




OPEN

## An evaluation of how connectopic mapping reveals visual field maps in V1


David M. Watson  & Timothy J. Andrews

Functional gradients, in which response properties change gradually across the cortical surface, have been proposed as a key organising principle of the brain. However, the presence of these gradients remains undetermined in many brain regions. Resting-state neuroimaging studies have suggested these gradients can be reconstructed from patterns of functional connectivity. Here we investigate the accuracy of these reconstructions and establish whether it is connectivity or the functional properties within a region that determine these “connectopic maps”. Different manifold learning techniques were used to recover visual field maps while participants were at rest or engaged in natural viewing. We benchmarked these reconstructions against maps measured by traditional visual field mapping. We report an initial exploratory experiment of a publicly available naturalistic imaging dataset, followed by a preregistered replication using larger resting-state and naturalistic imaging datasets from the Human Connectome Project. Connectopic mapping accurately predicted visual field maps in primary visual cortex, with better predictions for eccentricity than polar angle maps. Non-linear manifold learning methods outperformed simpler linear embeddings. We also found more accurate predictions during natural viewing compared to resting-state. Varying the source of the connectivity estimates had minimal impact on the connectopic maps, suggesting the key factor is the functional topography within a brain region. The application of these standardised methods for connectopic mapping will allow the discovery of functional gradients across the brain.

**Protocol registration** The stage 1 protocol for this Registered Report was accepted in principle on 19 April 2022. The protocol, as accepted by the journal, can be found at <https://doi.org/10.6084/m9.figshare.19771717>.

Functional gradients are an important organising principle in the brain<sup>1</sup>. The key feature underlying these gradients is a gradual change in preferred stimulus or response parameters that is topographically mapped across brain regions. Gradients are well described in primary sensory regions, such as retinotopic maps in visual cortex<sup>2–4</sup>, somatotopic maps in somatosensory cortex<sup>5,6</sup>, and tonotopic maps in auditory cortex<sup>7–9</sup>. Functional gradients have also been proposed outside of primary sensory regions, including topographic maps of real-world object size<sup>10</sup> and animacy<sup>11</sup> in the ventral visual stream, and maps of numerosity<sup>12</sup> and multisensory integration<sup>13</sup> in the dorsal visual stream. Nevertheless, the majority of research to date has regarded regions outside of primary sensory cortices as comprising distinct parcels or modules<sup>14</sup>, and the role of underlying functional gradients has been frequently overlooked (although see<sup>15,16</sup>). Consequently, functional gradients throughout many brain regions are poorly understood, while others yet may remain entirely undiscovered. While there are established paradigms for measuring maps in primary sensory regions, such as visual field mapping in visual cortex<sup>17,18</sup>, it is often unclear what tasks would be required to estimate gradients in other regions.

The function of a brain region is closely tied to its connectivity with other regions, such that functionally similar neurons may also be expected to show similar connectivity patterns<sup>19,20</sup>. Consequently, the topographic organisation of brain functions may be similarly reflected in a topographic organisation of connectivity<sup>21</sup>, yielding so-called “connectopic” maps<sup>22,23</sup>. As such, functionally similar points within a brain region may be represented close to one another within the corresponding connectivity space, while functionally dissimilar points may appear more distant. Thus, each of these points may be distributed along a manifold within the connectivity space, with the principal dimensions of that manifold corresponding to the principal functional gradients underpinning that

Department of Psychology and York Neuroimaging Centre, University of York, York YO10 5DD, UK. email: david.watson@york.ac.uk

brain region. This implies that functional gradients may be reconstructed from connectivity patterns by extracting the principal dimensions from within the connectivity space. Particular focus has been given to manifold learning techniques that extract non-linear dimensions<sup>24,25</sup>. By applying manifold learning to resting-state functional connectivity patterns, previous studies have identified coarse-scale gradients spanning the whole brain<sup>16,25</sup>, as well as local gradients within brain regions including the striatum<sup>26</sup>, hippocampus<sup>27</sup>, and primary sensory regions<sup>22,28</sup>. The same techniques can also be applied to recover functional gradients from structural connectivity<sup>29</sup>. Thus, connectopic mapping provides a potential method for characterising functional gradients without necessarily requiring prior knowledge of those gradients or a specific task to elicit them.

Nevertheless, there remain a number of open questions regarding the efficacy and practicalities of connectopic mapping. Firstly, previous studies have not provided ground-truth estimates of the functional gradients to benchmark the connectopic maps against. For instance, while Haak and colleagues<sup>22</sup> demonstrated that connectopic mapping can reconstruct retinotopic maps in primary visual cortex, they did not have access to maps estimated directly from traditional visual field mapping techniques to compare these reconstructions against. Consequently, it remains unclear how accurately the functional gradients are reconstructed in each individual. Secondly, previous connectopic mapping studies have exclusively measured functional connectivity in the brain at rest. While resting-state estimates are a common choice, they have also been criticised as representing an unnatural cognitive state that may not generalise well to more active states<sup>30</sup>. An alternative approach is to measure functional connectivity during naturalistic stimulation (e.g. movie-watching)<sup>31</sup>, and indeed such paradigms may yield more reliable connectivity patterns in some circumstances<sup>32,33</sup>. Finally, many algorithms exist for extracting the principal dimensions of the connectivity space, yet a comprehensive comparison between them remains lacking. Previous studies have primarily employed spectral embedding<sup>22,26,28,29</sup> and diffusion map<sup>16,27</sup> manifold learning techniques, but many other such techniques exist. Manifold learning techniques will likely outperform linear embeddings if the representation is non-linear, but linear techniques are typically simpler and may still perform well if the representation is linear.

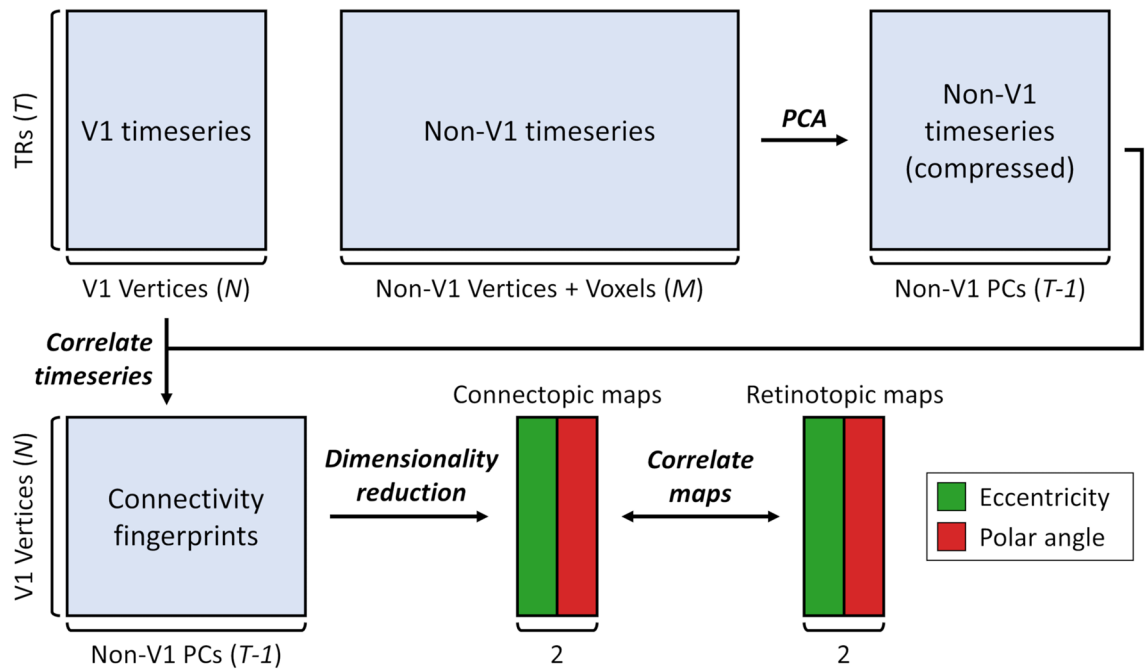
Here, we provide a comprehensive assessment of connectopic mapping for reconstructing local functional gradients, using the retinotopic maps in primary visual cortex as a test case. We compare the performance of multiple commonly used dimensionality reduction techniques by benchmarking the reconstructions against retinotopic maps derived from traditional visual field mapping techniques. We give particular focus to spectral embedding as this has previously been applied to reconstructing gradients in sensory regions<sup>22,28</sup>, though we also consider a number of other linear embeddings and manifold learning algorithms. In an initial experiment, we perform exploratory analyses of a publicly available naturalistic-imaging dataset<sup>34</sup>. We then provide a preregistered replication experiment using both resting-state and naturalistic-imaging data obtained from the Human Connectome Project<sup>35</sup>. If connectopic mapping provides a robust reconstruction of the underlying functional gradients, we would expect a high degree of correspondence between the maps estimated from connectivity and those measured directly by visual field mapping. By developing standardised connectivity methods, it will be possible to determine whether topographic maps are a ubiquitous organising principle across the brain.

## Results

**Experiment 1.** We first performed exploratory analyses of 15 subjects obtained from the publicly available StudyForrest naturalistic-imaging MRI dataset<sup>34,36,37</sup> (<https://www.studyforrest.org/>). This includes approximately 2 h of movie-watching data plus retinotopic mapping scans for each subject. We first extracted visual field maps in each subject via a travelling-wave analysis<sup>17</sup> of the retinotopy data. These were used to define individualised V1 regions of interest (ROIs) on the cortical surface, identified by tracing along the phase reversals in the polar angle map. The eccentricity and polar angle phase maps within V1 also served as ground-truth estimates of the visual field maps, which the connectopic maps could be compared against.

Next, we performed connectopic mapping within each V1 ROI in each subject, adapting the methods of Haak and colleagues<sup>22</sup>. The analysis pipeline is illustrated in Fig. 1. First, the movie-watching data were split into odd and even scan runs to allow cross-validated parameter selection for the dimensionality reduction algorithms (where necessary). For each data split, timeseries were concatenated over scan runs yielding approximately 1 h of data per split. We extracted the timeseries within all subcortical voxels and cortical surface vertices. These were then split into those surface vertices within the V1 ROI versus all surface vertices and subcortical voxels outside the ROI. Because the number of the non-ROI vertices and voxels exceeds the number of timepoints, we losslessly compressed this data via principal component analysis, retaining one fewer components than the number of timepoints so that 100% of the variance remained explained. This amounts to rotating the samples within the feature space and removing the unused dimensions, and aids the computational tractability of later processing stages. The timeseries were then correlated between all pairwise combinations of V1 vertices and non-V1 principal components. This yielded a matrix of connectivity fingerprints representing the functional connectivity profile of each V1 vertex.

