connectome Project database, we compute the truncated reconstruction signal $f_n(\mathbf{x}, t) = \sum_{i=1}^{n} c_i(t)\mathbf{e}_i$ containing the n first terms of the signal in that basis. The reconstructed signal is projected (by means of spatial averaging) onto the 180 regions of interest (ROIs) in the left hemisphere as defined by Glasser and colleagues [5].

By taking the pairwise Pearson correlation coefficient of the truncated reconstruction on each ROI one obtains an 180×180 correlation matrix $\mathbf{R}'$, which is known as the truncated functional connectivity matrix. Denoting by $\mathbf{R}$ the (full) functional connectivity matrix obtained similarly from the original signal $f(\mathbf{x}, t)$, one can compare $\mathbf{R}$ and $\mathbf{R}'$ in order to gauge the accuracy of the reconstructed signal. In particular, we compute the correlation between the sequences $\mathbf{R}_{i<j}$, $\mathbf{R}'_{i<j}$, i.e. the $180 \times 179/2$ coefficients in the lower triangle portion (we consider only the lower triangle portions of these matrices since they are symmetric and have units on the diagonal). In summary we define $r := \mathrm{Corr}(\mathbf{R}_{i<j}, \mathbf{R}'_{i<j})$, henceforth known as the reconstruction accuracy, which depends on the choice of basis $(\mathbf{e}_1, \ldots, \mathbf{e}_N)$ as well as on n, the number of modes.

2.4 Connectomic eigenmodes

The following idea, due to Atasoy and colleagues [2], is known as CHA, and consists of introducing diffusion MRI data in order to inform the choice of a functional decomposition basis $(\mathbf{e}_1, \ldots, \mathbf{e}_N)$.

From the HCP diffusion MRI dataset, Mansour and colleagues [8] derive a population-averaged structural connectome which applies spatial smoothing and correction for gyral bias. This connectome induces a weighted graph structure of $\mathcal{C}$ whose (weighted, degree-normalised) graph Laplacian $\triangle$ we go on to compute. Since it is a symmetric and positive semidefinite matrix, it has $N = |\mathcal{C}|$ eigenvectors $\mathbf{e}_1, \ldots, \mathbf{e}_N$ of $\triangle$, associated with the increasing sequence of nonnegative eigenvalues $0 \leqslant \lambda_1 \leqslant \cdots \leqslant \lambda_N$, which form an orthogonal basis for $\mathbb{R}^{\mathcal{C}}$, i.e. the space of all possible functional signals at a fixed time point. Note that the first few eigenvalues intuitively correspond to patterns with low spatial frequency, and as n increases higher frequency information tends to be captured by $f_n(\mathbf{x}, t)$. Note also that when $n = N$, $f_N(\mathbf{x}, t)$ perfectly reconstructs the original signal.

2.5 Spin Test

In the case of CHA we can formulate the Spin Test as follows: if $\rho : \mathcal{C} \to \mathcal{C}$ is a (bijective) transformation of the cortical surface, how accurately does the transformed eigenbasis $(\rho\mathbf{e}_1, \ldots, \rho\mathbf{e}_N)$ reconstruct resting state fMRI signals compared with the original basis $(\mathbf{e}_1, \ldots, \mathbf{e}_N)$?

As originally formulated in [1], the Spin Test framework makes use of a spherical projection of the cortical surface $\sigma : \mathcal{C} \to S^2$ where $S^2 := \left\{\mathbf{x} \in \mathbb{R}^3 \mid \|\mathbf{x}\| = 1\right\}$ denotes the (hollow) 2-sphere. A random orthogonal transformation $Q : \mathbb{R}^3 \to \mathbb{R}^3$ is applied to that embedding and a nearest neighbour search is used to match points between the embedding $\sigma(\mathcal{C})$ and the transformed embedding $Q(\sigma(\mathcal{C}))$. The crudeness of this method leads to several shortcomings:

1. The resulting transformation $\rho : \mathcal{C} \to \mathcal{C}$ need not be bijective because two points in the embedding may well be mapped to the same nearest neighbour.
2. Consequently, the transformed basis $(\rho\mathbf{e}_1, \ldots, \rho\mathbf{e}_N)$ need not be orthogonal (or indeed even a basis properly speaking).
3. Since the interhemispheric barrier is embedded onto an irregular patch of the sphere, the resulting transformation $\rho : \mathcal{C} \to \mathcal{C}$ violates spatial continuity.
4. The variance due to the random orthogonal transformation Q lacks an explicit parametrization.

