

This is a repository copy of *Connectopic mapping techniques do not reflect functional gradients in the brain*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/201613/>

Version: Published Version

Article:

Watson, David Mark and Andrews, Tim orcid.org/0000-0001-8255-9120 (2023)
Connectopic mapping techniques do not reflect functional gradients in the brain.
Neuroimage. 120228. ISSN 1053-8119

<https://doi.org/10.1016/j.neuroimage.2023.120228>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Connectopic mapping techniques do not reflect functional gradients in the brain

David M. Watson*, Timothy J. Andrews

Department of Psychology and York Neuroimaging Centre, University of York, York YO10 5DD, UK

ARTICLE INFO

Keywords:

Connectopic mapping
Neural gradients
Functional connectivity

ABSTRACT

Functional gradients, in which response properties change gradually across a brain region, have been proposed as a key organising principle of the brain. Recent studies using both resting-state and natural viewing paradigms have indicated that these gradients may be reconstructed from functional connectivity patterns via “connectopic mapping” analyses. However, local connectivity patterns may be confounded by spatial autocorrelations artificially introduced during data analysis, for instance by spatial smoothing or interpolation between coordinate spaces. Here, we investigate whether such confounds can produce illusory connectopic gradients. We generated datasets comprising random white noise in subjects’ functional volume spaces, then optionally applied spatial smoothing and/or interpolated the data to a different volume or surface space. Both smoothing and interpolation induced spatial autocorrelations sufficient for connectopic mapping to produce both volume- and surface-based local gradients in numerous brain regions. Furthermore, these gradients appeared highly similar to those obtained from real natural viewing data, although gradients generated from real and random data were statistically different in certain scenarios. We also reconstructed global gradients across the whole-brain – while these appeared less susceptible to artificial spatial autocorrelations, the ability to reproduce previously reported gradients was closely linked to specific features of the analysis pipeline. These results indicate that previously reported gradients identified by connectopic mapping techniques may be confounded by artificial spatial autocorrelations introduced during the analysis, and in some cases may reproduce poorly across different analysis pipelines. These findings imply that connectopic gradients need to be interpreted with caution.

1. Introduction

Functional gradients, marked by a topographic map of preferred stimulus or response features across brain regions, are proposed as a key organising principle of the brain (Bednar and Wilson, 2016). For instance, gradients have been described in the form of retinotopic maps in visual cortex (Wandell and Winawer, 2011), tonotopic maps in auditory cortex (Formisano et al., 2003), and somatotopic maps in somatosensory cortex (Sanchez-Panchuelo et al., 2010). Recently, “connectopic mapping” techniques have been proposed for recovering functional gradients from patterns of neural connectivity (Haak and Beckmann, 2020; Huntenburg et al., 2018). These approaches regard voxels or vertices within a brain region as being distributed along a high-dimensional manifold within the connectivity space, such that functionally similar locations are represented close to one another on the manifold surface. The functional gradients underlying a brain region can thus be revealed by applying non-linear manifold learning techniques to extract the principal dimensions of this manifold.

Connectopic mapping offers the advantage that gradients can be described even in brain regions where the underlying functional organisation is poorly understood or unknown. Such techniques have been applied to reconstruct local gradients within brain regions including the primary visual cortex (Haak et al., 2018; Watson and Andrews, 2022), somatosensory cortex (Ngo et al., 2021), entorhinal cortex (Navarro Schröder et al., 2015), insula (Wang et al., 2022), hippocampus (Borne et al., 2023; Przeździk et al., 2019; Vos de Wael et al., 2018), and striatum (Marquand et al., 2017). Connectopic mapping has also been applied to describe coarser-scale gradients over the whole brain (Huntenburg et al., 2018; Margulies et al., 2016).

Recently, we demonstrated that connectopic mapping could be used to accurately reconstruct the eccentricity and polar angle maps in primary visual cortex (V1) using connectivity measured either at rest or during movie watching (Watson and Andrews, 2022). Interestingly, we observed that the neural source of the connectivity estimates had minimal impact on the resulting connectopic maps, implying that the main driver in connectopic mapping is the functional topography within the

* Corresponding author.

E-mail address: david.watson@york.ac.uk (D.M. Watson).

<https://doi.org/10.1016/j.neuroimage.2023.120228>.

Received 4 April 2023; Received in revised form 8 June 2023; Accepted 12 June 2023

Available online 18 June 2023.

1053-8119/© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

brain region, rather than necessarily reflecting a topography embedded within the connectivity itself. We were able to confirm the accuracy of the connectopic maps in V1 by comparing them to retinotopic maps measured through traditional visual field mapping techniques. Nevertheless, the ground-truth functional organisation of many other brain regions is unknown or poorly understood, so there needs to be a degree of confidence that gradients reconstructed from connectopic mapping accurately reflect the genuine underlying topography of such regions.

It has been suggested that gradients could represent an inevitability of the connectopic mapping analysis, rather than accurately reflecting the genuine functional topographies. Ciantar and colleagues noted a confound introduced during data pre-processing which can give rise to illusory local functional connectivity patterns, which could in turn potentially bias connectopic mapping analyses (Ciantar et al., 2022). Specifically, interpolation of the timeseries data from the volume to the cortical surface is liable to induce local spatial autocorrelations amongst neighbouring vertices. Because connectopic gradients reflect the local similarity in connectivity patterns between voxels or vertices, it may be that such gradients are biased by these artificially induced spatial autocorrelations. Importantly, such autocorrelations could be introduced by other processing steps too, such as spatial smoothing or interpolation to other volume spaces. Nevertheless, it remains unclear to what extent this issue practically influences connectopic gradients.

Here, we provide a test of how local spatial autocorrelations affect connectopic mapping. We obtained naturalistic imaging data from a publicly available dataset. We then generated random white noise timeseries matched to the real data in each subject's functional volume. These random datasets are entirely biologically implausible, with data being both spatially and temporally uncorrelated, and so provide a strong test for whether gradients can be produced by artifacts introduced during the analysis pipeline alone. For both real and random datasets, we then optionally applied spatial smoothing and/or interpolated the data to a different volume space or to the cortical surface. We then identified both local and coarse-scale gradients via connectopic mapping. If smoothing and/or interpolation induces spatial autocorrelations sufficient to confound connectopic mapping, we would expect to reconstruct gradients even from random data and these should appear similar to equivalent gradients derived from real data.

2. Methods

2.1. Dataset

Movie-watching and visual field mapping MRI data were obtained for 15 subjects from the publicly available StudyForrest project (Hanke et al., 2016, 2014; <https://www.studyforrest.org>). Briefly, functional data were acquired on a 3T Philips Achieva MRI scanner via an EPI sequence (TR = 2 s, TE = 30 ms, voxel resolution = 3 mm isotropic). The dataset comprises approximately 2 h of MRI data of each subject watching the "Forrest Gump" movie.

2.2. Pre-processing

The same pre-processing steps were applied to movie-watching and visual field mapping datasets. Some light pre-processing was already applied by the StudyForrest project: motion correction using FSL's MCFLIRT tool (Jenkinson et al., 2002) and aligning each volume to a subject-specific EPI template. We then applied further pre-processing using FSL's FEAT v6.0 (Smith et al., 2004; Jenkinson et al., 2012; <https://fsl.fmrib.ox.ac.uk/fsl>): slice-timing correction using Fourier-space time-series phase-shifting, non-brain removal (Smith, 2002), grand-mean intensity normalisation of the entire 4D dataset by a single multiplicative factor, and high-pass temporal filtering (Gaussian-weighted least-squares straight line fitting with $\sigma = 50$ s). Spatial

smoothing was not yet applied at this stage. The timeseries were then standardised by converting to units of percentage signal change, and then finally both the mean ventricular and white-matter timeseries and motion parameters were regressed out.

We computed registrations from each subject's EPI space to both standard volume and surface spaces. First, we computed a volume-based registration between each subject's EPI and T1 anatomical images via boundary based registration (Greve and Fischl, 2009), and then further to the MNI152 standard brain via FSL's FNIRT tool (Andersson et al., 2010). Next, cortical surfaces were reconstructed from T1- and T2-weighted anatomical images using Freesurfer v6.0 (Dale et al., 1999; <https://surfer.nmr.mgh.harvard.edu>). Functional data were co-registered to each subject's native cortical surface via boundary based registration, and then further to the fsaverage surface via a surface-based registration (Fischl et al., 1999a, 1999b).

To synthesise random data, we generated random timeseries comprising Gaussian white-noise matched in mean and variance to the real movie-watching data in the functional EPI volume of each subject. These timeseries will approximate the signal amplitudes of the real data, but will lack any coherent spatial or temporal correlation structure. Both real and random datasets were then optionally transformed to the MNI volume or cortical surface using trilinear interpolation. Data were transformed to the fsaverage5 surface for whole-brain analyses, and the fsaverage6 surface for all other analyses. We then applied spatial smoothing in all three spaces (functional volume, MNI volume, fsaverage surface) at FWHM = 6 mm (twice the voxel resolution), using volume-based smoothing for volume spaces and surface-based smoothing for surface spaces. We also retained the unsmoothed datasets for comparison. Volume datasets were restricted to cortical and subcortical grey-matter masks, derived from individualised Freesurfer segmentations (Fischl et al., 2002) for the functional volumes (back-transformed from the T1 images), or from the Harvard-Oxford atlas for the MNI volume. Thus, each subject had a total of 12 movie-watching datasets: a real and complementary random dataset, with a smoothed and unsmoothed version of each, and with a version of each in the functional volume, MNI volume, and on the fsaverage surface.

2.3. Regions of interest

We measured connectopic gradients in a number of different regions of interest: (1) primary visual cortex (V1), (2) hippocampus, (3) striatum, and (4) Schaefer parcels. We first generated regions of interest (ROIs) for primary visual cortex in each subject. Visual field mapping data were either retained in the native functional volume or transformed to the fsaverage surface. Both spatially smoothed (FWHM = 6 mm) and unsmoothed versions were generated in each space. In each case, a travelling wave analysis (Engel et al., 1997) was conducted using the *3dRetinoPhase* command in AFNI (Cox, 1996; Saad et al., 2001). This produced phase maps of the eccentricity and polar angle tunings. The phase maps from the spatially smoothed surface analysis were used to define individualised V1 ROIs by tracing along the phase reversals. These surface ROIs were additionally back-projected to each subject's functional volume. These ROIs were then applied to the analysis of both smoothed and unsmoothed data in each subject's functional volume and surface spaces. The eccentricity and polar angle phase maps also provided a ground-truth estimate of the retinotopic maps within V1 to compare the connectopic maps against. Next, for analysis of the MNI volume data, we defined ROIs for the hippocampus and striatum from the Harvard-Oxford subcortical atlas. Finally, for further analysis of the surface data, we generated surface ROIs from the 100-area Schaefer parcellation (Schaefer et al., 2018). This parcellation scheme partitions the cortical surface along boundaries in local functional connectivity while simultaneously maximising the functional similarity within each parcel, and has been shown to predict functional boundaries as well as or better than comparable functional, anatomical, and multimodal parcellations (Zhi et al., 2022).