The intuition of connectopic mapping is that functional gradients within the ROI correspond to principal dimensions of variation within the connectivity space. Gradients can be reconstructed by embedding the high-dimensional connectivity fingerprints into a lower-dimensional space and extracting the initial components. Here, we retained the first two components, which are expected to correspond to the eccentricity and polar angle maps. Where necessary, parameter selection for a given algorithm was performed via cross-validation using a Bayesian optimisation routine to maximise the prediction accuracy. Figure 2 illustrates retinotopic and connectopic maps in an example subject. All of the dimensionality reduction techniques captured a component following an anterior–posterior gradient which corresponded well to the retinotopic eccentricity map (Fig. 2a), as well as a component following an inferior–superior gradient which corresponded to the retinotopic polar angle map (Fig. 2b).




**Figure 1.** Connectopic mapping pipeline. Functional MRI timeseries, comprising  $T$  timepoints, are split between the  $N$  vertices within V1 versus the  $M$  cortical vertices and subcortical voxels outside V1. The non-V1 timeseries are reduced in dimensionality via a lossless principal component analysis, retaining one fewer components than the number of timepoints ( $T-1$ ). The V1 and compressed non-V1 timeseries are then correlated to derive the connectivity space. Dimensionality reduction is applied to this space and two components are retained, yielding the connectopic maps which are expected to correspond to the eccentricity and polar angle maps respectively. Prediction accuracy is assessed by correlating the connectopic maps with retinotopic maps estimated by traditional visual field mapping.

We next tested the prediction accuracies of each algorithm by correlating the connectopic maps estimated from the movie-watching with the ground-truth retinotopic maps estimated from the visual field mapping (Fig. 3). Because there is a sign ambiguity in the connectopic maps we took the absolute correlation values. We first considered the spectral embedding algorithm, which has previously been shown to reconstruct gradients in sensory regions including primary visual cortex<sup>22</sup>. We ran variants based on both weighted and unweighted versions of both radius and nearest neighbours (kNN) graphs, as well as a weighted fully-connected graph (see Methods). The resulting prediction accuracies for each variant are illustrated in Fig. 3a. Accuracies appeared highest for the nearest neighbour approach, and did not appear substantially different between weighted and unweighted variants. In all cases, correlations were significantly greater than zero (one-sample t-tests; all  $p < 0.001$ ). To compare between the variants, we entered the correlations into a two-way repeated-measures ANOVA with factors of map (eccentricity, polar angle) and graph type. This revealed a significant main effect of map type due to overall higher correlations for eccentricity than polar angle ( $F(1, 14) = 8.20, p = 0.013, \eta_p^2 = 0.37, \eta_G^2 = 0.18$ ). There was also a significant main effect of graph type ( $F(1.59, 22.20) = 16.18, p < 0.001, \eta_p^2 = 0.54, \eta_G^2 = 0.11$ ); post-hoc Tukey contrasts revealed that correlations for the nearest neighbour variants were significantly higher than the radius and fully-connected variants (all  $p < 0.001$ ), while there were no significant differences between the radius and fully-connected variants (all  $p > 0.05$ ), nor between weighted and unweighted variants (all  $p > 0.05$ ). Finally, there was no significant interaction ( $F(1.94, 27.19) = 2.43, p = 0.108, \eta_p^2 = 0.15, \eta_G^2 = 0.01$ ). Thus, all variants of spectral embedding predicted the retinotopic maps above chance, but the nearest neighbour approach performed best, and there was no clear benefit to weighting the graphs.

We next considered a number of other commonly used dimensionality reduction techniques (Fig. 3b), including linear embeddings (PCA and ICA) and various other manifold learning approaches: kernel PCA, locally linear embedding (LLE), local tangent space alignment (LTSA), diffusion maps, Isomap, and t-distributed stochastic neighbourhood embedding (t-SNE). For the kernel PCA, we employed two commonly used kernels: a second-order polynomial kernel, and a radial basis function (RBF) kernel. Again, one-sample t-tests revealed that correlations with the retinotopic maps were significantly greater than zero in all cases (one-sample t-tests; all  $p < 0.001$ ). Furthermore, paired-samples t-tests revealed significantly higher correlations for the eccentricity than the polar angle maps for all algorithms (all  $p < 0.05$ ) except Isomap ( $t(14) = 1.27, p = 0.225$ ) and t-SNE ( $t(14) = 1.66, p = 0.119$ ).

Prediction accuracies generally appeared higher for the manifold learning than the linear embedding methods. To compare accuracies between algorithms, we entered the correlations into a two-way repeated measures ANOVA with factors of map and algorithm type. For spectral embedding, we selected the unweighted nearest neighbour approach as one of the better performing variants. The main effect of map type was significant ( $F(1,$



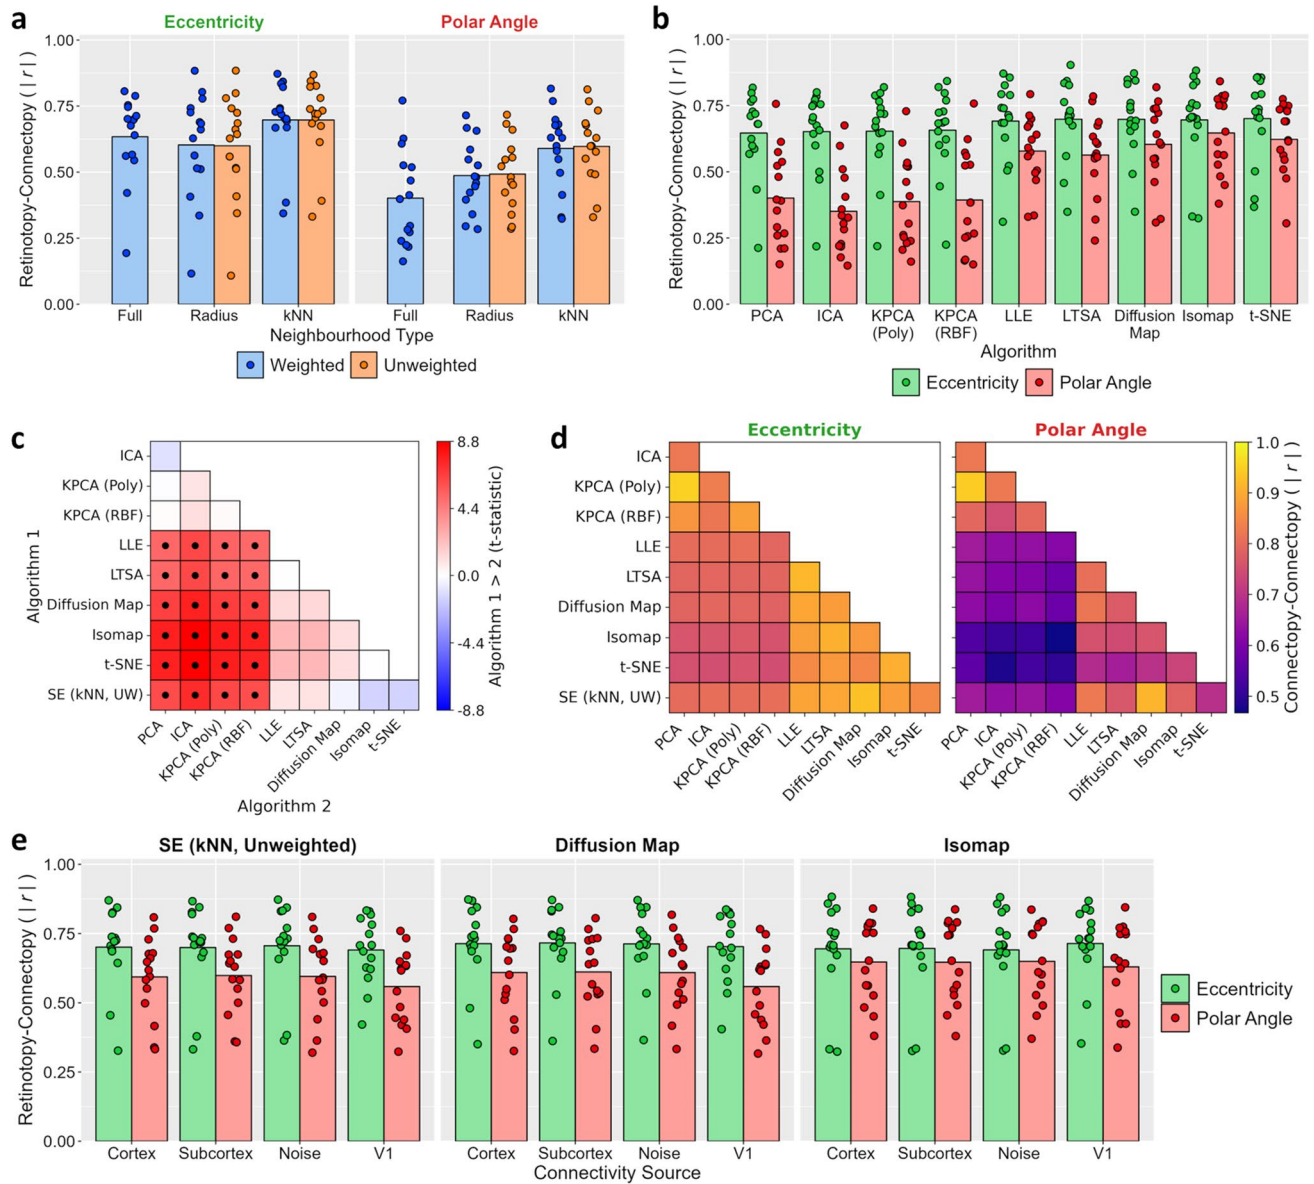
**Figure 2.** Experiment 1: Retinotopic and connectopic maps of (a) eccentricity and (b) polar angle in the left hemisphere of an example subject. The top-left plot in each panel illustrates retinotopic maps measured by traditional visual field mapping techniques. The remaining plots illustrate maps reconstructed from connectivity patterns measured during movie-watching. Spectral embedding is illustrated for the unweighted nearest neighbour variant.