We overcome these shortcomings using a precise and elegant alternative based on a regular planar embedding $\phi : \mathcal{C} \to D^2$ where $D^2 = \left\{\mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x}\| \leqslant 1\right\}$ denotes the 2-dimensional unit disc. By exploiting the symmetries of the D^2 we subsequently obtain a family of maps ρ_θ smoothly parametrised by an angle $\theta \in [0, 2\pi)$. The construction of ϕ is given as follows. Since $\mathcal{C}$ has the structure a triangle mesh, it automatically inherits a graph structure. Note that $\mathcal{C}$ has the topology of the hemispherical shell whose boundary (i.e. equator) corresponds anatomically to the interhemispheric barrier. By associating these points in the appropriate order to the vertices of regular 243-gon centred in the origin of $\mathbb{R}^2$ one can apply the Tutte Spring Embedding theorem [9] so as to obtain a planar embedding of $\mathcal{C}$ in the unit disc. In order to promote uniformity of the embedding we use springs with nonzero free length $\ell = 1/\sqrt{N} \approx 5.8 \times 10^{-3}$, and after the convergence of the physical model we apply a sweeping radius regularisation step. We illustrate this sequence of transformations in Fig. 1. More detailed descriptions of the algorithms are given in §5.

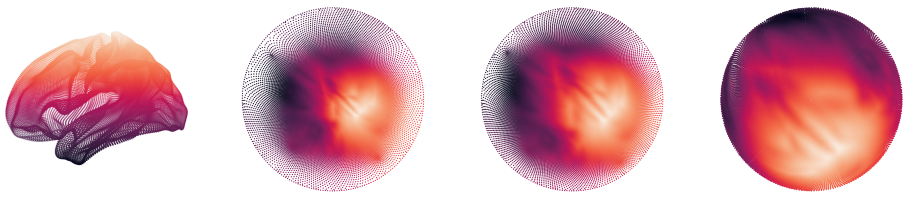

Fig. 1: From left to right, colour-coded by z-coordinate in brain space: (1) vertices of $\mathcal{C}$; (2) Tutte embedding of $\mathcal{C}$; (3) Tutte-like spring embedding with nonzero free spring length; (4) Spring-based embedding after radial regularisation

By passing through to D^2 via ϕ one easily computes a family of transformations of $\mathcal{C}$ parametrised by an angle $\theta \in [0, 2\pi)$. Indeed, each angle θ gives

rise to a transformation of D^2 in the form of a rotation ρ_θ by θ radians about the origin. One pulls back to a transformation of $\mathcal{C}$ (also denoted ρ_θ) via the composition $\phi^{-1} \circ \rho_\theta \circ \phi$, as illustrated in the following diagram.

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{\phi} & D^2 \\ {\scriptstyle \rho_\theta}\downarrow & & \downarrow{\scriptstyle \rho_\theta} \\ \mathcal{C} & \xleftarrow[\phi^{-1}]{} & D^2 \end{array}$$

In practice, since we lack an explicit form for ϕ^{-1}, we compute $\phi^{-1} \circ \rho_\theta \circ \phi$ by solving the Euclidean linear assignment problem (LAP) $(\rho_\theta \circ \phi(\mathbf{x}))_{\mathbf{x}\in\mathcal{C}} \rightsquigarrow (\phi(\mathbf{x}))_{\mathbf{x}\in\mathcal{C}}$. Owing to the computational complexity of the LAP (cubic in $N = 29{,}696$) we partition the embedding space $f(\mathcal{C})$ into three concentric annuli in order to apply a divide-and-conquer strategy by solving the associated LAP on each annulus. Since ρ_θ preserves annuli this yields a fair approximation to the exact solution.

With the transformation ρ_θ computed we can compute the transformed basis $(\rho_\theta \mathbf{e}_1, \ldots, \rho_\theta \mathbf{e}_N)$. To correct the artifacts in the computation of ρ_θ we select those points with assignment cost (in terms of the previous LAP step) greater than a threshold $\alpha = 0.1$ and replace their value by the average of their neighbors on the triangular mesh. We then apply a global smoothing process by repeatedly applying a 10-fold neighbour averaging process, and finally we orthogonalise the resulting basis using the Gram–Schmidt process.