2.4. Connectopic mapping

We tested two separate pipelines for performing the connectopic mapping analysis, based on subtly different methods described in the literature (Haak et al., 2018; Margulies et al., 2016). Both pipelines aim to reconstruct functional gradients from patterns of connectivity, but use differing approaches to achieve this.

2.4.1. Pipeline 1

Our first connectopic mapping pipeline was based on the methods of Haak et al. (2018). A schematic illustration of this pipeline for ROI analyses is shown in Fig. 1. For analyses of V1 data, timeseries were concatenated over odd and even scan runs separately to allow cross-validated parameter selection for the manifold learning. For all other analyses, timeseries were concatenated over all scan runs. Timeseries were split between grayordinates within the ROI and those outside the ROI. For surface-based analyses, this entailed splitting the timeseries between cortical vertices inside and outside of the ROI. For volume-based analyses, the split was between voxels inside and outside of the ROI, with non-ROI voxels selected from a grey-matter masque. Because the number of non-ROI grayordinates exceeds the number of timepoints, we reduced the dimensionality via a lossless PCA – retaining all available components (one fewer than the number of timepoints) so that 100% of the variance remained explained. This amounts to rotating the samples within the feature space and dropping the un-

used dimensions, and aids computational tractability for later processing stages.

Next, connectivity fingerprints were derived by correlating time-series between the ROI grayordinates and non-ROI principal components, and correlations were converted to units of Fisher’s z . We then measured the pairwise Euclidean distances between samples (ROI grayordinates) in the connectivity space, yielding a symmetrical distance matrix. These distances were then used to construct an unweighted and undirected nearest neighbour graph. This is represented as a symmetrical sparse affinity matrix, where two samples are assigned a value of one if either sample falls within the k nearest neighbours of the other, and a value of zero otherwise. This requires a parameter (k) to be selected for the neighbourhood size. For analysis of the V1 ROI, the parameter was selected via cross-validation. Specifically, a Bayesian optimisation algorithm (implemented using the *scikit-optimize* Python module; <https://scikit-optimize.github.io>) selected the parameter that maximised the correlation between retinotopic and connectopic maps in one of the data splits (odd or even runs), and the selected parameter was then applied to the analysis of the other data split. As such, the parameter selection remains statistically independent from the measure of the prediction accuracy. For the analysis of other regions, the ground-truth functional topography is unknown, so it wasn’t possible to optimise a prediction accuracy. Instead, the parameter was set to the log of the number of samples (rounded to the nearest integer) as this has previously been suggested as a reasonable default choice (von Luxburg, 2007). The affinity




Fig. 1. Schematic illustration of the first connectopic mapping pipeline for ROI analyses, based on the methods of Haak et al. (2018). Timeseries are divided between grayordinates within versus outside of the ROI. Non-ROI timeseries are reduced in dimensionality via lossless principal component analysis. The ROI and compressed non-ROI timeseries are cross-correlated to derive the connectivity space. Pairwise distances are measured between samples in the connectivity space, which are then used to derive a nearest neighbour network graph (represented by a sparse affinity matrix). Spectral embedding is then used to derive the connectopic maps. To perform connectopic mapping at the group level, the distance matrices are averaged over subjects, and then the remaining pipeline proceeds as normal.

matrix was then decomposed via spectral embedding to yield the connectopic maps. In all cases, we retained the first two components. Spectral embedding was implemented using the *scikit-learn* Python module (Pedregosa et al., 2011; <https://scikit-learn.org>). The analysis can either be performed at the individual subject level, or at the group-level by averaging the distance matrices over subjects then deriving an affinity matrix and applying spectral embedding as normal. We conducted the V1 ROI analyses at the individual-level, the Schaefer parcel analyses at the group-level, and the MNI volume ROI analyses (hippocampus, striatum) at both individual- and group-levels.

We also applied a modified version of this pipeline to connectopic mapping of whole-brain gradients. Connectivity fingerprints were estimated for all vertices within each hemisphere (excluding the medial wall) by correlating the timeseries in those vertices against principal components derived from those same vertices (via a lossless PCA). From here, the remaining analysis pipeline proceeded as described above. This yielded whole-brain connectopic maps for each hemisphere. We performed the whole-brain analyses at the group-level.

2.4.2. Pipeline 2

We also employed an alternative connectopic mapping pipeline based on the methods of Margulies et al. (2016), primarily for the analysis of whole-brain gradients (Fig. 2). Timeseries were correlated between all vertices in each hemisphere (excluding the medial wall), yielding a symmetrical connectivity fingerprints matrix. Correlations were left in units of Pearson's r because the matrix diagonal comprises perfect correlations which would invalidate a Fisher's z transformation. Next, the connectivity fingerprints were made sparse by thresholding the values at the 90th percentile for each sample (row) in the matrix. Where necessary, values were further thresholded at zero to ensure that all values were non-negative. Next, the thresholded connectivity matrix was converted into an affinity matrix using a similarity kernel: we employed kernels based on cosine, normalised angle, and Pearson's correlation similarity. Although the connectivity matrix is square in the whole-brain case, it is still necessary to derive a separate symmetric affinity matrix because the connectivity matrix is rendered asymmetric by the thresholding. This affinity matrix is dense, such that the similarity is measured from every vertex to every other vertex, and hence represents a weighted and fully-connected network graph. Spectral embedding requires the affinity matrix to be non-negative: for cosine and angular similarity this is already guaranteed by the non-negativity of

the thresholded connectivity fingerprints, but for correlation similarity the affinity matrix was further thresholded at zero. Finally, the affinity matrix was decomposed via spectral embedding, and the first two components were retained. This can be performed at the individual-subject level, or at the group-level by averaging the connectivity fingerprints over subjects before thresholding and then proceeding with the remaining pipeline as normal. To obtain an unbiased average, the connectivity fingerprints were converted to units of Fisher's z before averaging, and then converted back to Pearson's r after averaging. The principal difference between the analysis pipelines is that the first pipeline uses dense connectivity fingerprints and a sparse affinity matrix, while the second pipeline uses sparse connectivity fingerprints and a dense affinity matrix.

A modified version of the second pipeline was also applied to surface- and volume-based ROI analyses of V1. Here, connectivity fingerprints were derived by correlating timeseries between grayordinates within the V1 ROI and all other cortical grayordinates outside of the ROI. Beyond this point, the remaining pipeline stages proceeded as described above. We conducted the V1 ROI analyses at the individual-level, and the whole-brain analyses at the group-level.

2.5. Statistical analyses

For the V1 ROI analyses, we estimated retinotopic and connectopic maps in each subject. Prediction accuracies were measured by the absolute correlation between the eccentricity and polar angle retinotopic and connectopic maps; the absolute value was taken to account for the sign ambiguity in the connectopic maps. Correlations were converted to units of Fisher's z and averaged within each subject over hemispheres and the two cross-validation splits. These correlations were then entered into a four-way repeated-measures ANOVA with factors for the co-ordinate space (volume/surface), randomness (real/random), smoothness (smoothed/unsmoothed) and map (eccentricity / polar angle). Next, we measured the absolute correlations between the real and random connectopic maps. Values were again converted to Fisher's z and averaged within subjects. These correlations were then submitted to a three-way repeated-measures ANOVA with factors for the co-ordinate space, smoothness, and map.

Connectopic maps in the MNI volume ROIs (hippocampus and striatum) were estimated at both the group- and individual-level. In each case, we measured the absolute correlations between the real and ran-

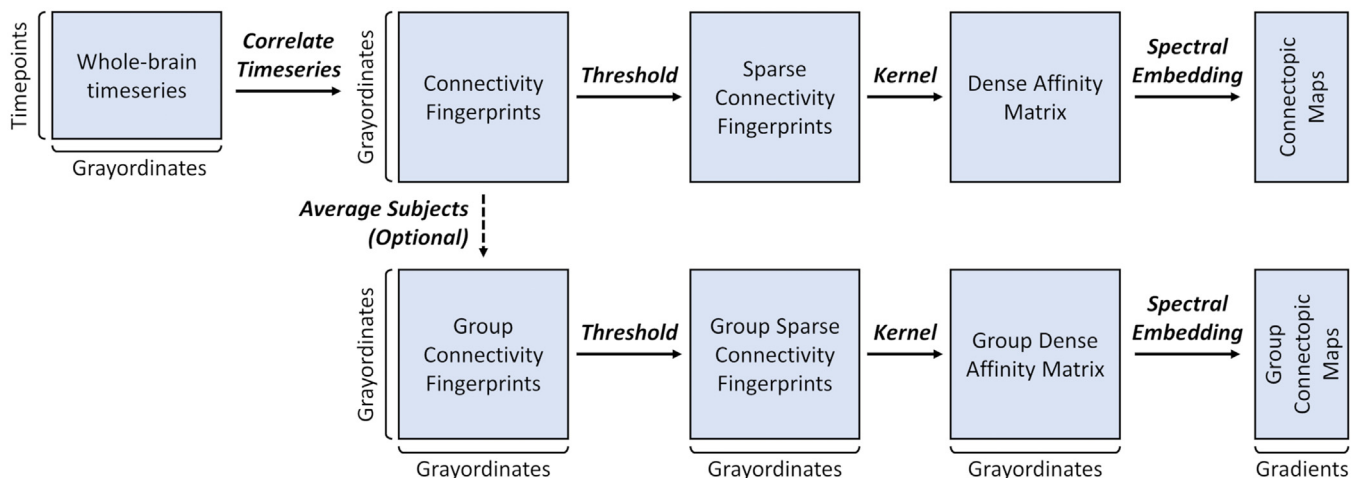


Fig. 2. Schematic illustration of the second connectopic mapping pipeline for whole-brain analyses, based on the methods of Margulies et al. (2016). Whole-brain timeseries are cross-correlated to derive the connectivity space. The connectivity space is made sparse by thresholding at the 90th percentile for each sample (row) in the matrix. This is then converted to a dense (fully-connected) affinity matrix via a similarity kernel. Spectral embedding is then used to derive the connectopic maps. To perform connectopic mapping at the group level, the initial dense connectivity fingerprints are averaged over subjects, and then the remaining pipeline proceeds as normal.

dom connectopic maps, converted the values to Fisher's z , and averaged over hemispheres. Individual correlations were entered into a two-way repeated-measures ANOVA with factors for the smoothness and component (primary and secondary gradients). To assess the consistency of the maps between individuals, we also calculated the pairwise correlations of individual-level gradients between subjects, as well as the correlation of each subject's gradient to the group-level gradient. We additionally estimated noise ceilings from these values by averaging the between-subject correlations over pairwise combinations for each real connectopic map – this estimates the average similarity of each real gradient to itself between subjects. Correlations were converted to Fisher's z before averaging to obtain an unbiased mean, which was then converted back to Pearson's r for visualisation. We compared the Fisher transformed within-subject correlations (between real and random connectopic maps) against these noise ceilings via a series of one-sample t -tests, including Bayes factors calculated with the *BayesFactor* R package (<https://cran.r-project.org/package=BayesFactor>).