14) = 11.00,  $p = 0.005$ ,  $\eta_p^2 = 0.44$ ,  $\eta_G^2 = 0.24$ ) due to higher correlations for eccentricity than polar angle maps. There was also a significant main effect of the algorithm type ( $F(3.01, 42.08) = 25.23$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.64$ ,  $\eta_G^2 = 0.16$ ); post-hoc Tukey contrasts (Fig. 3c) revealed that prediction accuracies were significantly lower for the PCA, ICA, and kernel PCA algorithms than for the other manifold learning techniques (all  $p < 0.001$ ). There was also an interaction between the map and algorithm type ( $F(2.47, 34.61) = 6.61$ ,  $p = 0.002$ ,  $\eta_p^2 = 0.32$ ,  $\eta_G^2 = 0.05$ ) reflecting a greater difference between algorithms for the polar angle than for the eccentricity maps. We also assessed the relative similarities of the algorithms by correlating the connectopic maps between them (Fig. 3d). Similar to the prediction accuracies, this indicated that the algorithms fell into two broad groups: the linear embeddings and kernel PCAs produced relatively similar gradients to one another, and the remaining manifold learning methods also appeared similar to one another, but gradients were less similar between the two groups. In summary, the manifold learning approaches generally outperformed the linear approaches (PCA and ICA), especially when predicting the polar angle maps. Both variants of the kernel PCA (polynomial and RBF) performed more poorly and did not differ significantly from the linear PCA, suggesting that neither of these kernels were able to sufficiently capture the non-linearities in the data.

In all of the previous analyses, the connectivity fingerprints were estimated between V1 vertices and all cortical vertices plus subcortical voxels throughout the rest of the brain. To test how the source of connectivity information affects the estimated connectopic maps, we recalculated the connectivity fingerprints by correlating the V1 timeseries with the timeseries from either just the cortical vertices or just the subcortical voxels. We repeated our analyses with the spectral embedding (unweighted nearest neighbour), diffusion map, and Isomap algorithms based on these new fingerprints (Fig. 3e). Prediction accuracies remained high (one-sample t-tests: all  $p < 0.001$ ) and did not appear substantially different to the original analyses. This was surprising as we expected the connectopic map with the cortical vertices to be a better predictor of the functional organisation.

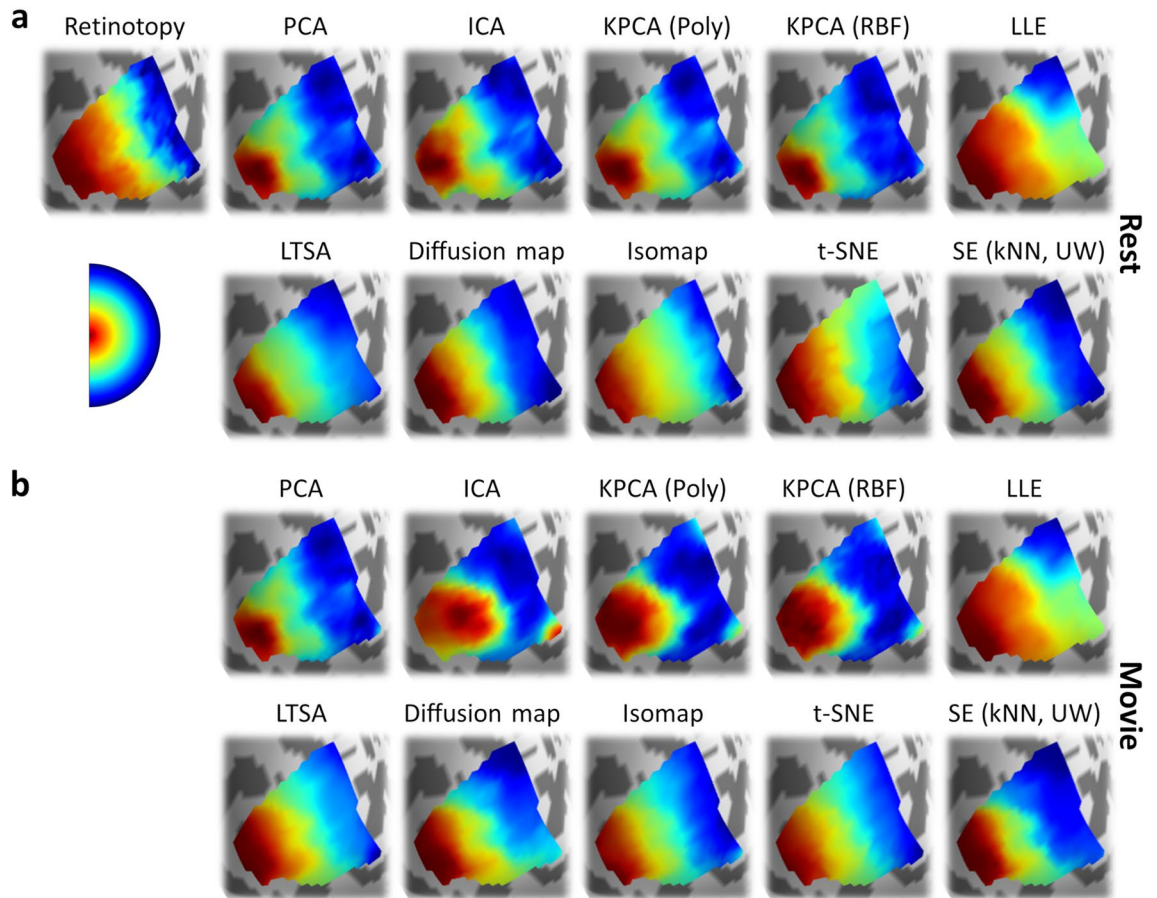
A possible explanation for this result is that the key factor in the connectopic mapping is the topographic organisation of the functional responses within the region of interest, rather than connectivity with the rest of the brain. To test this possibility, we performed two further variants of our analyses. Firstly, we estimated connectivity fingerprints using signals from all cortical vertices plus subcortical voxels (as per the main analyses), but first substituted the non-V1 timeseries with normally distributed noise matched in mean and variance to the original timeseries. Secondly, we estimated within-ROI fingerprints by correlating each of the V1 timeseries





**Figure 3.** Experiment 1 results. (a) Prediction accuracies for all variants of spectral embedding, measured by absolute correlations between retinotopic and connectopic maps. Dot markers indicate per-subject means and bars indicate group means. (b) Prediction accuracies for other algorithms. (c) Post-hoc Tukey contrasts of prediction accuracies between algorithms. For spectral embedding, the unweighted nearest neighbour variant is selected. Dot markers indicate significant contrasts ( $p < 0.05$ ). (d) Correlations of connectopic maps between algorithms. (e) Prediction accuracies for unweighted nearest neighbour variant of spectral embedding, diffusion map, and Isomap algorithms when varying the source of the connectivity estimates: V1 timeseries are correlated with timeseries taken from just non-V1 cortical vertices, just subcortical voxels, non-V1 cortical vertices plus subcortical voxels substituted with Gaussian noise, or the same V1 timeseries.

with each other, entirely omitting timeseries outside of V1. Again, both variants produced similar results to the original analyses (Fig. 3e). In both cases, prediction accuracies remained high (one-sample t-tests: all  $p < 0.001$ ). We entered the correlations into a three-way repeated-measures ANOVA with factors of map and algorithm type, and connectivity source (cortex + subcortex, cortex, subcortex, noise, and within-V1). This revealed significant main effects of map type ( $F(1, 14) = 6.06, p = 0.027, \eta_p^2 = 0.30, \eta_G^2 = 0.13$ ) and algorithm type ( $F(1.14, 16.01) = 8.31, p = 0.008, \eta_p^2 = 0.37, \eta_G^2 = 0.01$ ); post-hoc Tukey contrasts revealed overall higher correlations for Isomap than spectral embedding ( $p = 0.001$ ) and diffusion maps ( $p = 0.028$ ). However, the main effect of the connectivity source was not significant ( $F(1.20, 16.87) = 3.01, p = 0.095, \eta_p^2 = 0.18, \eta_G^2 < 0.01$ ). Finally, there were significant interactions between the map and algorithm type ( $F(1.20, 16.79) = 6.35, p = 0.018, \eta_p^2 = 0.31, \eta_G^2 = 0.01$ ) and between the algorithm type and connectivity source ( $F(2.72, 38.04) = 5.92, p = 0.003, \eta_p^2 = 0.30, \eta_G^2 < 0.01$ ). In summary, varying the source of the connectivity information had minimal impact on the connectopic maps



**Figure 4.** Experiment 2: Retinotopic and connectopic eccentricity maps in the left hemisphere of an example subject. (a) Retinotopic maps (top-left plot) and connectopic maps reconstructed from resting-state data (remaining plots). (b) Connectopic maps reconstructed from movie-watching data. Spectral embedding is illustrated for the unweighted nearest neighbour variant.


and in all cases prediction accuracies remained high. This suggests the principal source of information in ‘connectopic’ mapping is the functional topography within the ROI.

Finally, we tested the stability of the connectopic mapping algorithms by correlating the connectopic maps between the two cross-validation splits (Supplementary Fig. S1). The split-half correlations appeared high in all cases, indicating good internal reliability within each algorithm and suggesting that the connectopic maps are at least partially independent from the precise content of the movie stimulus.

**Experiment 2.** Our first experiment reported exploratory analyses of one naturalistic imaging dataset. To determine the robustness of these results we next conducted a preregistered replication experiment using a larger 7T MRI dataset obtained from the Human Connectome Project<sup>35</sup>. This additionally allowed us to compare connectopic mapping between connectivity measured during both rest and movie watching. The dataset comprised 174 subjects each with approximately one hour of resting-state and movie-watching data in addition to visual field mapping scans.

We first extracted visual field maps in each subject via population receptive field modelling<sup>18</sup>. As before, these maps were used to define individualised V1 ROIs as well as forming ground-truth estimates to compare the connectopic maps against. We then performed connectopic mapping in each subject, following the same methods as for Experiment 1. Retinotopic and connectopic maps of eccentricity (Fig. 4) and polar angle (Fig. 5) are shown for an example subject during the resting-state and movie-watching tasks. All of the dimensionality reduction techniques captured an anterior–posterior gradient which corresponded well to the retinotopic eccentricity map (Fig. 4), and most captured an inferior–superior gradient corresponding to the retinotopic polar angle map (Fig. 5).