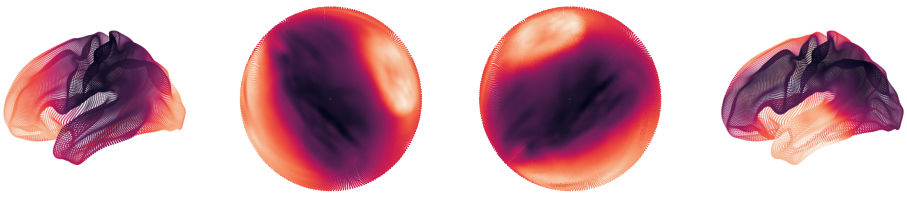

Fig. 2: From left to right: (1) third connectome eigenmode $\mathbf{e}_3$ projected in brain space; (2) $\mathbf{e}_3$ in embedding space; (3) $\mathbf{e}_3$ rotated by $\pi/2$ and smoothed; (4) $\mathbf{e}_3$ rotated and smoothed, in brain space

3 Results

Using the method described in §2.5 we compute, for 243 regularly spaced angles $\theta \in [0, 2\pi)$, the transformed connectomic eigenbasis $(\rho_\theta \mathbf{e}_1, \ldots, \rho_\theta \mathbf{e}_N)$. For each of the 252 preprocessed resting state fMRI sequences used to compute the average structural connectome, we compute the values of the reconstruction accuracy $r = r(n, \theta)$ as a function of both the number of modes $n \in \{25, 50, 75, 100\}$ and the transformation parameter θ, using the formula described in §2.3. We plot the effect of θ on r for the selected values of n in the left plot in Fig. 3. For fixed n and θ, treating subjects as random effects, we compute the p-value of the

t-test corresponding to the statistical hypothesis $r(0, n) > r(\theta, n)$; we plot the results on the right side of the same figure. The legend includes an estimation of the proportion of those angles θ with $p > 0.05$, as a function of n, i.e. the probability that the connectomic eigenbasis transformed through a random cortical transformation not be significantly worse than the untransformed eigenbasis.

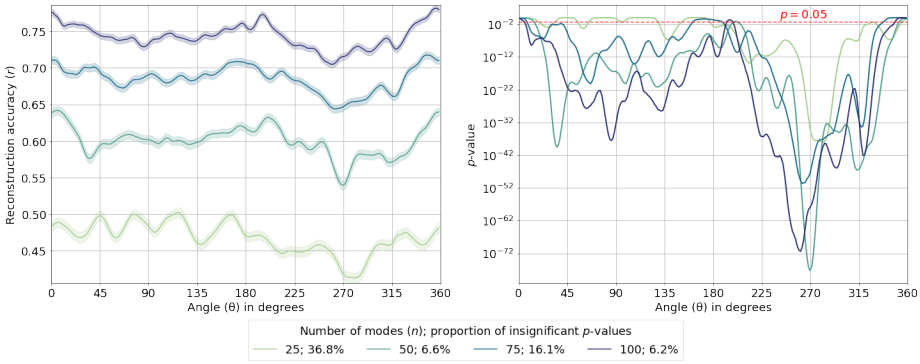

Fig. 3: Left: reconstruction accuracy as a function of the parameter θ of the cortical surface transformations, for selected values of n. Right: p-value corresponding to the t-test for $r(0, n) > r(\theta, n)$ as a function of θ, for fixed n. The dashed red line corresponds to the significance threshold $p = 0.05$. The legend includes an estimation, for each n, of the proportion of those values of θ with insignificant p-values.

4 Discussion

Were the connectomic eigenbasis to be truly contingent upon anatomical priors, one would expect an arbitrary rotation of the cortical surface to induce a sharp drop-off in the prescribed metric. This does not appear to be the case (left plot of Fig. 3), and our intuition is rigorously confirmed by the plot of the p-values corresponding to the t-test for the statistical hypothesis $r(0, n) > r(\theta, n)$. Indeed, a non-negligible proportion of the transformed eigenbases rotated by a nonzero angle θ yield reconstruction accuracies that are not significantly inferior to the untransformed eigenbasis. The proportion of those angles are reported in the legend of Fig. 3 and in all cases one can see that this proportion is between 6 and 37%.