We also estimated group-level connectopic maps within each of the 100 Schaefer parcel surface ROIs. Absolute correlations were measured between real and random connectopic maps for each parcel, converted to Fisher's z , and entered into a two-way repeated-measures ANOVA with factors for the smoothness and component.

Finally, we estimated group-level connectopic maps over the whole-brain within each hemisphere. At the whole-brain level, we observed more variable degrees of similarity in real and random connectopic maps within and between components. We therefore measured the full 2×2 pairwise absolute correlations between real and random maps across both components within each hemisphere. We assessed the significance of these correlations using spin permutation tests (1000 permutations per hemisphere) implemented using the BrainSpace toolbox (Vos de Wael et al., 2020; <https://brainspace.readthedocs.io>). On each permutation, the real gradients were randomly rotated over the cortical sphere, and the absolute correlation was measured between the rotated real gradient and original random gradient (excluding invalid vertices within the medial wall). Repeating this process over all permutations produces an empirical null distribution of correlation values. Statistical significance was determined by identifying the proportion of correlations in the null distribution that were greater than the genuine correlation value.

In all ANOVAs, we report effect sizes in units of both partial and generalised eta-squared (Bakeman, 2005; Olejnik and Algina, 2003). All statistical tests employed an alpha criterion of 0.05 for determining significance.

2.6. Connectivity searchlight

To assess how the pre-processing pipeline affects spatial autocorrelations, we used a searchlight analysis (Ciantar et al., 2022; Kriegeskorte et al., 2006) to measure local functional connectivity. We employed spherical searchlights for volume spaces (5 mm / 2 voxel radius for functional volumes, 6 mm / 2 voxel radius for the MNI volume) and a disc searchlight (6 mm radius) for the surface space. Volume searchlights were restricted to a grey-matter mask. For each searchlight, we calculated the pairwise correlations in timeseries between all grayordinates within the searchlight. These correlations were converted to units of Fisher's z and averaged over grayordinates, and then this average was converted back to Pearson's r and assigned to the central grayordinate of the searchlight. This resulted in a map of the local functional connectivity over the whole brain. We performed searchlight analyses in each individual subject for all 12 versions of the dataset: real and random data, smoothed and unsmoothed, and in the functional volume, MNI volume, and surface spaces. Volume searchlights were implemented with the *PyMVPA* toolbox (Hanke et al., 2009; <http://www.pympva.org/>), and surface searchlights were implemented using custom code leveraging the *surfdist* toolbox (Margulies et al., 2016; <https://github.com/NeuroanatomyAndConnectivity/surfdist>).

3. Results

3.1. ROI connectopic mapping

We employed our first connectopic mapping pipeline to analyses of local regions of interest. We first examined gradients estimated by retinotopic and connectopic mapping in primary visual cortex (Fig. 3). We employed both surface- and volume-based analyses. For surface analyses, the timeseries were interpolated to the cortical surface and retinotopic and connectopic mapping were performed on the surface. For volume analyses, retinotopic and connectopic mapping were performed within the native functional volume, and the resulting maps were interpolated to the surface for visualisation. We observed highly similar connectopic maps between the real and random data, indicating that spatial autocorrelations induced via spatial smoothing and/or interpolation to the surface were sufficient to produce illusory gradients. Retinotopic and connectopic maps are shown for a representative subject in Fig. 3a. As previously reported (Haak et al., 2018; Watson and Andrews, 2022), connectopic mapping of the real movie-watching data identified a posterior-anterior gradient corresponding to the eccentricity map, and a superior-inferior gradient corresponding to the polar angle map. This was true for both surface and volume analyses and using both smoothed and unsmoothed data. However, similar gradients were also obtained from connectopic mapping of the random data. The only exception to this was the case of the unsmoothed random data in the functional volume, where the connectopic mapping is operating directly on the raw random data with no further manipulation.

We measured prediction accuracies in all subjects by correlating the retinotopic and connectopic maps (Fig. 3b, 3d). Although there was a slight reduction in prediction accuracies for connectopic mapping of random compared to real data, correlations were generally high in all cases – with the exception of the unsmoothed random data in the functional volume. Prediction accuracies were also generally higher for eccentricity than polar angle maps. A four-way repeated-measures ANOVA revealed significant main effects of the co-ordinate space ($F(1,14) = 105.42, p < .001, \eta_p^2 = 0.88, \eta_G^2 = 0.39$) due to higher correlations for the surface pipeline, of randomness ($F(1,14) = 47.06, p < .001, \eta_p^2 = 0.77, \eta_G^2 = 0.18$) due to higher correlations for real data, of smoothness ($F(1,14) = 204.66, p < .001, \eta_p^2 = 0.94, \eta_G^2 = 0.39$) due to higher correlations for smoothed data, and of map-type ($F(1,14) = 28.31, p < .001, \eta_p^2 = 0.67, \eta_G^2 = 0.30$) due to higher correlations for eccentricity maps. Additionally, all interactions were significant (all $p < .05$) except for the two-way smoothness by map-type interaction ($F(1,14) = 1.36, p = .263, \eta_p^2 = 0.09, \eta_G^2 < 0.01$) and three-way randomness by smoothness by map-type interaction ($F(1,14) = 0.03, p = .873, \eta_p^2 < 0.01, \eta_G^2 < 0.01$).

We also tested the consistency between the connectopic maps themselves by correlating the real and random connectopic maps (Fig. 3c, 3e). Correlations were generally high, except for those of the unsmoothed volume analyses, and appeared comparable to or higher than the retinotopic prediction accuracies. A three-way repeated-measures ANOVA indicated significant main effects of the co-ordinate space ($F(1,14) = 33.36, p < .001, \eta_p^2 = 0.70, \eta_G^2 = 0.33$) due to higher correlations for the surface analyses, of smoothness ($F(1,14) = 168.90, p < .001, \eta_p^2 = 0.92, \eta_G^2 = 0.47$) due to higher correlations for smoothed data, and of map-type ($F(1,14) = 33.63, p < .001, \eta_p^2 = 0.71, \eta_G^2 = 0.16$) due to higher correlations for the eccentricity maps. All interactions were significant (all $p < .05$) except for the two-way smoothness by map-type interaction ($F(1,14) = 0.07, p = .794, \eta_p^2 < 0.01, \eta_G^2 < 0.01$). These results indicate that spatial autocorrelations alone, induced via spatial smoothing and/or interpolation to the surface, are sufficient to produce illusory connectopic maps resembling retinotopic gradients in V1.

We next tested volume-based connectopic gradients in the hippocampus and striatum. We again observed very similar connectopic maps between real and random data, indicating that spatial smoothing and




Fig. 3. First connectopic mapping pipeline results for primary visual cortex. (a) Eccentricity and polar angle retinotopic and connectopic maps in the left hemisphere of a representative subject derived from both surface and volume pipelines. Volume maps were projected to the surface for visualisation. (b,d) Prediction accuracies for (b) surface and (d) volume pipelines, measured by absolute correlations between retinotopic and connectopic maps. (c,e) Absolute correlations between real and random connectopic maps for (c) surface and (e) volume pipelines. Dot markers indicate per-subject correlations and bars indicate group means.




Fig. 4. First connectopic mapping pipeline results for MNI volume regions. (a,c) Primary and secondary group-level gradients in the (a) hippocampus and (c) striatum. (b,d) Absolute correlations between real and random gradients for the (b) hippocampus and (d) striatum. Dot markers indicate per-subject correlations for individual-level gradients, bars indicate group means of per-subject correlations, and diamond markers indicate correlations for group-level gradients. Dotted lines indicate noise ceilings estimated as the average pairwise correlation of each real individual-level gradient between all subjects.

interpolation to the MNI volume were sufficient to produce illusory gradients. Furthermore, these gradients were highly consistent between subjects. Group-level connectopic maps of the hippocampus are shown in Fig. 4a, and individual-level gradients are shown for a representative subject in Supplementary Fig. 2a. Connectopic mapping of the real data identified a primary gradient running posterior-anterior along the long axis of the hippocampus, replicating previously reported gradients (Borne et al., 2023; Przeździk et al., 2019; Vos de Wael et al., 2018) in this region, and a secondary gradient that sub-divided the primary gradient. However, we also observed similar gradients using the random data. Unlike the V1 analyses in subjects' functional volumes, these gradients were now also apparent for unsmoothed random data – this indicates that interpolation to the MNI space was sufficient to induce the necessary spatial autocorrelations even without further spatial smoothing. Correlations between the real and random connectopic maps appeared high at both the group and individual levels (Fig. 4b). Entering the individual correlations into a two-way repeated-measures ANOVA revealed no significant main effect of smoothness ($F(1,14) = 0.02, p = .883, \eta_p^2 < 0.01, \eta_G^2 < 0.01$), but did indicate a significant main effect of the component ($F(1,14) = 535.13, p < .001, \eta_p^2 = 0.97, \eta_G^2 = 0.70$) due to higher correlations for the primary gradient, and a significant smoothness by component interaction ($F(1,14) = 23.68, p < .001, \eta_p^2 = 0.63, \eta_G^2 = 0.21$) due to a bigger difference between gradients for smoothed data. We additionally compared individual-level gradients between subjects (Supplementary Fig. 3a) – these revealed very high spatial correlations both between subjects and to the group-level for both real and random gradients. We estimated noise ceilings by averaging over the between-subject correlations for the real gradients. The within-subject correlations between real and random gradients (Fig. 4b) approached the noise ceilings. The correlations fell significantly below the noise ceiling for the unsmoothed

primary ($t(14) = 4.76, p = .001, \text{Hedges' } g = 1.16, \text{BF}_{10} = 109.00$) and smoothed secondary gradients ($t(14) = 4.11, p = .003, \text{Hedges' } g = 1.00, \text{BF}_{10} = 36.54$). However, the correlations for the unsmoothed secondary ($t(14) = 0.95, p = .717, \text{Hedges' } g = 0.23, \text{BF}_{10} = 0.39$) and smoothed primary gradients ($t(14) = 0.93, p = .717, \text{Hedges' } g = 0.23, \text{BF}_{10} = 0.38$) did not differ significantly from the noise ceilings, and Bayes factors indicated slight support for the null hypothesis. Thus, very similar hippocampal gradients were obtained from real and random data, and these were highly consistent over subjects.