We next tested the prediction accuracies of each algorithm by taking the absolute correlations between the retinotopic and connectopic maps (Fig. 6). We first considered the spectral embedding algorithm (Fig. 6a). As per Experiment 1, accuracies appeared highest for the nearest neighbour approach and were minimally affected by weighting the graphs. In addition, accuracies appeared higher for movie-watching than resting-state estimates, particularly for polar angle maps. In all cases, correlations were significantly greater than zero (one-sample t-tests; all  $p < 0.001$ ). To compare between the variants, we entered the correlations into a three-way repeated-measures



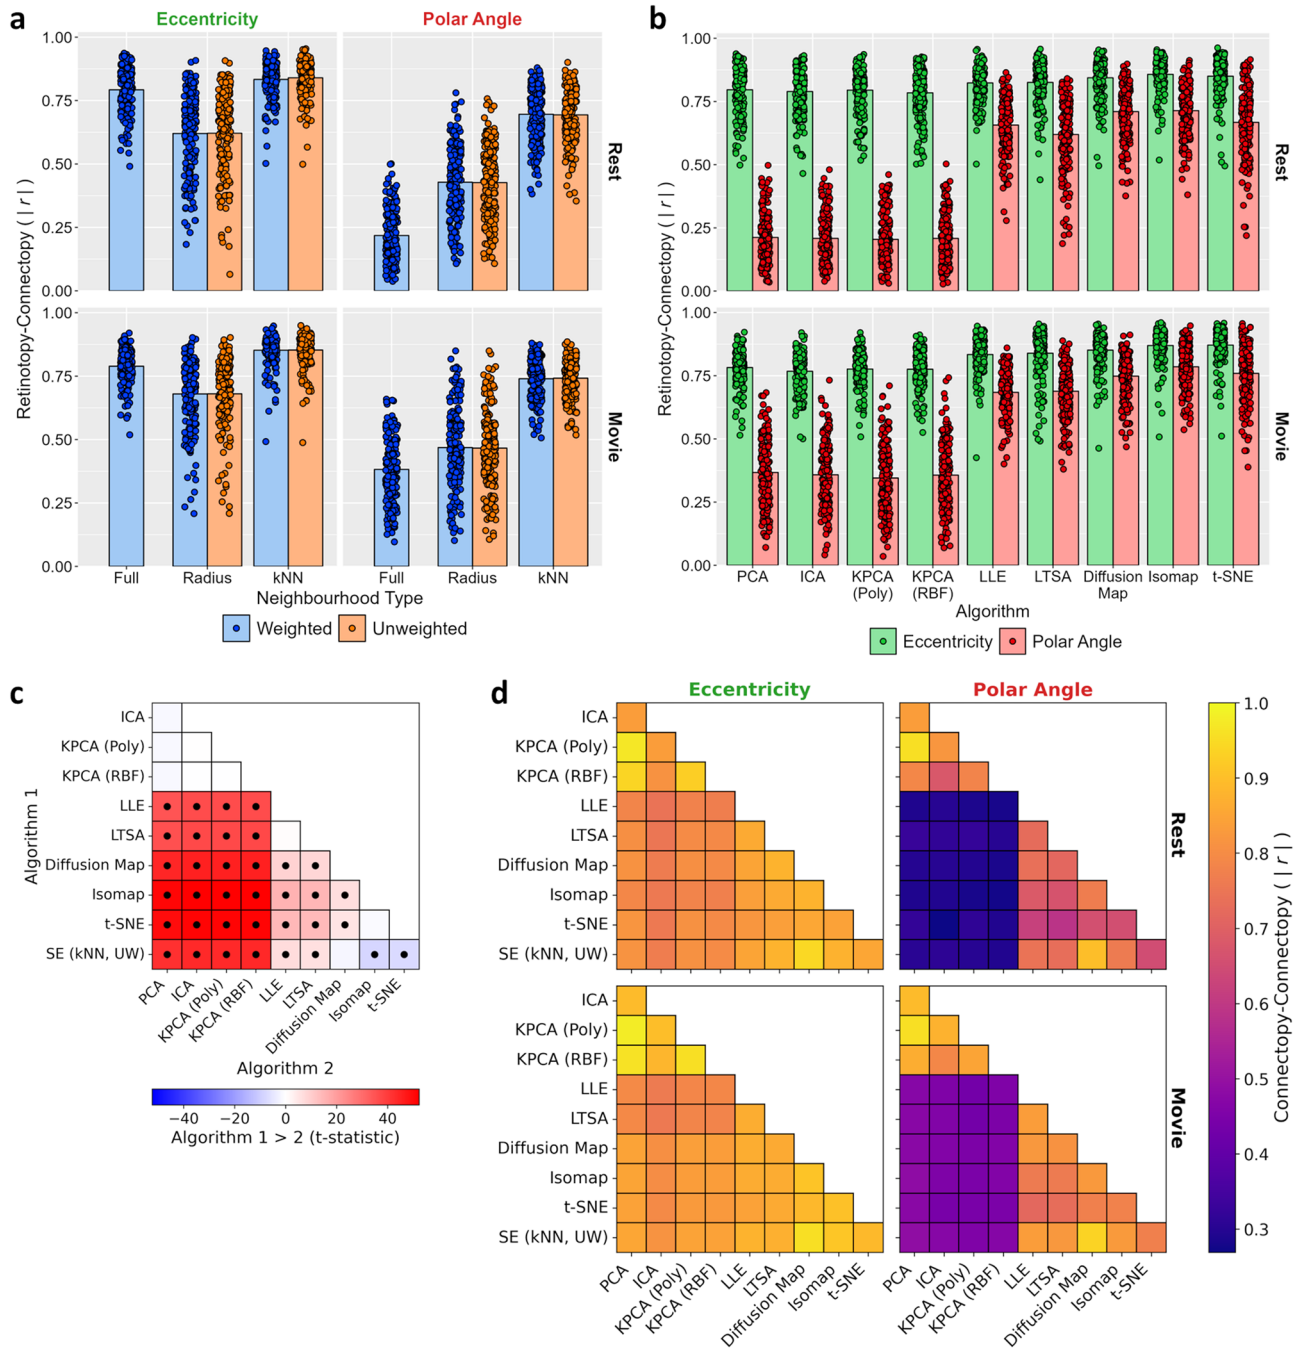
**Figure 5.** Experiment 2: Retinotopic and connectopic polar angle maps in the left hemisphere of an example subject. (a) Retinotopic maps (top-left plot) and connectopic maps reconstructed from resting-state data (remaining plots). (b) Connectopic maps reconstructed from movie-watching data. Spectral embedding is illustrated for the unweighted nearest neighbour variant.

ANOVA with factors of map type (eccentricity, polar angle), graph type, and task (resting-state, movie-watching). This revealed a significant main effect of map type reflecting overall higher correlations for eccentricity than polar angle maps ( $F(1, 173) = 883.07, p < 0.001, \eta_p^2 = 0.84, \eta_G^2 = 0.52$ ). There was also a significant main effect of graph type ( $F(1.96, 339.30) = 1068.47, p < 0.001, \eta_p^2 = 0.86, \eta_G^2 = 0.50$ ); post-hoc Tukey contrasts revealed that correlations were significantly higher for the nearest neighbour than the radius or fully-connected variants (all  $p < 0.001$ ), and the fully-connected variant also outperformed both the weighted ( $p = 0.049$ ) and unweighted ( $p = 0.035$ ) radius neighbour variants, while there were no significant differences between weighted and unweighted radius neighbourhood or nearest neighbour variants (all  $p > 0.05$ ). The main effect of task was also significant, reflecting higher correlations for movie-watching than resting-state datasets ( $F(1, 173) = 55.89, p < 0.001, \eta_p^2 = 0.24, \eta_G^2 = 0.03$ ). There were significant two-way interactions of map by graph type ( $F(2.14, 369.87) = 441.75, p < 0.001, \eta_p^2 = 0.72, \eta_G^2 = 0.16$ ) and map type by task ( $F(1, 173) = 14.81, p < 0.001, \eta_p^2 = 0.08, \eta_G^2 < 0.01$ ) reflecting greater differences between both graph types and tasks for polar angle than eccentricity maps. The graph type by task interaction was not significant ( $F(1.76, 304.31) = 1.34, p = 0.263, \eta_p^2 = 0.01, \eta_G^2 < 0.01$ ). Finally, there was also a significant three-way map type by graph type by task interaction ( $F(2.02, 349.57) = 38.02, p < 0.001, \eta_p^2 = 0.18, \eta_G^2 = 0.01$ ).

We next considered the other connectopic mapping algorithms (Fig. 6b). As before, prediction accuracies appeared higher for the manifold learning than the linear embedding methods, and higher for movie-watching than resting-state estimates. One-sample t-tests revealed that correlations with the retinotopic maps were significantly greater than zero in all cases (one-sample t-tests; all  $p < 0.001$ ). We further entered the correlations for each algorithm into a series of two-way repeated-measures ANOVAs with factors of map type and task – the results of these analyses are listed in Table 1. All algorithms showed a significant main effect of map type due to higher correlations for eccentricity than polar angle maps, as well as a significant main effect of task due to higher correlations for movie-watching than resting-state datasets. Additionally, all algorithms except LLE and LTSA showed a significant map type by task interaction.