Our findings seem to suggest that in the scope of this study, that is to say the reconstruction of resting state fMRI data in a small number of connectomic modes, the CHA theory does not pass the Spin Test. In other words, the reconstruction accuracy in terms of one commonly accepted metric does not appear to be truly contingent upon the anatomical data and instead results from generic analytic properties of that basis.

We emphasise the limited scope of our study, and the fact that additional tests should be carried out e.g. on task-evoked data to confirm this hypothesis. Moreover, our conclusion breaks down when the number of modes increases ($n \geqslant 200$) which suggests that CHA encodes anatomically relevant information in the high-frequency modes. Further investigation over more diverse cohorts would constitute an essential next step to confirming our results. Also, the reconstruction metric employed in this paper is a somewhat crude measure of an eigenbasis's meaningfulness and the understanding of this topic would benefit from the elaboration of more sophisticated metrics.

5 Appendix: embedding algorithms

5.1 Tutte embedding

Let $\mathcal{G}$ denote the graph structure on $\mathcal{C}$ induced by the triangular mesh (*not* by the population-averaged connectome). The Tutte embedding [9] of $\mathcal{C}$ is defined by the following rules:

1. The points on the interhemispheric barrier are mapped in an appropriate order to the boundary of the regular unit polygon centred in the origin.
2. Each point not on the boundary is mapped to the barycentre of its neighbours.

An *ad hoc* way of determining the 243 points that make up the interhemispheric boundary is by finding the largest connected component of the subgraph of vertices of degree 5 and less. The Tutte embedding can then be computed by solving two sparse linear systems (one for each dimension of the unit disc) involving the Laplacian matrix $\triangle$ of $\mathcal{G}$. Theoretical properties of this embedding, such as a sufficient condition for planarity, are discussed in [9]. We make no effort to verify this condition in our case as the resulting embedding is manifestly planar.

5.2 Spring embedding

It is shown in [9] that the Tutte embedding is the equilibrium state of the physical system given by attaching springs (with some fixed spring constant k and free length $\ell = 0$) along each edge of the graph, while "pinning" the interhemispheric boundary to the boundary of the unit polygon. By letting the free length ℓ be positive one can promote uniformity in the unit disc and remove some of the "clumps" that appear in the original embedding.

In this case, the equilibrium of the system cannot be obtained as the solution of linear systems in the data, as is the case with the Tutte embedding. Instead, one must numerically integrate the system of spring forces. It is helpful to set the initial values to the Tutte embedding to promote faster convergence. Setting the springs' free length to $\ell = 1/\sqrt{N}$, where N is the number of points in $\mathcal{C}$, produces satisfactory results.

5.3 Radial regularisation

In a final attempt to promote uniformity in the unit disc, we supplement the spring embedding with a subsequent radial regularisation step. The idea is to traverse the embedding with a "sliding sector" and to uniformise the radii within each sector while uniformising polar angles globally. This is described more precisely in the following pseudocode. In our application we set sliding sector size parameter α to 0.3.

Algorithm 1 Radial regularisation

Require: Points $\mathbf{x}_1, \ldots, \mathbf{x}_N \in \mathbb{R}^2$, sliding sector size parameter $\alpha \in (0, 1]$
Ensure: Radially regularised points $\mathbf{y}_1, \ldots, \mathbf{y}_N \in \mathbb{R}^2$
$\mu \leftarrow (1/N) \sum_{i=1}^{N} \mathbf{x}_i$ {Barycentre}
$\mathbf{c} \leftarrow \arg\min_{i=1}^{N} \|\mathbf{x}_i - \mu\|$ {Centre point}
$n \leftarrow \lfloor \alpha \cdot N \rfloor$ {Size of sliding sector}
$(r_i, \theta_i) \leftarrow \text{polar}(\mathbf{x}_i - \mathbf{c})$ for $i = 1, \ldots, N$ {Convert to polar coordinates}
$\sigma \leftarrow \arg\text{sort}_{i=1}^{N}(\theta_i)$ {Sort by angle}
$r'_j \leftarrow r_{\sigma(j)}$ for $j = 1, \ldots, N$
for $k = 1$ to N **do**
  $\widehat{\theta'} \leftarrow 2\pi k/(N-1)$ {Uniformised angle}
  if $r'_k = 0$ **then**
    $\widehat{r'} \leftarrow 0$
  else
    $\rho \leftarrow \#\{j \in [n] \mid \{r'_{(k-\lfloor n/2 \rfloor + j) \bmod N} \leqslant r'_k\}$ {Radius rank within sector}
    $\widehat{r'} \leftarrow \sqrt{\rho/n}$ {Uniformised radius in $(0, 1]$}
  end if
  $\mathbf{y}_{\sigma(k)} \leftarrow \widehat{r'} \cdot (\cos(\widehat{\theta'}), \sin(\widehat{\theta'}))$
end for
return $\mathbf{y}_1, \ldots, \mathbf{y}_N$