Group-level connectopic maps of the striatum (Fig. 4c) indicated a primary gradient that curved around the striatum and a secondary gradient running superior-inferior, replicating previously reported gradients (Marquand et al., 2017). However, these gradients were again observed for both the real and random data with and without spatial smoothing. Similar gradients were also found at the individual-level (Supplementary Fig. 2b). Correlations between real and random gradients (Fig. 4d) again appeared high at both the group and individual-level. A two-way repeated-measures ANOVA on the individual correlations indicated significant main effects of smoothness ($F(1,14) = 53.15, p < .001, \eta_p^2 = 0.79, \eta_G^2 = 0.52$) due to higher correlations for smoothed data, and of the component ($F(1,14) = 400.06, p < .001, \eta_p^2 = 0.97, \eta_G^2 = 0.74$) due to higher correlations for the primary gradient, while the smoothness by component interaction was not significant ($F(1,14) = 3.58, p = .079, \eta_p^2 = 0.20, \eta_G^2 = 0.03$). We again observed very high spatial correlations between subjects for both real and random gradients (Supplementary Fig. 3b). We also found that the within-subject correlations between real and random gradients approached the noise ceilings estimated from the between-subject correlations (Fig. 4b). Within-subject correlations fell significantly below the noise ceilings for the unsmoothed primary ($t(14) = 9.09, p < .001, \text{Hedges' } g = 2.22, \text{BF}_{10} = 5.20 \times 10^4$), unsmoothed secondary ($t(14) = 6.31, p < .001, \text{Hedges' } g = 1.54,$

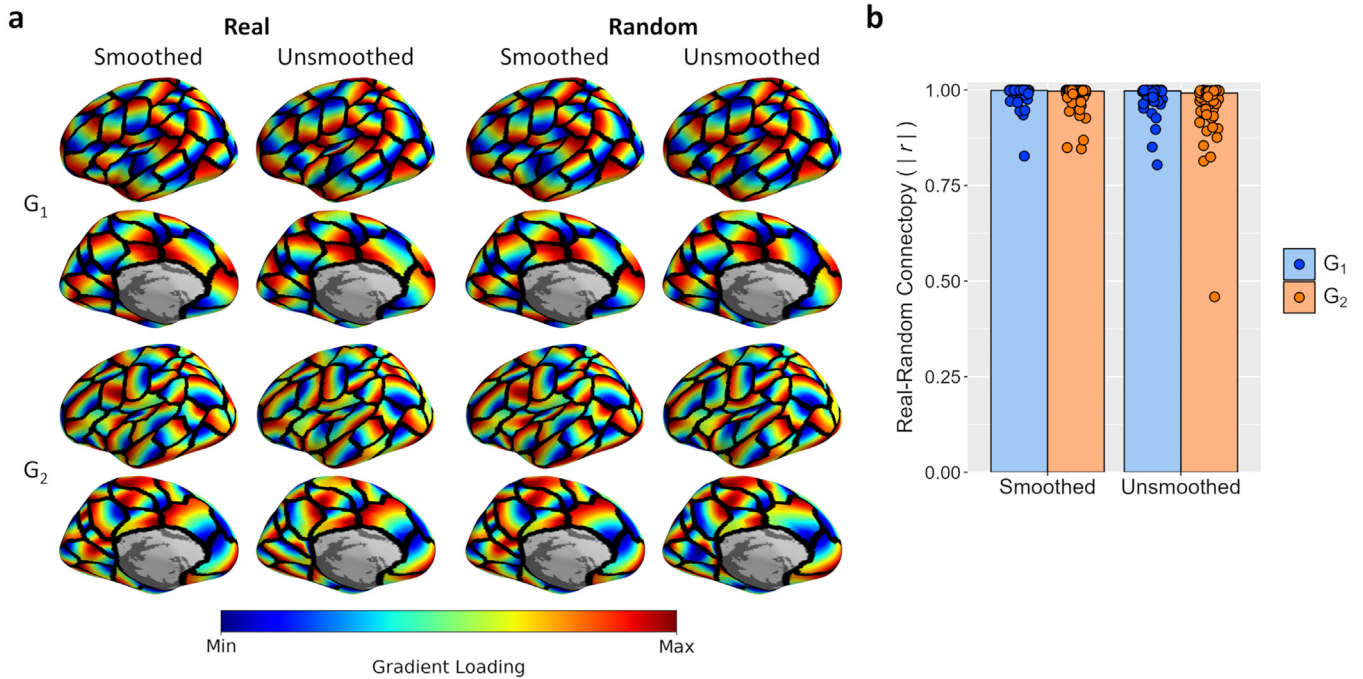


Fig. 5. First connectopic mapping pipeline results for group-level analyses of the 100-area Schaefer surface parcellation. (a) Primary and secondary gradients within each parcel in the left hemisphere. (b) Absolute correlations between real and random group-level gradients. Dot markers indicate correlations for each parcel and bars indicate means over parcels.

$BF_{10} = 1231.97$), and smoothed primary gradients ($t(14) = 10.55$, $p < .001$, Hedges' $g = 2.58$, $BF_{10} = 2.76 \times 10^5$). However, correlations did not differ significantly from the noise ceiling for the smoothed secondary gradient, and the Bayes factor indicated slight support for the null hypothesis ($t(14) = 0.95$, $p = .356$, Hedges' $g = 0.23$, $B_{10} = 0.39$). Therefore, again, striatal gradients appeared very similar between real and random data and highly consistent over subjects. Thus, both spatial smoothing and interpolation to a different volume space (the MNI brain) induced sufficient spatial autocorrelations to obtain local connectopic gradients in the hippocampus and striatum, and these replicated gradients previously reported in the literature.

We next examined local surface-based gradients across the whole brain. We estimated connectopic maps for each region of the 100-area Schaefer surface parcellation (Schaefer et al., 2018). Highly similar local connectopic maps were observed between real and random data throughout the brain, again indicating that spatial smoothing and/or interpolation to the surface can induce illusory gradients. Group-level connectopic maps are illustrated for the left hemisphere in Fig. 5a, and for both hemispheres in Supplementary Fig. 4. Gradients typically ran smoothly and continuously over the extent of each region, with primary and secondary gradients usually appearing orthogonal to one another. These gradients were again obtained using both real and random smoothed and unsmoothed data. Correlations between real and random gradients appeared high in all parcels (Fig. 5b). Entering the correlations over parcels into a two-way repeated-measures ANOVA revealed a significant main effect of smoothing ($F(1,99) = 34.77$, $p < .001$, $\eta_p^2 = 0.26$, $\eta_G^2 = 0.09$) due to higher correlations for smoothed data, and of the component ($F(1,99) = 101.65$, $p < .001$, $\eta_p^2 = 0.51$, $\eta_G^2 = 0.13$) due to higher correlations for the primary gradient, while the smoothness by component interaction was not significant ($F(1,99) = 0.62$, $p = .434$, $\eta_p^2 < 0.01$, $\eta_G^2 < 0.01$). Thus, spatial smoothing and/or interpolation to the surface induced sufficient spatial autocorrelations to generate local connectopic gradients throughout the brain. Similar results were obtained using an alternative multi-modal brain parcellation (Glasser et al., 2016; Supplementary Fig. 5).

3.2. Whole brain connectopic mapping

All of the gradients we have considered so far have been on a local scale within regions of interest. However, connectopic mapping has also been used to describe coarse-scale gradients spanning the whole brain (Margulies et al., 2016). Such studies have typically reported a primary gradient that runs between unimodal and transmodal regions, and a secondary gradient that runs between primary visual and somatosensory regions. To test how such coarse-scale gradients are affected by spatial autocorrelations, we ran both our connectopic mapping pipelines for the whole-brain within each hemisphere. Similar to the region of interest analyses, the first connectopic pipeline (using a sparse affinity matrix) identified gradients from both real and random data. However, none of these gradients appeared biologically plausible and they did not replicate previously reported whole-brain gradients. The second pipeline (using a dense affinity matrix) successfully replicated previous gradients, and appeared less susceptible to spatial autocorrelations, but was sensitive to the choice of similarity kernel.

Group-level connectopic maps derived from the first connectopic mapping pipeline are shown for the left hemisphere in Fig. 6a (data for both hemispheres are shown in Supplementary Fig. 6). Using this pipeline, we failed to replicate previously reported gradients even with the real data, and instead identified biologically implausible gradients that simply ran continuously across the brain along one axis (anterior-posterior and superior-inferior for the primary and secondary gradients respectively). Similar gradients were obtained between real and random spatially smoothed data. Gradients were still obtained for unsmoothed random data, but appeared less similar to their real data counterparts. We measured the pairwise absolute correlations between real and random gradients (Fig. 6b), and determined statistical significance via spin permutations (Vos de Wael et al., 2020). With spatial smoothing, real and random gradients correlated strongly and significantly within each component. Without spatial smoothing, moderate correlations were obtained both within and between components, though spin permutations indicated none were significant.




Fig. 6. Group-level connectopic mapping results for whole-brain surface analyses. (a,c) Primary and secondary gradients in the left hemisphere for the first and second pipelines respectively. (b,d) Absolute correlations between real and random gradients for the first and second pipelines respectively. Pairwise correlations were measured between both gradients, and significance was determined via spin permutation tests. Results for the second pipeline are illustrated for the cosine similarity kernel.

Our first connectopic mapping pipeline is based on the methods described by Haak et al. (2018). It is possible that the failure to replicate previously reported whole-brain gradients reflects subtle methodological differences between this pipeline and the one employed by Margulies et al. (2016). One key difference lies in the treatment of the connectivity fingerprints and generation of the affinity matrix (Figs. 1 & 2). Haak et al. measure dense connectivity fingerprints (from all grayordinates to all principal components), then construct a sparse affinity matrix (measuring affinity only between neighbouring samples in the connectivity space). By contrast, Margulies et al. make the connectivity fingerprints sparse by thresholding them, then construct a dense / fully-connected affinity matrix (measuring the affinity between every sample and every other sample). To test the influence of these methodological differences on the gradients, we re-ran our whole-brain analyses using our second pipeline based on the methods of Margulies et al. We first constructed the affinity matrix using a cosine kernel, as per Margulies et al. (2016). Group-level connectopic maps are shown for the left-hemisphere in Fig. 6c and for both hemispheres in Supplementary Figure 7. Using either smoothed or unsmoothed real data, we were

now able to replicate the unimodal-transmodal primary gradient and the visual-somatosensory secondary gradient. Using smoothed random data produced gradient-like patterns, but these did not correspond well to the gradients obtained with the real data and did not appear biologically plausible. Using unsmoothed random data did not produce convincing gradients – the maps instead converged on local brain regions and appeared largely uniform elsewhere. Correlations between real and random maps were generally low, and spin permutations did not indicate any of them to be significant. To ensure that the key factor in producing these gradients was the sparseness of the affinity matrix, we also ran a variant of the second pipeline in which the connectivity fingerprints were again made sparse by thresholding, but these were then used to generate a sparse affinity matrix as per the first pipeline. The resulting gradients again ran continuously along one axis and appeared similar to those obtained from the first pipeline (Supplementary Fig. 8). Thus, the ability to reproduce previously reported whole-brain gradients critically depends on using a dense rather than sparse affinity matrix.