To compare prediction accuracies between algorithms, we further entered the correlations into a three-way repeated measures ANOVA with factors of map type, task, and algorithm type. For spectral embedding, we selected the unweighted nearest neighbour approach as one of the better performing variants. There were significant main effects of map type ( $F(1, 173) = 1614.28, p < 0.001, \eta_p^2 = 0.90, \eta_G^2 = 0.65$ ) reflecting higher correlations





**Figure 6.** Experiment 2 main results. **(a)** Prediction accuracies for all variants of spectral embedding, measured by absolute correlations between retinotopic and connectopic maps. Dot markers indicate per-subject means, and bars indicate group means. Resting-state and movie-watching results are illustrated on top and bottom rows respectively. **(b)** Prediction accuracies for other algorithms. **(c)** Post-hoc Tukey contrasts of prediction accuracies between algorithms. For spectral embedding, the unweighted nearest neighbour variant is selected. Dot markers indicate significant contrasts ( $p < 0.05$ ). **(d)** Correlations of connectopic maps between algorithms.

for eccentricity than polar angle maps, of task ( $F(1, 173) = 87.33, p < 0.001, \eta_p^2 = 0.34, \eta_G^2 = 0.03$ ) reflecting higher correlations for movie-watching than resting-state, and of algorithm type ( $F(3.40, 587.99) = 1084.96, p < 0.001, \eta_p^2 = 0.86, \eta_G^2 = 0.55$ ). Post-hoc Tukey contrasts (Fig. 6c) revealed that Isomap and t-SNE outperformed spectral embedding and diffusion maps, which in turn outperformed LLE and LTSA, which in turn outperformed PCA, ICA, and the kernel PCAs (all  $p < 0.001$ ). There were no significant differences between Isomap and t-SNE, or between spectral embedding and diffusion maps, or between LLE and LTSA, or between PCA, ICA, and the kernel PCAs (all  $p > 0.05$ ). There were significant two-way interactions of map type by task ( $F(1, 173) = 111.59, p < 0.001, \eta_p^2 = 0.39, \eta_G^2 = 0.03$ ) and of map by algorithm type ( $F(3.95, 682.74) = 502.74, p < 0.001, \eta_p^2 = 0.74,$



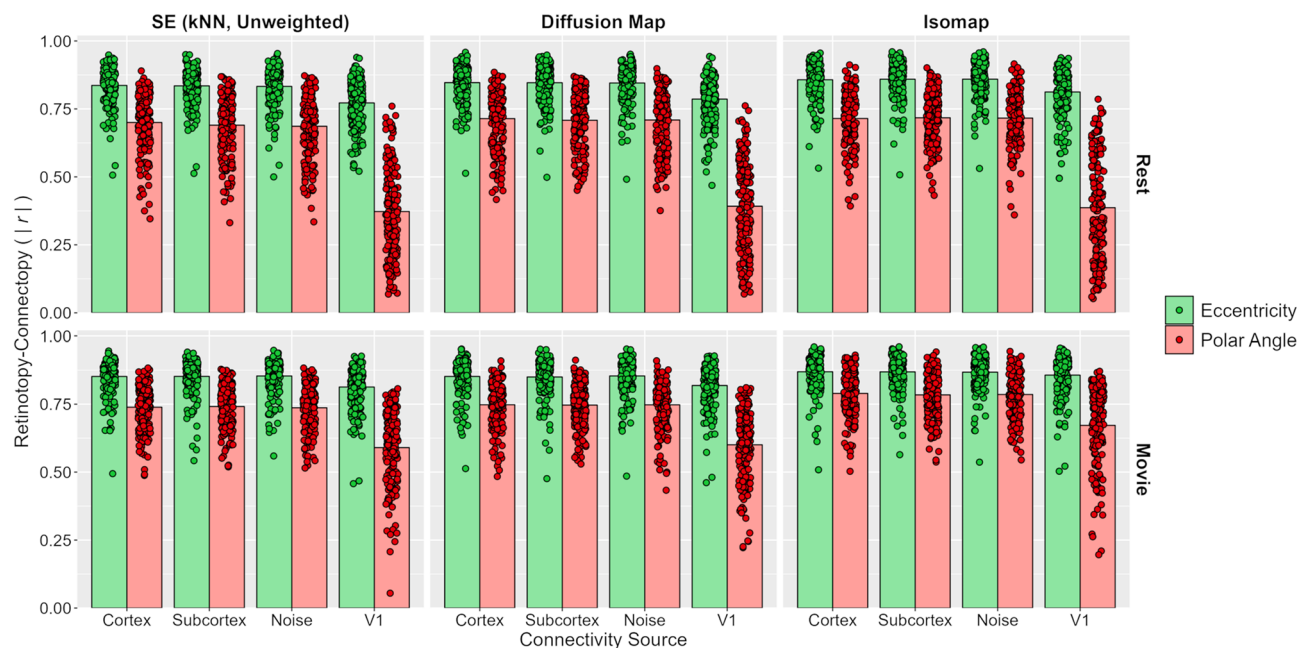
Algorithm	Factor	F	DoF	$p$	$\eta_p^2$	$\eta_G^2$
PCA	Map type	2381.73	1, 173	<0.001	0.93	0.84
	Task	27.91	1, 173	<0.001	0.14	0.02
	Interaction	129.24	1, 173	<0.001	0.43	0.11
ICA	Map type	2522.95	1, 173	<0.001	0.94	0.84
	Task	16.13	1, 173	<0.001	0.09	0.01
	Interaction	133.63	1, 173	<0.001	0.44	0.13
KPCA (Poly)	Map type	2553.89	1, 173	<0.001	0.94	0.84
	Task	16.48	1, 173	<0.001	0.09	0.01
	Interaction	113.19	1, 173	<0.001	0.40	0.11
KPCA (RBF)	Map type	2516.78	1, 173	<0.001	0.94	0.83
	Task	32.93	1, 173	<0.001	0.16	0.03
	Interaction	97.49	1, 173	<0.001	0.36	0.10
LLE	Map type	686.93	1, 173	<0.001	0.80	0.53
	Task	6.51	1, 173	<b>0.012</b>	0.04	0.01
	Interaction	0.07	1, 173	0.797	<0.01	<0.01
LTSA	Map type	544.03	1, 173	<0.001	0.76	0.51
	Task	36.86	1, 173	<0.001	0.18	0.03
	Interaction	3.39	1, 173	0.067	0.02	<0.01
Diffusion Map	Map type	451.98	1, 173	<0.001	0.72	0.41
	Task	12.56	1, 173	<0.001	0.07	0.01
	Interaction	9.63	1, 173	<b>0.002</b>	0.05	<0.01
Isomap	Map type	424.30	1, 173	<0.001	0.71	0.38
	Task	50.07	1, 173	<0.001	0.22	0.05
	Interaction	64.41	1, 173	<0.001	0.27	0.02
t-SNE	Map type	373.70	1, 173	<0.001	0.68	0.35
	Task	57.68	1, 173	<0.001	0.25	0.07
	Interaction	27.38	1, 173	<0.001	0.14	0.02

**Table 1.** Experiment 2: Two-way repeated-measures ANOVAs of prediction accuracies with factors for map type (eccentricity, polar angle) and task (rest, movie-watching) for all connectopic mapping algorithms except spectral embedding. Significant  $p$ -values are highlighted in bold.

$\eta_G^2 = 0.23$ ) because the difference between both tasks and algorithms was greater for polar angle than eccentricity maps, and a significant task by algorithm type interaction ( $F(5.23, 905.24) = 8.84, p < 0.001, \eta_p^2 = 0.05, \eta_G^2 = 0.01$ ) because the difference between tasks was reduced for better performing algorithms. Finally, there was a significant three-way map type by task by algorithm type interaction ( $F(4.30, 744.17) = 39.58, p < 0.001, \eta_p^2 = 0.19, \eta_G^2 = 0.02$ ).

We also tested the consistency of the connectopic maps between algorithms by correlating the maps (Fig. 6d). As before, the linear embeddings and kernel PCAs produced relatively similar gradients, and the remaining manifold learning methods also yielded similar gradients, but gradients appeared less similar between these two groups. Additionally, gradients appeared more consistent when based on movie-watching than on resting-state data. In summary, the manifold learning algorithms outperformed the linear embeddings and both kernel PCAs (especially when predicting the polar angle maps), with Isomap and t-SNE performing best, and performance was better when connectivity was estimated from movie-watching than resting-state data.

We next tested how the source of the connectivity information affects the connectopic mapping. We recalculated the connectivity fingerprints by correlating the V1 timeseries with the timeseries from either just the cortical vertices, just the subcortical voxels, the cortical vertices and subcortical voxels substituted with Gaussian noise, or the same V1 timeseries. We repeated the connectopic mapping with the spectral embedding (unweighted nearest neighbour), diffusion map, and Isomap algorithms using these new fingerprints (Fig. 7). Similar to Experiment 1, the pattern of results appeared largely similar across the different connectivity sources. As before, all prediction accuracies remained high (one-sample  $t$ -tests: all  $p < 0.001$ ). We entered the correlations into a four-way repeated-measures ANOVA with factors of map type, task, algorithm type, and connectivity source (cortex + subcortex, cortex, subcortex, noise, and within-V1). This revealed significant main effects of map type ( $F(1, 173) = 745.28, p < 0.001, \eta_p^2 = 0.81, \eta_G^2 = 0.47$ ), task ( $F(1, 173) = 132.43, p < 0.001, \eta_p^2 = 0.43, \eta_G^2 = 0.05$ ), and algorithm type ( $F(1.42, 245.83) = 190.49, p < 0.001, \eta_p^2 = 0.52, \eta_G^2 = 0.04$ ); post-hoc Tukey contrasts indicated correlations were significantly higher for Isomap than diffusion maps, and both were higher than spectral embedding (all  $p < 0.001$ ). Unlike Experiment 1, the main effect of connectivity source was now also significant ( $F(1.22, 211.02) = 597.45, p < 0.001, \eta_p^2 = 0.78, \eta_G^2 = 0.20$ ); post-hoc Tukey contrasts indicated this was due to significantly lower correlations for the within-V1 analyses than the other sources (all  $p < 0.001$ ), while the remaining sources did not differ significantly (all  $p > 0.05$ ). Note that this result may reflect differences in the pre-processing pipeline for the within-V1 analyses compared to other sources (see [Unregistered Analyses](#)). Finally, all interactions were significant (all  $p < 0.001$ ).



**Figure 7.** Experiment 2: Prediction accuracies for unweighted nearest neighbour variant of spectral embedding, diffusion map, and Isomap algorithms when varying the source of the connectivity fingerprints. V1 timeseries are correlated with timeseries taken from just non-V1 cortical grayordinates, just subcortical grayordinates, non-V1 cortical plus subcortical grayordinates substituted with Gaussian noise, or the same V1 timeseries. Resting-state and movie-watching results are illustrated on top and bottom rows respectively.