Code availability. The code needed to replicate this study is made available in the form of an online repository accessible through the following persistent object identifier: https://doi.org/10.5281/zenodo.15748862.

Acknowledgments. This research was generously supported by the Robert-Debré Child Brain Institute (Paris) under grant ANR-23-IAHU-0010 and the research program in precision psychiatry (PEPR PROPSY, ANR-22-EXPR-0001), both of which are funded by the France 2030 program and the French National Research Agency (ANR). Additional support was provided by two "Investissements d'Avenir" Hospital-University Research in Health initiatives: RHU-PsyCARE (ANR-18-RHUS-0014) and FAME (ANR-21-RHUS-0009). Data were provided by the Human Connectome Project, MGH-USC Consortium (Principal Investigators: Bruce R. Rosen, Arthur W. Toga and Van Wedeen; U01MH093765) funded by the NIH Blueprint Initiative for Neuroscience Research grant; the National Institutes of Health grant P41EB015896; and the Instrumentation Grants S10RR023043, 1S10RR023401, 1S10RR019307.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alexander-Bloch, A.F., Shou, H., Liu, S., Satterthwaite, T.D., Glahn, D.C., Shinohara, R.T., Vandekar, S.N., Raznahan, A.: On testing for spatial correspondence between maps of human brain structure and function. Neuroimage **178**, 540–551 (2018)
2. Atasoy, S., Donnelly, I., Pearson, J.: Human brain networks function in connectome-specific harmonic waves. Nature communications **7**(1), 10340 (2016)
3. Atasoy, S., Roseman, L., Kaelen, M., Kringelbach, M.L., Deco, G., Carhart-Harris, R.L.: Connectome-harmonic decomposition of human brain activity reveals dynamical repertoire re-organization under lsd. Scientific reports **7**(1), 17661 (2017)
4. Fischl, B.: Freesurfer. Neuroimage **62**(2), 774–781 (2012)
5. Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., et al.: A multimodal parcellation of human cerebral cortex. Nature **536**(7615), 171–178 (2016)
6. Luppi, A.I., Vohryzek, J., Kringelbach, M.L., Mediano, P.A., Craig, M.M., Adapa, R., Carhart-Harris, R.L., Roseman, L., Pappas, I., Peattie, A.R., et al.: Distributed harmonic patterns of structure-function dependence orchestrate human consciousness. Communications biology **6**(1), 117 (2023)
7. Pang, J.C., Aquino, K.M., Oldehinkel, M., Robinson, P.A., Fulcher, B.D., Breakspear, M., Fornito, A.: Geometric constraints on human brain function. Nature **618**(7965), 566–574 (2023)
8. Sina Mansour, L., Behjat, H., De Ville, D.V., Smith, R.E., Yeo, B.T., Zalesky, A.: Eigenmodes of the brain: revisiting connectomics and geometry. bioRxiv (2024). https://doi.org/10.1101/2024.04.16.589843, https://www.biorxiv.org/content/early/2024/04/20/2024.04.16.589843
9. Tutte, W.T.: How to draw a graph. Proceedings of the London Mathematical Society **3**(1), 743–767 (1963)
10. Vohryzek, J., Luppi, A.I., Atasoy, S., Deco, G., Carhart-Harris, R.L., Timmermann, C., Kringelbach, M.L.: Time-resolved coupling between connectome harmonics and subjective experience under the psychedelic dmt. bioRxiv (2024)