For the second pipeline, various similarity metrics can be used to construct the affinity matrix. In the above analyses, we used a cosine

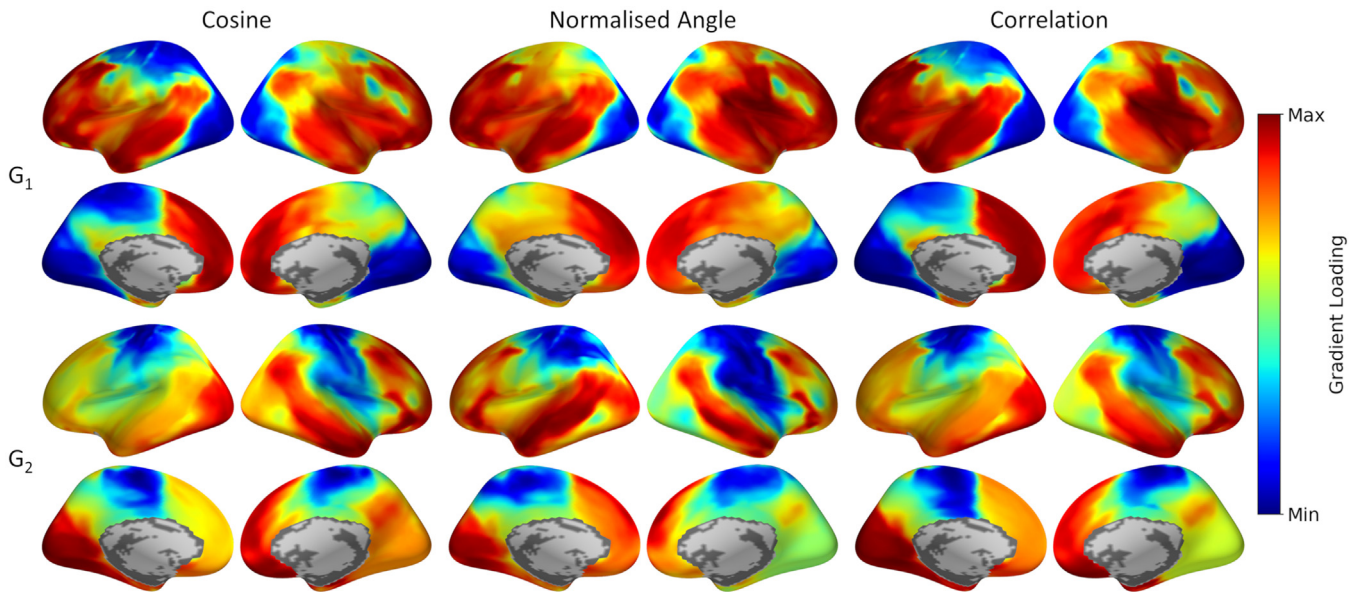


Fig. 7. Effect of varying the similarity kernel on the second connectopic mapping pipeline for group-level whole-brain surface analyses of the spatially smoothed real data. Plots illustrate primary and secondary gradients. Affinity matrices were derived using cosine, normalised angle, or Pearson's correlation similarity kernels.

kernel as per Margulies et al. (2016), however other metrics are available. To test how the choice of kernel affects the resulting gradients, we re-ran our analyses of the real and spatially smoothed data using cosine, normalised angle, and Pearson's correlation similarity metrics. Group-level gradients are illustrated in Fig. 7. While gradients do appear similar in all cases, they also differ in some key features that alter the interpretation of the gradients. For instance, using a normalised angle metric, the primary gradient represented a divide between primary visual cortex and the rest of the brain, rather than a unimodal-transmodal distinction. The secondary gradient under normalised angle similarity more strongly emphasised somatosensory versus superior and anterior temporal regions, rather than visual versus somatosensory. Gradients from the correlation kernel appeared to fall between the cosine and normalised angle solutions. Thus, although we can reproduce previously reported whole-brain gradients from this dataset, they are sensitive to the choice of analysis pipeline. Specifically, they critically depend on using a dense rather than sparse affinity matrix, and are sensitive to the choice of kernel used to construct the affinity matrix.

Because the second pipeline appears less susceptible to spatial correlations at the whole-brain level, we also tested how this pipeline (using a cosine similarity kernel) performed at a local level by re-running our analyses of primary visual cortex (Supplementary Fig. 9). In general, the connectopic maps provided a poorer reconstruction of the retinotopic maps than those obtained with the first pipeline. At the same time, while the degree of similarity between real and random connectopic maps decreased relative to the first pipeline, the correlations were still moderate and appeared comparable to the retinotopic prediction accuracies. Thus, while the second pipeline appears to reduce the influence of spatial autocorrelations on local connectopic maps, it does not eliminate the problem, and meanwhile produces poorer reconstructions of the ground-truth retinotopic maps compared to the first pipeline.

3.3. The effect of smoothing and interpolation on local functional connectivity

Finally, to test how the analysis pipeline affects spatial autocorrelations and local functional connectivity, we performed a series of connectivity searchlight analyses (Fig. 8). These indicated that spatial autocorrelations could be induced even in random data by applying spa-

tial smoothing and/or interpolating between volume or surface spaces. Searchlight connectivity maps for a representative subject are shown in Fig. 8a-c, and distributions of correlations over all subjects are illustrated in Fig. 8d. As expected, local functional connectivity was near zero for the unsmoothed random data in the functional volume, confirming that the random data as generated are spatially uncorrelated. However, both spatial smoothing and interpolation to another co-ordinate space (either volume- or surface-based) increased spatial correlations in both real and random data. It is likely that these artificially induced spatial autocorrelations form the basis of the connectopic maps identified in the real and random data. As previously reported (Ciantar et al., 2022), the distribution of local functional connectivity over the surface covaried with the surface curvature – correlations increased in regions with high curvature (Fig. 8e). We also correlated the average surface searchlight maps between the different datasets (Fig. 8f). Moderate-to-strong positive correlations were found between most real and random smoothed and unsmoothed datasets (with the exception of random-smoothed against real-unsmoothed) indicating some commonality in the pattern of local connectivity between datasets. Thus, both spatial smoothing and interpolation to either the surface or a different volume space induce artificial spatial autocorrelations sufficient to produce illusory gradients with connectopic mapping.

4. Discussion

In this study we investigated how local spatial autocorrelations introduced during data analysis influence gradients produced by connectopic mapping. We synthesised spatially uncorrelated random data in each subject's native functional volume, then explored the effect of applying spatial smoothing and interpolation to other co-ordinate spaces. We found that artificial spatial autocorrelations could be induced by applying spatial smoothing and/or interpolating to the cortical surface or MNI volume space. These spatial autocorrelations were sufficient to derive local gradients from connectopic mapping of both volume- and surface-based brain regions. Connectopic maps derived from the random data appeared highly similar to those produced from the real data and replicated gradients previously reported in the literature. Coarser-scale gradients over the whole-brain appeared less susceptible to artificial spatial autocorrelations but were sensitive to specific features of the analysis pipeline.