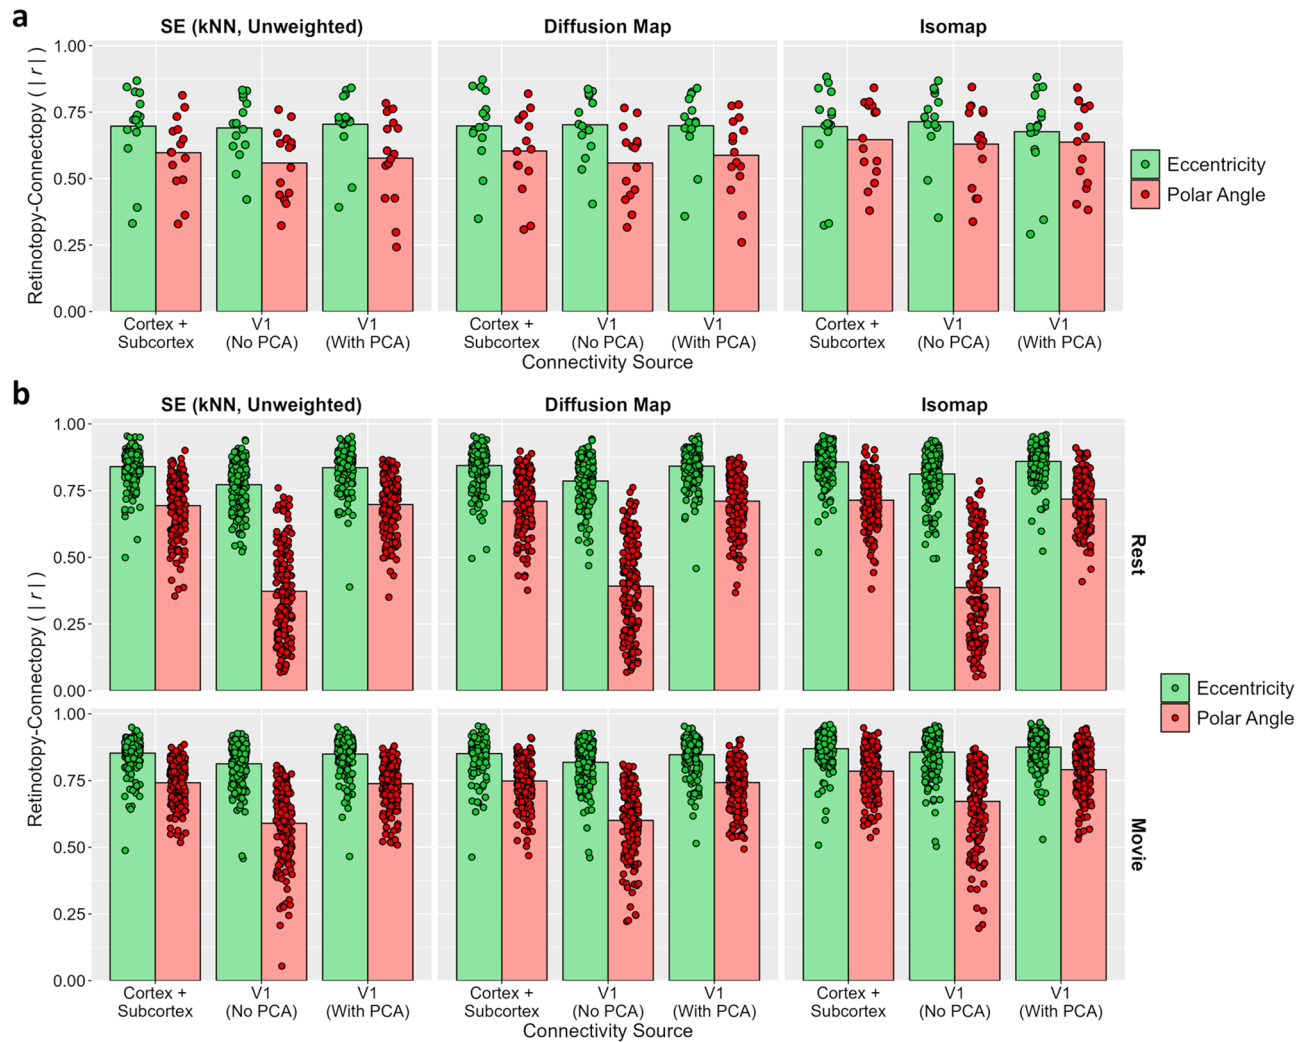
except for the two-way map type by task interaction ( $p=0.443$ ) and the four-way map type by task by algorithm type by connectivity source interaction ( $p=0.392$ ).

Finally, we conducted two further analyses to examine the stability of the connectopic mapping algorithms. Firstly, we investigated the split-half reliability by correlating the connectopic maps between the two cross-validation splits. The correlations appeared high in all cases (Supplementary Fig. S2). Secondly, we correlated the connectopic maps between the cross-validation splits *and* tasks (e.g. correlating maps for odd resting-state runs with even movie-watching runs). Again, the correlations appeared high in all cases (Supplementary Fig. S3). This indicates good internal reliability within each algorithm, that the connectopic maps are partially independent of the precise content of the movie stimulus, and that the maps generalise well between resting-state and movie-watching tasks.

**Unregistered analyses.** In our preregistered analyses, we observed an apparent reduction in prediction accuracies for the within-V1 analyses compared to those of other connectivity sources. One key difference between the within-V1 analyses and those of the other sources is that we omit the PCA transformation of the timeseries prior to calculating the connectivity fingerprints (Fig. 1). This is because the number of timepoints already exceeds the number of vertices within the ROI and hence the dimensionality cannot be reduced further without losing information. Nevertheless, it is possible that the reduction in prediction accuracy for the within-V1 analyses relative to other sources reflects the omission of this processing stage rather than the connectivity information represented within the region. To test this possibility, we conducted exploratory analyses repeating the within-V1 analyses but including the PCA transformation stage, such that the V1 timeseries were now correlated with principal components derived from those same timeseries. The results of these analyses are shown in Fig. 8. While the inclusion of the PCA transformation had minimal effect on the StudyForrest data (Experiment 1), including it for the HCP data (Experiment 2) recovered performance to levels comparable with the main analyses. This suggests the apparent reduction in performance for the within-V1 analyses mostly reflects methodological differences in the processing pipeline compared to the other neural sources. In general, as per Experiment 1, varying the connectivity source had minimal effect on the connectopic mapping, suggesting that the functional topography within the region is the key factor.

## Discussion

This study aimed to determine how connectopic mapping predicts functional gradients in primary visual cortex by benchmarking reconstructions against ground-truth estimates of retinotopic maps measured by visual field mapping. We reported results from an exploratory analysis using data from natural viewing (movie-watching), followed by a preregistered replication experiment using both resting-state and natural viewing data. In both experiments, connectopic mapping accurately reconstructed retinotopic maps across subjects. More advanced manifold learning methods outperformed simpler methods including linear embeddings (especially for polar angle maps), with Isomap and t-SNE performing best. Additionally, prediction accuracy was better for



**Figure 8.** Exploratory analyses testing effect of applying lossless PCA compression to timeseries prior to calculating connectivity fingerprints for within-V1 analyses. Plots illustrate prediction accuracies, measured by absolute correlations between retinotopic and connectopic maps, for the unweighted nearest neighbour variant of spectral embedding, diffusion map, and Isomap algorithms. Results are illustrated for (a) Experiment 1 and (b) Experiment 2: resting-state and movie-watching results are illustrated on top and bottom rows respectively. Left and middle bar groups duplicate results from Figs. 3, 6, and 7: V1 timeseries are correlated with timeseries taken from non-V1 cortical vertices and subcortical voxels (following PCA compression) or the same V1 timeseries (without PCA compression). Right bar groups illustrate accuracies when V1 timeseries are correlated with principal components derived from the same V1 timeseries.

eccentricity than polar angle maps, and better when using connectivity estimated during natural viewing than during rest. Varying the neural source of the connectivity estimates had minimal impact on the connectopic mapping, suggesting the key factor is the functional topography within the region of interest.

Previously, Haak and colleagues demonstrated that connectopic mapping (via spectral embedding) using resting-state data reconstructed gradients resembling retinotopic maps in primary visual cortex<sup>22</sup>. We replicated these findings, and further compared the connectopic maps against ground-truth estimates of the retinotopic maps measured by traditional visual field mapping techniques. We confirmed that connectopic maps reconstructed from both resting-state and movie-watching data accurately predicted retinotopic maps across subjects. The significance of these findings is that they provide validated methods for exploring gradients in other brain regions, and indeed connectopic mapping techniques have been applied to other regions including motor cortex<sup>22</sup>, the striatum<sup>26</sup>, and the hippocampus<sup>27</sup>. Such approaches require that a region of interest is first defined to perform the connectopic mapping within. This could, for instance, be done functionally (e.g. via visual field mapping or a functional localiser), or by using an atlas- or parcellation-based approach. Connectopic mapping has also been used to recover coarser scale gradients over the whole brain<sup>16,25</sup>, although ascertaining the accuracy of such gradients may be more challenging as the ground-truth is often unknown. We obtained remarkably consistent results between both experiments, suggesting that connectopic mapping is robust to differences in stimuli and tasks (e.g. continuous movie sequences in Experiment 1, compared to rest or compilations of short



movie clips in Experiment 2), data acquisition parameters (e.g. 3T vs. 7T, single- vs. multi-band, etc.), and pre-processing pipelines. This approach offers the possibility of deriving functional gradients throughout the brain, including in regions where the organising principles are currently poorly understood or unknown. In regions where functional gradients have previously been described, they are typically observed to be topographically organised such that there is a gradual change in response properties which is mapped smoothly across a brain region. Using connectopic mapping, it will therefore be possible to determine if topographic maps are a ubiquitous organising principle throughout the brain<sup>1,15</sup>.

Previous studies employing connectopic mapping have exclusively measured connectivity in the brain at rest. Although resting-state is a common choice, it has been criticised as an unnatural cognitive state<sup>30</sup>. Here, we compared the performance of connectopic mapping when estimating connectivity at rest or during natural viewing (movie-watching). Not only did connectopic mapping perform well using movie-watching data, but prediction accuracies exceeded those obtained from resting-state. Previous studies have indicated movie-watching elicits functional connectivity that is less widespread and reduced in magnitude compared to resting-state<sup>38–40</sup>, likely due to stimulus driven responses disrupting spontaneous or intrinsic brain activity. However, connectivity from movie-watching may be more predictive of function in some brain regions including visual cortex<sup>32,41</sup>, and our results are consistent with this conclusion. Movie-watching may also offer other advantages such as improved reliability of connectivity estimates<sup>39,42</sup>, reproducibility of brain activity across participants and data acquisition sites<sup>43</sup>, and suitability for studying developmental populations<sup>44,45</sup>. Nevertheless, the gradients could still be accurately reconstructed from connectivity at rest, and there may be scenarios where resting-state data would be advantageous, for instance if studying visually impaired populations. Furthermore, resting-state data may be better predictive of function in non-sensory regions such as within the default mode network<sup>32</sup>. The spatial topography of connectivity is typically similar between rest and movie-watching<sup>38,40</sup> (but see<sup>46</sup>), and resting-state connectivity is predictive of brain activity during movie-watching<sup>47</sup>, thus both approaches are likely to reveal similar aspects of functional connectivity.

One key question is the extent to which the neural source of the connectivity signals affects connectopic mapping within the region of interest. To test this, we repeated our analyses while varying the source of the connectivity signals. We first measured connectivity between vertices in V1 and either just the cortical vertices or just the subcortical voxels (as opposed to the combination of both in the main analyses). Surprisingly, this had minimal effect on the connectopic maps: prediction accuracies remained comparable to those of the main analyses. This suggests that the main driver of connectopic mapping is the functional topography within the region itself. For instance, two functionally similar voxels within the region displaying relatively similar timecourses of activation will produce relatively similar patterns of connectivity with the rest of the brain. Conversely, two functionally dissimilar voxels within the region will yield dissimilar connectivity fingerprints. Consequently, the functional topography within the region alone may prove sufficient to drive connectopic mapping. Indeed, we also found that performance was minimally affected when measuring connectivity purely within V1 or when substituting the non-V1 timeseries with random noise. A slight decrement in performance was observed for the within-V1 analyses in Experiment 2, however exploratory analyses indicated this likely resulted from omitting the lossless PCA transformation of the source timeseries prior to calculating the connectivity fingerprints. The PCA transformation generally reduces the magnitude of correlations in the connectivity fingerprints and biases the fingerprints to have stronger correlations along earlier than later dimensions, and potentially these properties may prove beneficial to the manifold learning. Regardless, our analyses tentatively suggest that a lossless PCA transformation of the source timeseries may be a useful standard processing step to include in connectopic mapping pipelines.

Previous studies have interpreted connectopic mapping as reflecting a topographic organisation of the connectivity itself<sup>21,22</sup>. In contrast, our results suggest a different interpretation in which the key factor is the functional topography within the target region, and the connectivity provides an index into this topography but the source of that connectivity is relatively unimportant. This is not to say that connectivity is not topographically organised, but rather that such organisation is not necessary for connectopic mapping to accurately predict functional gradients. Further investigations will be required to determine if this principle holds for other brain regions – for instance, whether the source of the connectivity would be more important for mapping higher-level and/or non-sensory regions, or whether the within-region functional topography would remain sufficient.

A key aim of this study was to compare the efficacy of different dimensionality reduction techniques in performing connectopic mapping. We tested a number of commonly used and widely implemented algorithms including both non-linear manifold learning and linear embedding methods. We found that Isomap and t-SNE achieved the best prediction accuracies; there was no significant difference between these algorithms, although Isomap may be preferable as t-SNE is considerably more computationally expensive. Spectral embedding and diffusion maps (both of which have previously been applied to connectopic mapping<sup>16,22</sup>) performed next best and only slightly behind Isomap and t-SNE. Within spectral embedding, we observed best performance using a nearest-neighbourhood graph (as opposed to a radius-neighbourhood or fully-connected graph) and found little benefit to weighting the graph. Locally linear embedding and local tangent space alignment performed next best, and the linear embeddings (PCA and ICA) and kernel PCAs (with polynomial and radial basis function kernels) performed the worst. Thus, more advanced manifold learning methods outperformed the simpler linear embeddings suggesting some degree of non-linearity to the connectivity manifold—particularly in the case of polar angle maps.

Contrary to our findings, Haak and colleagues found relatively poor performance of the Isomap algorithm<sup>22</sup>. One possible reason for this discrepancy may be the selection of the neighbourhood parameter for the algorithm: we optimised this parameter to maximise the prediction accuracy, whereas Haak and colleagues set it to the minimum value that permitted a connected graph. Thus, appropriate parameter selection may be important to the success of connectopic mapping in deriving accurate gradients. Most manifold learning algorithms require some

form of parameter selection. In our case the parameters were optimised based on the ground-truth estimates provided by the visual field mapping. However, if the ground-truth is unknown then appropriate parameter selection may be more challenging. Nevertheless, some heuristics may aid parameter selection. For example, one proposal for Isomap is to set the parameter to maximise the correlation between geodesic distances in the original space and Euclidean distances in the embedded space<sup>48</sup>. However, such heuristics may not be available or appropriate in all cases. If no better options are available, it may be preferable to simply observe how robust the connectopic maps are to changes in the parameter(s) over a range of values. An additional issue is that the connectopic maps may not emerge consistently or discretely along each component. For instance, the maps may emerge across different components over subjects, or could load on more than one component. One solution may be to apply a further linear transformation (e.g. Procrustes) after the manifold learning to optimise the alignment between the connectopic maps and some other reference maps (such as ground-truth estimates, or connectopic maps derived from a different participant or at the group level)<sup>49</sup>. We did not apply such an approach here as further transformation of the maps could have obscured differences between the dimensionality reduction algorithms.

In conclusion, we provide an appraisal of connectopic mapping methods for reconstructing functional gradients in primary visual cortex. We find that connectopic mapping can accurately predict retinotopic maps across subjects, with better performance for eccentricity than polar angle maps. We find that non-linear manifold learning algorithms outperform linear dimensionality reduction techniques, with Isomap and t-SNE performing especially well. Performance was best when using connectivity estimated during natural viewing, although prediction accuracies remained high using resting-state estimates too. The neural source of the connectivity estimates had minimal effect on the performance of the connectopic mapping suggesting that the key factor is the functional topography within the region of interest rather than a topography embedded within the connectivity itself. Developing standardised methods for performing connectopic mapping opens the possibility of applying this technique to explore topographic maps throughout the brain.

## Methods

**Experiment 1.** Our first experiment presents exploratory analyses of movie-watching data obtained from a publicly available MRI dataset.

*Dataset.* We obtained movie-watching 3T MRI data from the publicly available StudyForrest dataset<sup>34,36</sup> (<https://www.studyforrest.org/>). We used a subset comprising 15 participants who also have visual field mapping data<sup>37</sup>: S1–S6, S9, S10, and S14–S20. In brief, functional data were acquired on a 3 T Philips Achieva MRI scanner via an EPI sequence (TR=2 s, TE=30 ms, voxel resolution=3 mm isotropic). The movie-watching stimulus comprised approximately 2 h of the “Forrest Gump” movie, and the visual field mapping stimulus comprised a flickering chequerboard presented within an aperture displaying expanding/contracting rings or rotating wedges. Full details are provided in<sup>37</sup>.

*Pre-processing.* Movie-watching and visual field mapping data were subjected to the same pre-processing pipeline. Some light pre-processing had already been applied by the StudyForrest project: this included motion correction using FSL’s MCFLIRT tool<sup>50</sup>, and aligning each volume to a common subject-specific reference volume shared across all runs. We then applied additional pre-processing using FSL’s FEAT v6.0<sup>51,52</sup> (<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/>): slice-timing correction using Fourier-space time-series phase-shifting, non-brain removal<sup>53</sup>, 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 applied at this stage. The functional timeseries were then normalised by first converting to units of percentage signal change, and then finally regressing out both the mean ventricular and white-matter timeseries and motion parameters.

The pre-processed timeseries were co-registered to anatomical spaces. Cortical surfaces in each subject were reconstructed from T1- and T2-weighted anatomical images using Freesurfer v6.0<sup>54</sup> (<https://surfer.nmr.mgh.harvard.edu/>). Cortical data were co-registered to each subject’s native surface via boundary based registration<sup>55</sup>, then further transformed to the fsaverage6 surface via a surface-based registration<sup>56,57</sup>. Surface-based spatial smoothing was then applied at FWHM = 6 mm (twice the voxel resolution). For the movie-watching runs only, volumetric data were also co-registered to each subject’s anatomical T1 image via boundary based registration and then onto the MNI152 standard brain via FSL’s FNIRT tool<sup>58</sup>. Volume-based spatial smoothing was then applied at FWHM = 6 mm. The volumetric data were finally restricted to a subcortical grey-matter mask generated from Freesurfer’s Aseg atlas<sup>59</sup> comprising the following labels: cerebellum grey matter, thalamus, caudate, putamen, pallidum, hippocampus, amygdala, accumbens area, and ventral diencephalon.

*Visual field mapping.* We performed a travelling wave analysis<sup>17</sup> of the retinotopy data registered to the cortical surface using the *3dRetinoPhase*<sup>60</sup> command in AFNI<sup>61</sup> (<https://afni.nimh.nih.gov/>). This yielded phase maps representing the eccentricity and polar angle tunings of each surface vertex. The polar angle map was used to define individualised V1 regions of interest (ROIs) by tracing along the phase reversals. The eccentricity and polar angle maps restricted to the V1 ROIs also provided a ground-truth estimate of the retinotopic maps against which the connectopic maps could be benchmarked.

**Experiment 2.** Our second experiment presents preregistered analyses of a larger dataset comprising both movie-watching and resting-state data obtained from the Human Connectome Project.

**Dataset.** We obtained visual field mapping, resting-state, and movie-watching 7T MRI data from the Human Connectome Project<sup>35</sup>. We used a subset of 176 subjects from the S1200 release who have fully completed all three tasks. Two subjects (126,931 and 745,555) were removed as MSMAll-aligned versions of their datasets were unavailable, leaving a total of 174 subjects; note that this represents a deviation from our preregistration. A list of the subject IDs is provided in Supplementary Table S1. In brief, functional data were acquired on a 7T Siemens Magnetom MRI scanner via a multiband EPI sequence (TR = 1 s, TE = 22.2 ms, voxel resolution = 1.6 mm isotropic). The visual field mapping stimuli comprised a flickering “mashfast” texture of objects superimposed on an amplitude mask background<sup>62</sup> presented within apertures displaying rotating wedges, expanding/contracting rings, or drifting bars. Resting-state data were acquired in four scan runs each approximately 16 min in duration. The movie-watching stimuli comprised approximately 1 h of short clips, taken from independent films and Hollywood movies<sup>63</sup>, presented across four scan runs. See the WU-Minn HCP S1200 Data Release reference manual for full details.