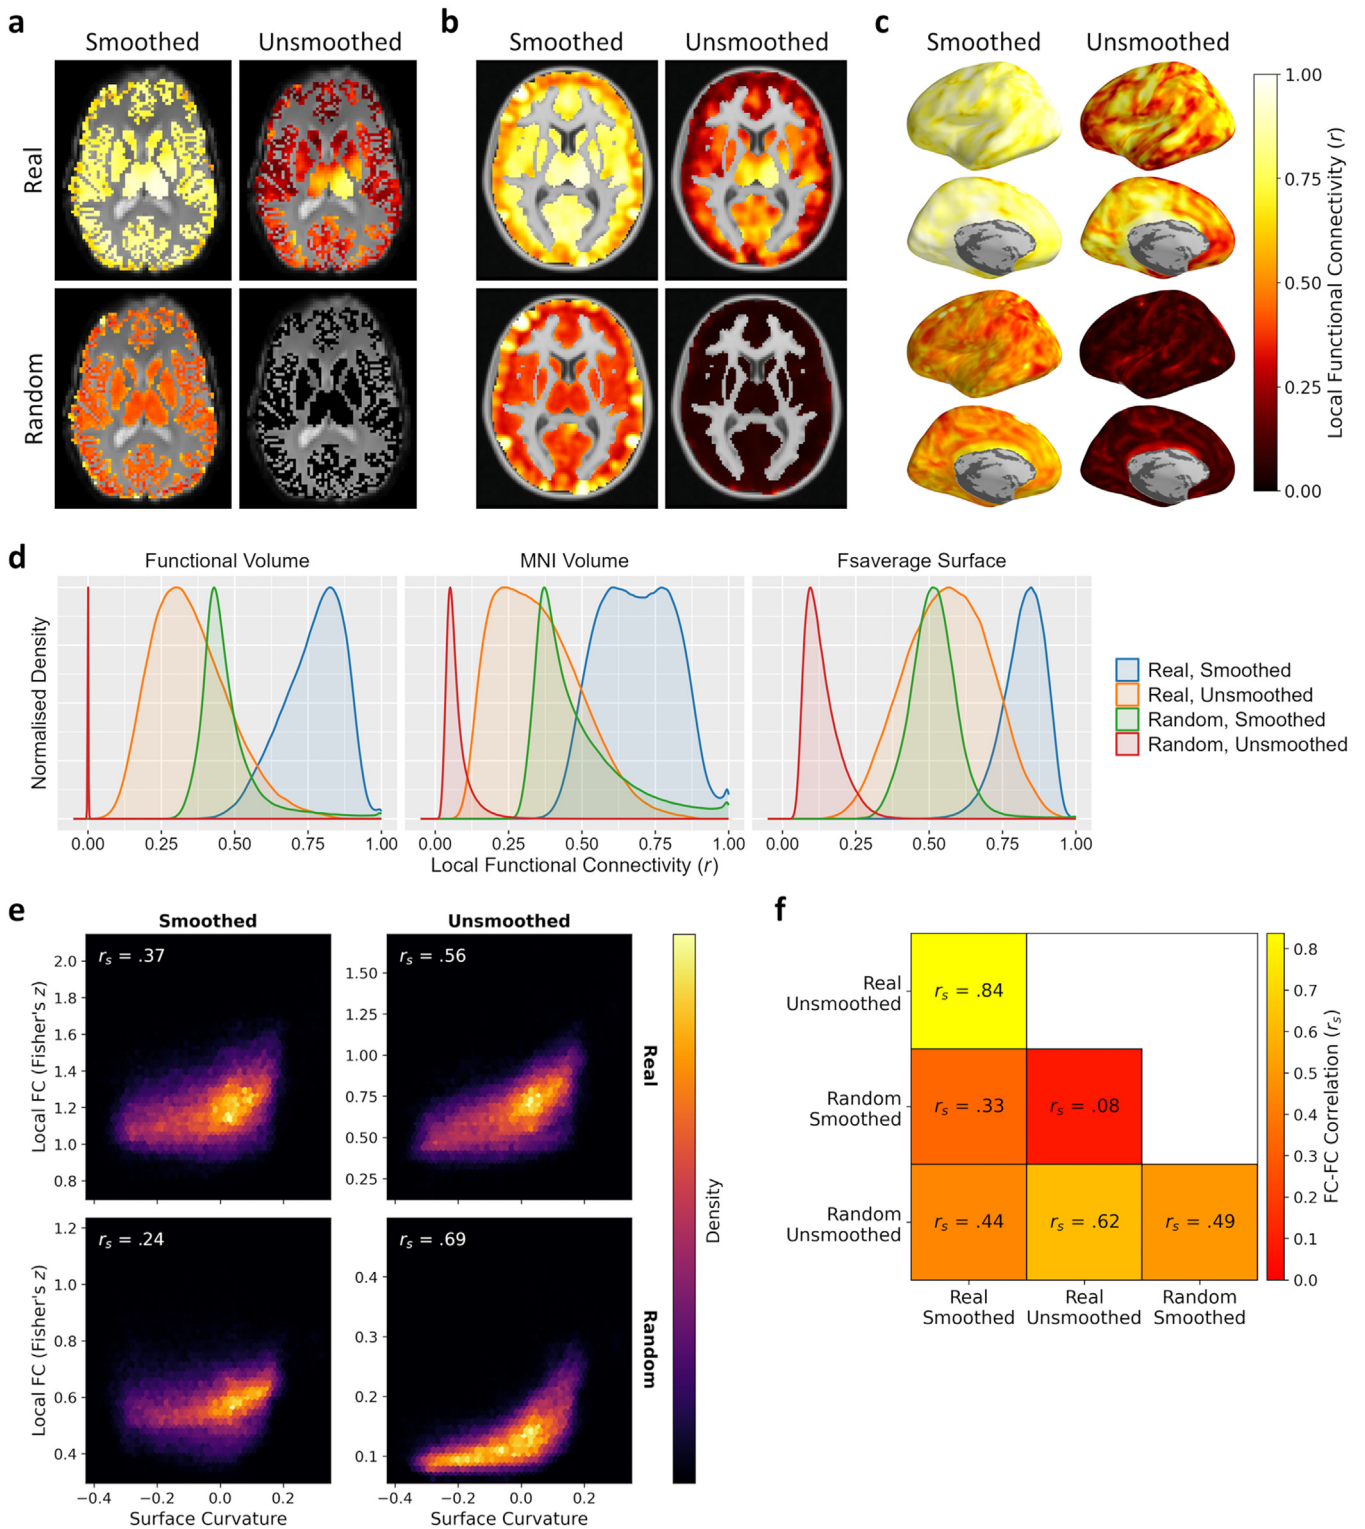


Fig. 8. Spatial smoothing and interpolation induce local spatial autocorrelations within random data. (a-c) Local connectivity searchlight maps for a representative subject based on data (a) retained in the native functional volume, (b) transformed to the MNI volume space, or (c) transformed to the fsaverage surface (illustrated for left hemisphere only). (d) Kernel density estimates illustrating distributions of searchlight correlations for all subjects. (e) Hexbin plots illustrating relationship between surface curvature and group average surface searchlight maps. Spatial autocorrelations increase in regions of high surface curvature. (f) Correlations between different variants of group average surface searchlight maps.

We first investigated local gradients within both volume- and surface-based brain regions using our first connectopic mapping pipeline. Connectopic mapping of real movie-watching data reproduced previously reported volume- and surface-based gradients in primary visual cortex (Haak et al., 2018; Watson and Andrews, 2022), the hippocampus (Borne et al., 2023; Przeździk et al., 2019; Vos de Wael et al., 2018), and the striatum (Marquand et al., 2017). We also reported novel surface-based local gradients throughout the brain within regions of the Schaefer and Glasser parcellations (Glasser et al., 2016; Schaefer et al., 2018). However, we were also able to obtain highly similar gradients from random data following either spatial smoothing or interpolation to other volume or surface spaces. This suggests the manifold structure within the connectivity space is biased by spatial autocorrelations (naturally present in the data and/or introduced during data processing), such that neighbouring samples within the connectivity space tend to also be anatomically proximal. Consequently, the connectopic mapping tends to yield gradients that flow continuously over the extent of the region, even when derived from random data. The only combination that did not yield gradients was applying connectopic mapping directly to the raw unsmoothed random data in the native functional volume without any further manipulation. However, the complete lack of any spatial correlations in the raw random data is biologically implausible – some degree of spatial correlations will always be present in real data, even if data are retained in the native volume and no artificial spatial smoothing is applied. While a functional gradient does necessitate some degree of spatial smoothness in the data, the reverse is not true and spatially smooth data could still reflect non-topographic organisations. Thus, with real data it will be challenging to disentangle the effects of genuine functional topographies versus non-topographic spatial correlations on local connectopic gradients.

We note that many analyses in this work rely on similarity in the visual appearance of gradients, and it is possible that connectopic maps in real data may be distinguishable from smoothed/interpolated random data. Indeed, in primary visual cortex we observed significantly higher prediction accuracies of the retinotopic maps for real compared to random connectopic maps, providing support for real connectopic gradients being distinct from random gradients induced by smoothing and/or interpolation. This may indicate contributions of genuine functional signals from the real data, or could reflect the higher spatial autocorrelations in the real compared to random data – the connectivity searchlight analysis revealed that spatial correlations were generally higher in the real than the random data even after smoothing or interpolation (see Fig. 8). Furthermore, the correlations between the real and random connectopic maps were typically comparable to or higher than their correlations with the retinotopic maps, so the greatest similarity was observed between the real and random gradients themselves. As such, further work is needed to establish to what degree smoothing/interpolation explain connectopic mapping gradient findings. For instance, future studies could titrate the degree of spatial autocorrelations to more fully determine the extent of their influence on connectopic gradients. Nevertheless, our results indicate that standard data preprocessing steps, such as spatial smoothing and interpolation between co-ordinate spaces, are sufficient to induce illusory connectopic gradients.

The organisation of the gradients also appeared highly consistent between regions, with the primary gradient typically running continuously along a single axis from one side of the region to the other, and the secondary gradient usually either subdividing or running orthogonal to the primary gradient. This suggests that these local gradients may be an inevitability of the connectopic mapping technique, rather than necessarily reflecting genuine underlying functional topographies. The specific direction of the gradients could be influenced by interactions with the surface curvature (Ciantar et al., 2022), or could be biased by the shape of the brain region – for instance being orientated relative to the long axis of the region. Such anatomical constraints might influence both connectopic mapping and functional brain organisation so

that connectopic and functional gradients may coincide, as was found in primary visual cortex. However, this is not always guaranteed to be the case, and indeed connectopic mapping failed to reproduce the phase reversals in the retinotopic polar angle map (which delineate the boundaries between retinotopic regions) of an expanded early visual region including V1, V2, and V3 (Supplementary Fig. 1). Even where functional and connectopic gradients do coincide, there remains the issue that the connectopic maps are likely to be only partially driven by the underlying functional topography. In a region where the functional topography is unknown or poorly understood, it would therefore be difficult to ascertain how accurately the connectopic mapping is able to reproduce such topographies. Previously, we noted that the source of connectivity signals had minimal impact on the gradients reconstructed from connectopic mapping (Watson and Andrews, 2022), and interpreted this as indicating that connectopic mapping largely reflects the functional topography within a brain region, rather than necessarily a topography embedded within the connectivity itself. Our current results partially accord with this principle, but further imply that the topography within a region need not even have a functional origin. Instead, illusory local gradients may be driven by an artificial topography induced by spatial autocorrelations introduced during the analysis pipeline.

One key issue is how connectopic mapping might be expected to perform under the null hypothesis. Spatial autocorrelations within the data make it unlikely that a random white noise pattern would be produced even when derived from spatially correlated random data. However, it is not guaranteed that the technique would yield similar gradients from both the real and random data. The spatial extent of the autocorrelations will be less than the size of the brain regions, so spatially correlated random data might have produced a smooth but still disorganised pattern rather than a continuous gradient. Alternatively, gradients might have been produced but following a different axis or pattern between the real and random data. In practice, we did obtain continuous gradients in each region, presumably as the connectopic mapping chained the local spatial correlations together over the extent of the region, and these appeared highly similar between the real and random data. For the hippocampus and striatum regions, we observed that individual-level gradients were highly consistent both between subjects and between each subject and the group-level gradients (Supplementary Fig. 3). We estimated noise ceilings by averaging the between-subject correlations for the real gradients, indicating how well each real gradient correlated with itself between subjects, and observed that the within-subject correlations between real and random gradients approached these noise ceilings. This not only indicates that the real and random hippocampal and striatal gradients were highly consistent over subjects, but also that the similarity between the real and random gradients was close to the ceiling level. No genuine gradients are present within the random data itself, so any gradients produced by the connectopic mapping must be illusory. The high similarity observed between real and random connectopic gradients therefore indicates that the spatial autocorrelations driving the random gradients likely also strongly influence the gradients obtained from real data.

We also tested how spatial autocorrelations influence coarser-scale surface-based gradients over the whole-brain. Using our first connectopic mapping pipeline, based on the methods of Haak et al. (2018), we obtained whole-brain gradients that appeared highly consistent between spatially smoothed real and random data. Without spatial smoothing, gradients were again obtained for random data, though appeared less similar to their real data counterparts. However, in all cases, gradients simply ran smoothly and continuously along a single axis along the cortical surface. Such gradients appear biologically implausible, and failed to replicate the unimodal-transmodal and visual-somatosensory gradients previously reported by Margulies et al. (2016). This result was surprising, as the Margulies et al. gradients have been replicated many times, including from resting-state functional connectivity (Bethlehem et al., 2020; Cross et al., 2021; Gao et al., 2022; Hong et al., 2019; Larivière et al., 2020; Lee et al., 2023; Samara et al.,

2023; Xia et al., 2022; Zhang et al., 2022), structural connectivity (Park et al., 2021a, 2021b), and brain microstructure (Paquola et al., 2019). However, all previous replications used a connectopic mapping pipeline based on the methods of Margulies et al. (2016), which differs subtly from Haak et al.'s (2018) methods. Specifically, Haak et al. estimate a sparse affinity matrix in which the affinity is only measured between samples within local neighbourhoods of the connectivity space, while Margulies et al. estimate a dense affinity matrix in which the affinity is measured between all samples. We were able to reproduce the unimodal-transmodal and visual-somatosensory gradients when repeating our analyses with a second connectopic mapping pipeline based on Margulies et al.'s methods using a dense affinity matrix. Furthermore, these gradients were not produced by the random data, indicating that this pipeline is less susceptible to the spatial autocorrelations. Employing a sparse affinity matrix is intended to aid in identifying non-linearities within the geometric structure of the manifold by restricting the network graph to connections between neighbouring samples lying along the manifold surface, and ignoring connections that may jump across gaps or pass through the manifold surface. However, if spatial autocorrelations in the data make it so that neighbouring samples in the connectivity space also tend to be neighbouring in the anatomical space, then a sparse affinity matrix may be more prone to yielding connectopic maps in which gradients simply run continuously over the extent of the brain region. By contrast, a dense affinity matrix, representing a fully-connected and weighted network graph, provides a coarser-scale view of the manifold that may be less biased by local spatial correlations. However, such a representation may also make it more challenging to identify non-linearities in the manifold and will hence promote more linear solutions. Indeed, the second pipeline did perform more poorly at reconstructing retinotopic maps in primary visual cortex compared to the first pipeline (Supplementary Fig. 9).

We further tested how the choice of similarity kernel affects the whole-brain gradients produced by the second connectopic mapping pipeline. While Margulies et al. (2016) constructed their affinity matrix using cosine similarity, other metrics such as normalised angle or correlation similarity are also available (Vos de Wael et al., 2020). While the gradients appeared similar in all cases, they differed in a number of important ways that could alter the interpretation of the gradients. For instance, using the normalised angle kernel, the primary gradient shifted to run between visual cortices and the rest of the brain, rather than providing a unimodal-transmodal distinction. Similarly, the secondary gradient under the normalised angle kernel ran between somatosensory versus frontal and superior and anterior temporal cortices, rather than between visual and somatosensory regions. A correlation similarity kernel produced gradients appearing between the cosine and normalised angle kernel solutions. Thus, although the unimodal-transmodal and visual-somatosensory gradients appear reproducible across studies and datasets, they may be less robust to changes in the connectopic mapping pipeline, and indeed this has been described previously (Bajada et al., 2020). Specifically, reproducing these specific gradients depends on using a dense affinity matrix representing a fully-connected network graph, and the gradients are sensitive to the choice of similarity kernel.

Functional topographies are unknown or poorly understood in many brain regions, so it is important to be confident that the gradients produced by connectopic mapping accurately reflect the genuine underlying topographies. Given the issues discussed above, how confident can we be in connectopic gradients? There are multiple features of previously identified connectopic maps that would be difficult to explain away purely with spatial autocorrelations or specific features of the analysis pipeline. Previously, we reported better reconstructions of retinotopic gradients from connectivity measured during movie watching than at rest (Watson and Andrews, 2022), demonstrating connectopic maps are influenced by the stimulus. Furthermore, individual variability in local gradients in both the hippocampus and striatum has been linked to human behaviour (Borne et al., 2023; Marquand et al., 2017; Przeździk et al., 2019). Meanwhile, the topology of whole-brain

gradients corresponds well with the organisation of large-scale resting-state networks (Margulies et al., 2016). Whole-brain gradients have also been shown to modulate with task demands (Gao et al., 2022; Samara et al., 2023; Zhang et al., 2022), correlate with brain microstructure and structural connectivity (Larivière et al., 2020), change with ageing (Bethlehem et al., 2020; Park et al., 2021a), and differ between neurotypical and neurodiverse populations (Hong et al., 2019; Lee et al., 2023; Park et al., 2021b; Xia et al., 2022). In some cases, these differences may interact with spatial autocorrelations – for instance different stimuli or tasks may evoke responses with differing signal to noise ratios which in turn alter the degree of spatial autocorrelation. Nevertheless, it seems likely that connectopic maps do at least partially reflect genuine underlying functional topographies and connectivity. However, it is also clear that connectopic maps can be strongly influenced by artificially induced spatial autocorrelations and specific features of the analysis pipeline. It may therefore be challenging to disentangle the influence of functional versus artificial signals on connectopic maps, thus the extent to which these gradients truly capture underlying functional topographies remains unclear. Connectopic mapping techniques may be able to reveal genuine functional gradients, but this may require further methods to account for spatial correlations within the data or at least a consideration of how such correlations may influence the resulting gradients. Comparing connectopic maps to behaviour or connectivity patterns with other brain regions may be informative (Marquand et al., 2017; Przeździk et al., 2019), though such measures remain correlational. Finally, while we discuss these issues in relation to connectopic mapping, it is possible that other analysis techniques concerned with the topographic organisation of the brain could be similarly affected.

In this study, we identified a potential flaw in connectopic mapping analyses. Local connectopic maps can be substantially biased by local spatial autocorrelations, for instance as introduced by spatial smoothing or interpolation between co-ordinate spaces, to the extent that illusory gradients can be produced even from random data. Gradients appear highly similar between real and random data, and replicate gradients previously reported in the literature. Coarser-scale whole-brain gradients can be less susceptible to local spatial autocorrelations, but nevertheless appear sensitive to specific features of the analysis pipeline. These results indicate that previously reported gradients from connectopic mapping studies may need to be interpreted with a degree of caution. Future connectopic mapping studies may aim to develop methods to address these issues, or provide a consideration of how gradients may be influenced them.

Data availability statement

All analysis code is available on the Open Science Framework (<https://osf.io/4dpuw/>). All MRI data were obtained from the already publicly available StudyForrest project (<https://www.studyforrest.org/>).

Declaration of Competing Interest

None.

Credit authorship contribution statement

David M. Watson: Conceptualization, Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Timothy J. Andrews:** Conceptualization, Methodology, Writing – original draft, Writing – review & editing, Supervision.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.neuroimage.2023.120228](https://doi.org/10.1016/j.neuroimage.2023.120228).

References

- Andersson, J.L.R., Jenkinson, M., Smith, S., 2010. Non-linear registration, aka spatial normalization (Publication No. TR07JA2). FMRIB Centre, Oxford, United Kingdom. <https://www.fmrib.ox.ac.uk/datasets/techrep/tr07ja2/tr07ja2.pdf>.
- Bajada, C.J., Costa Campos, L.Q., Caspers, S., Muscat, R., Parker, G.J.M., Lambon Ralph, M.A., Cloutman, L.L., Trujillo-Barreto, N.J., 2020. A tutorial and tool for exploring feature similarity gradients with MRI data. *Neuroimage* 221, 117140. doi:10.1016/j.neuroimage.2020.117140.
- Bakeman, R., 2005. Recommended effect size statistics for repeated measures designs. *Behav. Res. Methods* 37, 379–384. doi:10.3758/BF03192707.
- Bednar, J.A., Wilson, S.P., 2016. Cortical Maps. *Neurosci.* 22 (6), 604–617. doi:10.1177/1073858415597645.
- Bethlehem, R.A.I., Paquola, C., Seidlitz, J., Ronan, L., Bernhardt, B., Consortium, C.C., Tsvetanov, K.A., 2020. Dispersion of functional gradients across the adult lifespan. *Neuroimage* 222, 117299. doi:10.1016/j.neuroimage.2020.117299.
- Borne, L., Tian, Y., Lupton, M.K., van der Meer, J.N., Jeganathan, J., Paton, B., Koussis, N., Guo, C.C., Robinson, G.A., Frripp, J., Zalesky, A., Breakspear, M., 2023. Functional reorganization of hippocampal-cortical gradients during naturalistic memory processes. *Neuroimage* 271, 119996. doi:10.1016/j.neuroimage.2023.119996.
- Ciantar, K.G., Farrugia, C., Galdi, P., Scerri, K., Xu, T., Bajada, C.J., 2022. Geometric effects of volume-to-surface mapping of fMRI data. *Brain Struct. Funct.* 227, 2457–2464. doi:10.1007/s00429-022-02536-4.
- Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* 29, 162–173.
- Cross, N., Paquola, C., Pomares, F.B., Perrault, A.A., Jegou, A., Nguyen, A., Aydin, U., Bernhardt, B.C., Grova, C., Dang-Vu, T.T., 2021. Cortical gradients of functional connectivity are robust to state-dependent changes following sleep deprivation. *Neuroimage* 226, 117547. doi:10.1016/j.neuroimage.2020.117547.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis I. segmentation and surface reconstruction. *Neuroimage* 9, 179–194. doi:10.1006/nimg.1998.0395.
- Engel, S., Glover, G.H., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex* 7, 181–192. doi:10.1093/cercor/7.2.181.
- Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., Dale, A.M., 2002. Whole Brain Segmentation. *Neuron* 33, 341–355. doi:10.1016/S0896-6273(02)00569-X.
- Fischl, B., Sereno, M.I., Dale, A.M., 1999a. Cortical surface-based analysis II: inflation, flattening, and a surface-based coordinate system. *Neuroimage* 9, 195–207. doi:10.1006/nimg.1998.0396.
- Fischl, B., Sereno, M.I., Tootell, R.B.H., Dale, A.M., 1999b. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.* 8, 272–284. doi:10.1002/(SICI)1097-0193(1999)8:4<272::AID-HBM10>3.0.CO;2-4.
- Formisano, E., Kim, D.S., Di Salle, F., van de Moortele, P.F., Ugurbil, K., Goebel, R., 2003. Mirror-symmetric tonotopic maps in human primary auditory cortex. *Neuron* 40, 859–869. doi:10.1016/S0896-6273(03)00669-X.
- Gao, Z., Zheng, L., Krieger-Redwood, K., Halai, A., Margulies, D.S., Smallwood, J., Jeffries, E., 2022. Flexing the principal gradient of the cerebral cortex to suit changing semantic task demands. *Elife* 11. doi:10.7554/eLife.80368.
- 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., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536, 171–178. doi:10.1038/nature18933.
- Greve, D.N., Fischl, B., 2009. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* 48, 63–72. doi:10.1016/j.neuroimage.2009.06.060.
- Haak, K.V., Beckmann, C.F., 2020. Understanding brain organisation in the face of functional heterogeneity and functional multiplicity. *Neuroimage* 220, 117061. doi:10.1016/j.neuroimage.2020.117061.
- Haak, K.V., Marquand, A.F., Beckmann, C.F., 2018. Connectopic mapping with resting-state fMRI. *Neuroimage* 170, 83–94. doi:10.1016/j.neuroimage.2017.06.075.
- Hanke, M., Adelhöfer, N., Kottke, D., Iacovella, V., Sengupta, A., Kaule, F.R., Nigbur, R., Waite, A.Q., Baumgartner, F., Stadler, J., 2016. A study forrest extension, simultaneous fMRI and eye gaze recordings during prolonged natural stimulation. *Sci. Data* 3, 160092. doi:10.1038/sdata.2016.92.
- Hanke, M., Baumgartner, F.J., Ibe, P., Kaule, F.R., Pollmann, S., Speck, O., Zinke, W., Stadler, J., 2014. A high-resolution 7-Tesla fMRI dataset from complex natural stimulation with an audio movie. *Sci. Data* 1, 140003. doi:10.1038/sdata.2014.3.
- Hanke, M., Halchenko, Y.O., Sederberg, P.B., Hanson, S.J., Haxby, J.V., Pollmann, S., 2009. PyMPPA: a python toolbox for multivariate pattern analysis of fMRI data. *Neuroinformatics* 7, 37–53. doi:10.1007/s12021-008-9041-y.
- Hong, S.J., Vos de Wael, R., Bethlehem, R.A.I., Larivière, S., Paquola, C., Valk, S.L., Milham, M.P., Di Martino, A., Margulies, D.S., Smallwood, J., Bernhardt, B.C., 2019. Atypical functional connectome hierarchy in autism. *Nat. Commun.* 10, 1022. doi:10.1038/s41467-019-08944-1.
- Huntenburg, J.M., Bazin, P.L., Margulies, D.S., 2018. Large-scale gradients in human cortical organization. *Trends Cogn. Sci.* 22, 21–31. doi:10.1016/j.tics.2017.11.002.
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17, 825–841. doi:10.1006/nimg.2002.1132.
- Jenkinson, M., Beckmann, C.F., Behrens, T.E.J., Woolrich, M.W., Smith, S.M., 2012. FSL. *Neuroimage* 62, 782–790. doi:10.1016/j.neuroimage.2011.09.015.
- Kriegeskorte, N., Goebel, R., Bandettini, P., 2006. Information-based functional brain mapping. *Proc. Natl. Acad. Sci.* 103 (10), 3863–3868. doi:10.1073/pnas.0600244103.
- Larivière, S., Vos de Wael, R., Hong, S.J., Paquola, C., Tavakol, S., Lowe, A.J., Schrader, D.V., Bernhardt, B.C., 2020. Multiscale structure–function gradients in the neonatal connectome. *Cereb. Cortex* 30, 47–58. doi:10.1093/cercor/bhz069.
- Lee, C.H., Park, H., Lee, M.J., Park, B., 2023. Whole-brain functional gradients reveal cortical and subcortical alterations in patients with episodic migraine. *Hum. Brain Mapp.* 44, 2224–2233. doi:10.1002/hbm.26204.
- Margulies, D.S., Ghosh, S.S., Goulas, A., Falkiewicz, M., Huntenburg, J.M., Langs, G., Bezdin, G., Eickhoff, S.B., Castellanos, F.X., Petrides, M., Jefferies, E., Smallwood, J., 2016. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc. Natl. Acad. Sci.* 113, 12574–12579. doi:10.1073/pnas.1608282113.
- Margulies, D.S., Falkiewicz, M., Huntenburg, J.M., 2016. A cortical surface-based geodesic distance package for Python. *GigaScience* 5 (suppl_1), 1–2. doi:10.1186/s13742-016-0147-0-q.
- Marquand, A.F., Haak, K.V., Beckmann, C.F., 2017. Functional corticostriatal connection topographies predict goal-directed behaviour in humans. *Nat. Hum. Behav.* 1, 0146. doi:10.1038/s41562-017-0146-0.
- Navarro Schröder, T., Haak, K.V., Zaragoza Jimenez, N.I., Beckmann, C.F., Doeller, C.F., 2015. Functional topography of the human entorhinal cortex. *Elife* 4. doi:10.7554/eLife.06738.
- Ngo, G.N., Haak, K.V., Beckmann, C.F., Menon, R.S., 2021. Mesoscale hierarchical organization of primary somatosensory cortex captured by resting-state-fMRI in humans. *Neuroimage* 235, 118031. doi:10.1016/j.neuroimage.2021.118031.
- Olejnik, S., Algina, J., 2003. Generalized eta and omega squared statistics: measures of effect size for some common research designs. *Psychol. Methods* 8, 434–447. doi:10.1037/1082-989X.8.4.434.
- Paquola, C., Vos de Wael, R., Wagstyl, K., Bethlehem, R.A.I., Hong, S.J., Seidlitz, J., Bullmore, E.T., Evans, A.C., Misić, B., Margulies, D.S., Smallwood, J., Bernhardt, B.C., 2019. Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLOS Biol.* 17, e3000284. doi:10.1371/journal.pbio.3000284.
- Park, B., Bethlehem, R.A.I., Paquola, C., Larivière, S., Rodríguez-Cruces, R., Vos de Wael, R., Bullmore, E., Dolan, R., Goodyer, I., Fonagy, P., Jones, P., Moutoussis, K., Hauser, T., Neufeld, S., Romero-Garcia, R., St Clair, M., Vértes, P., Whitaker, K., Inkster, B., Prabhu, G., Ooi, C., Toseeb, U., Widmer, B., Bhatti, J., Villis, L., Alru-maithi, A., Birt, S., Bowler, A., Cleridou, K., Dadabhoy, H., Davies, E., Firkins, A., Granville, S., Harding, E., Hopkins, A., Isaacs, D., King, J., Kokorikou, D., Maurice, C., McIntosh, C., Memarzia, J., Mills, H., O'Donnell, C., Pantaleone, S., Scott, J., Kiddle, B., Polek, E., Fearon, P., Suckling, J., van Harmelen, A.L., Kievit, R., Chamberlain, S., Bullmore, E.T., Bernhardt, B.C., 2021a. An expanding manifold in transmodal regions characterizes adolescent reconfiguration of structural connectome organization. *Elife* 10, 1–32. doi:10.7554/eLife.64694.
- Park, B., Hong, S.J., Valk, S.L., Paquola, C., Benkarim, O., Bethlehem, R.A.I., Di Martino, A., Milham, M.P., Gozzi, A., Yeo, B.T.T., Smallwood, J., Bernhardt, B.C., 2021b. Differences in subcortico-cortical interactions identified from connectome and microcircuit models in autism. *Nat. Commun.* 12, 2225. doi:10.1038/s41467-021-21732-0.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine Learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Przeźdźik, I., Faber, M., Fernández, G., Beckmann, C.F., Haak, K.V., 2019. The functional organisation of the hippocampus along its long axis is gradual and predicts recollection. *Cortex* 119, 324–335. doi:10.1016/j.cortex.2019.04.015.
- Saad, Z.S., Repella, K.M., Cox, R.W., DeYoe, E.A., 2001. Analysis and use of fMRI response delays. *Hum. Brain Mapp.* 13, 74–93. doi:10.1002/hbm.1026.
- Samara, A., Eilbott, J., Margulies, D.S., Xu, T., Vandervel, T., 2023. Cortical gradients during encephalic processing are hierarchical and modality-specific. *Neuroimage* 271, 120023. doi:10.1016/j.neuroimage.2023.120023.
- Sanchez-Panchuelo, R.M., Francis, S., Bowtell, R., Schluppeck, D., 2010. Mapping human somatosensory cortex in individual subjects with 7T functional MRI. *J. Neurophysiol.* 103, 2544–2556. doi:10.1152/jn.01017.2009.
- Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T.T., 2018. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. *Cereb. Cortex* 28, 3095–3114. doi:10.1093/cercor/bhx179.
- Smith, S.M., 2002. Fast robust automated brain extraction. *Hum. Brain Mapp.* 17, 143–155. doi:10.1002/hbm.10062.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23, S208–S219. doi:10.1016/j.neuroimage.2004.07.051.
- von Luxburg, U., 2007. A tutorial on spectral clustering. *Stat. Comput.* 17, 395–416. doi:10.1007/s11222-007-9033-z.
- Vos de Wael, R., Benkarim, O., Paquola, C., Larivière, S., Royer, J., Tavakol, S., Xu, T., Hong, S.J., Langs, G., Valk, S., Misić, B., Milham, M., Margulies, D., Smallwood, J., Bernhardt, B.C., 2020. BrainSpace: a toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. *Commun. Biol.* 3, 103. doi:10.1038/s42003-020-0794-7.
- Vos de Wael, R., Larivière, S., Caldauro, B., Hong, S.J., Margulies, D.S., Jefferies, E., Bernasconi, A., Smallwood, J., Bernasconi, N., Bernhardt, B.C., 2018. Anatomical and microstructural determinants of hippocampal subfield functional connectome embedding. *Proc. Natl. Acad. Sci.* 115, 10154–10159. doi:10.1073/pnas.1803667115.
- Wandell, B.A., Winawer, J., 2011. Imaging retinotopic maps in the human brain. *Vision Res.* 51, 718–737. doi:10.1016/j.visres.2010.08.004.
- Wang, R., Mo, F., Shen, Y., Song, Y., Cai, H., Zhu, J., 2022. Functional connectivity gradients of the insula to different cerebral systems. *Hum. Brain Mapp.* doi:10.1002/hbm.26099.

- Watson, D.M., Andrews, T.J., 2022. An evaluation of how connectopic mapping reveals visual field maps in V1. *Sci. Rep.* 12, 16249. doi:[10.1038/s41598-022-20322-4](https://doi.org/10.1038/s41598-022-20322-4).
- Xia, M., Liu, J., Mechelli, A., Sun, X., Ma, Q., Wang, X., Wei, D., Chen, Y., Liu, B., Huang, C.C., Zheng, Y., Wu, Y., Chen, T., Cheng, Y., Xu, X., Gong, Q., Si, T., Qiu, S., Lin, C.P., Cheng, J., Tang, Y., Wang, F., Qiu, J., Xie, P., Li, L., He, Y., 2022. Connectome gradient dysfunction in major depression and its association with gene expression profiles and treatment outcomes. *Mol. Psychiatry* 27, 1384–1393. doi:[10.1038/s41380-022-01519-5](https://doi.org/10.1038/s41380-022-01519-5).
- Zhang, H., Zhao, R., Hu, X., Guan, S., Margulies, D.S., Meng, C., Biswal, B.B., 2022. Cortical connectivity gradients and local timescales during cognitive states are modulated by cognitive loads. *Brain Struct. Funct.* doi:[10.1007/s00429-022-02564-0](https://doi.org/10.1007/s00429-022-02564-0).
- Zhi, D., King, M., Hernandez-Castillo, C.R., Diedrichsen, J., 2022. Evaluating brain parcellations using the distance-controlled boundary coefficient. *Hum. Brain Mapp.* 1–15. doi:[10.1002/hbm.25878](https://doi.org/10.1002/hbm.25878).