We performed a series of power analyses using G\*Power v3.1<sup>64</sup> to ensure the sample size would be appropriate. A basic test of our hypotheses is whether the connectopic mapping can accurately predict the retinotopic maps. This can be tested by measuring the prediction accuracies as the absolute correlation between retinotopic and connectopic maps and using a one-sample t-test to compare the values over subjects against zero (see “Methods”—“Statistical Analyses” section below). The smallest effect size obtained from these tests within our exploratory analyses (see “Results”—“Experiment 1”) was Cohen’s  $d_z = 1.76$  (for the prediction of polar angle maps by the RBF kernel PCA). Entering this effect size into G\*Power, with a sample size of 174, and an alpha criterion of 0.05, indicated a power of 100% (to within numerical precision) would be achieved. A further critical test is the ability to discriminate prediction accuracies between the different connectopic mapping algorithms. To test this in our exploratory analyses we conducted a two-way repeated-measures ANOVA on the retinotopy correlations with factors for the map and algorithm type. This revealed an effect size of  $\eta_p^2 = 0.64$  (Cohen’s  $f = 1.33$ ) for the main effect of algorithm type. We entered this effect size into G\*Power, listing the effect size specification “as in SPSS”, with 10 measurements (factor levels—the number of algorithms), an alpha criterion of 0.05, and a sample size of 174. This again indicated a power of 100% (to within numerical precision) would be achieved.

**Pre-processing.** Visual field mapping, resting-state, and movie-watching datasets all followed the HCP minimal pre-processing pipeline including FIX-denoising. In brief, this includes gradient distortion correction, motion correction, high-pass temporal filtering with a cutoff of FWHM = 2355 s, and ICA denoising via FSL’s MELODIC tool. No slice timing correction was applied due to the relatively fast TR. Cortical data were registered to the HCP’s fsLR32k standard surface<sup>65</sup>, while subcortical data were registered to the MNI brain. Surface registration was performed via a multimodal alignment procedure (MSMAll) which aligns surfaces based on cortical folding plus additional areal features derived from myelin maps, resting-state network maps, and resting-state visuo-topic maps<sup>66,67</sup>. Full details of the HCP pre-processing can be found in<sup>68,69</sup>. We then applied the following further pre-processing steps: timeseries were converted to percentage signal change, and the data were spatially smoothed at FWHM = 3.2 mm (twice the voxel resolution; surface-based for cortical grayordinates and volume-based for subcortical grayordinates). Unlike with the StudyForrest dataset, we did not regress out the mean white-matter and ventricular timeseries as these structures are not included within the grayordinates.

**Visual field mapping.** We fit a population receptive field model to the retinotopy data using the MATLAB analyzePRF toolbox (<https://github.com/cvnlab/analyzePRF>) and code adapted from<sup>70</sup>. Phase maps were extracted representing the eccentricity and polar angle tunings of each surface vertex, and the polar angle maps were used to define individualised V1 ROIs by tracing along the phase reversals. As before, the eccentricity and polar angle maps also served as ground-truth estimates to benchmark the connectopic maps against.

**Deviations from preregistration.** We note the following deviations from our pre-registered protocol:

- We intended to use all 176 subjects who fully completed the resting-state, movie-watching, and visual field mapping scans. However, we removed two subjects (126,931, 745,555) because MSMAll-aligned versions of their datasets were unavailable, leaving a total of 174 subjects.
- We include additional exploratory analyses of the within-V1 connectopic mapping analyses (Fig. 8; see *Unregistered Analyses*). These explore the effect of including or omitting a PCA transformation of the timeseries prior to calculating the connectivity fingerprints.

**Estimating functional connectivity.** Further connectivity analyses followed the same procedures for both experiments. A schematic illustration of the connectivity analysis pipeline is shown in Fig. 1. Movie-watching data (and resting-state data for the HCP dataset) were partitioned into odd and even scan runs to allow cross-validated parameter selection for the dimensionality reduction algorithms (where applicable). Within each data split, the normalised and spatially smoothed timeseries were concatenated over scan runs. Connectivity fingerprints were estimated for each data split following the methods of<sup>22</sup>. The timeseries were split between surface vertices within the V1 ROI versus all surface vertices and subcortical voxels outside the ROI. For the non-V1 timeseries, the number of dimensions (surface vertices plus subcortical voxels) exceeds the number of samples (timepoints). We therefore reduced the dimensionality via a lossless PCA, retaining all available components (one fewer than the number of timepoints) to explain 100% of the variance. This operation amounts to rotating the samples within the feature space and removing the unused dimensions—this aids the computational tractability of later processing stages, but does not incur any loss of information from the non-V1 data. The timeseries were correlated between the V1 vertices and non-V1 principal components, and the correlations



were Fisher transformed. This yielded a set of connectivity fingerprints describing the pattern of functional connectivity between each V1 vertex and the rest of the brain. Within this connectivity space V1 vertices are represented as samples and non-V1 principal components as dimensions. The dimensionality of this space was then reduced to extract the connectopic maps (see below).

We also generated further connectivity estimates to test the influence of the source of connectivity signals. In the first two variants, we derived non-V1 timeseries from only the cortical surface vertices or from only the subcortical voxels. In a further two variants, we tested if the functional topography within V1 alone is sufficient to drive the connectopic mapping. Firstly, we repeated the main analyses with non-V1 timeseries derived from surface vertices and subcortical voxels together, but we replaced the timeseries with normally distributed random noise matched in mean and variance to the real timeseries. These noise timeseries will reproduce any global amplitude differences over vertices/voxels but will not contain any consistent temporal variation. Secondly, we conducted a purely within-ROI analysis by correlating timeseries between pairwise combinations of V1 vertices, disregarding the non-V1 timeseries entirely. Note that the lossless PCA compression stage was omitted for this analysis as the number of vertices within V1 is already less than the number of timepoints. Because each V1 vertex is trivially connected to itself (and the resulting perfect correlation would invalidate the Fisher transformation), the diagonal elements of this connectivity matrix were instead set to zero. For all variants, we then repeated our connectopic mapping with the spectral embedding (nearest-neighbour, unweighted), diffusion map, and Isomap algorithms.

**Connectopic mapping.** The connectivity fingerprints were reduced in dimensionality to extract the connectopic maps. We always retained the first two components, which are expected to correspond to the eccentricity and polar angle retinotopic maps. Prediction accuracy was assessed by taking the absolute correlation between the retinotopic and connectopic maps—the absolute value was used to account for the sign ambiguity inherent in the connectopic maps. The order of the components is not guaranteed to be consistent over different analyses, so we re-ordered the first two components according to their correlations with the retinotopic maps. Specifically, we calculated the  $2 \times 2$  pairwise correlations between both retinotopic and connectopic maps, then matched the retinotopic/connectopic map pair with the largest absolute correlation, then further matched the remaining map pair. Note that if both connectopic maps are maximally correlated with the same retinotopic map (or vice versa) then the first match will be based on whichever pairing yields the highest correlation—this ensures that only one connectopic map is matched to each retinotopic map.

We employed a number of dimensionality reduction algorithms, including both linear embeddings and non-linear manifold learning techniques. We provide more detail on each of these algorithms in the sections below. Although not an exhaustive list, we have included some of the most commonly used and widely implemented manifold learning algorithms, including those previously used for connectopic mapping such as spectral embedding<sup>22</sup> and diffusion maps<sup>16</sup>. In the case of spectral embedding, we also considered a number of variants of the algorithm that have been variously employed in the literature<sup>22,24,49</sup>.

All algorithms were implemented using the *scikit-learn* Python module<sup>71</sup> (<https://scikit-learn.org/stable/index.html>), except for the diffusion maps algorithm which was implemented using the *pydiffmap* Python module (<https://github.com/DiffusionMapsAcademics/pyDiffMap>). Where necessary, parameter selection was performed via cross-validation: a Bayesian optimisation algorithm, implemented using the *scikit-optimize* Python module (<https://scikit-optimize.github.io/stable/>), selected the parameter value(s) that maximised the absolute correlation between retinotopic and connectopic maps (averaged over maps) within one of the data splits. The selected parameters were then applied to connectopic mapping of the other data split. The Bayesian optimisation performed 100 iterations (including 10 random starts) per parameter to be optimised.

**Spectral embedding.** Spectral embedding, also referred to as Laplacian eigenmaps, is a manifold learning technique that finds a lower dimensional embedding preserving local neighbourhood relationships from the original feature space<sup>72</sup>. This technique has previously been used for connectopic mapping in a number of sensory brain regions<sup>22,24,26,28</sup>, including primary visual cortex, and we provide particular focus to it here. First, a neighbourhood graph was constructed denoting which samples fall within the local neighbourhoods of each other sample. This can be performed using a radius approach, in which each neighbourhood is defined by a fixed distance from each sample (denoted by the parameter  $\epsilon$ ). However, it may be challenging to find an appropriate radius to apply to all neighbourhoods if samples are not evenly distributed throughout the feature space<sup>73</sup>. Alternatively, a nearest-neighbours approach can be used in which each neighbourhood comprises the  $k$  nearest samples to a given sample—this has the advantage that the neighbourhood extent can vary with the density of the samples. Here, we construct both radius and nearest neighbour graphs based on the squared Euclidean distance between samples in the connectivity space. The neighbourhood parameter ( $\epsilon$  or  $k$ ) was selected via cross-validation. Nearest neighbour graphs were made undirected by considering a pair of samples as being within each other's neighbourhood if either sample is within the  $k$  nearest neighbours of the other. This produces *unweighted* graphs, in which each sample is discretely labelled as being inside or outside a given neighbourhood. The graph may also be made *weighted* by setting the non-zero elements of the unweighted graph to some measure of similarity between the samples. Many metrics are available for weighting, though values must be non-negative. Here, we simply use a Pearson's correlation between samples, rescaled to a zero to one range (i.e., zero represents a perfect negative correlation). We also include a fully-connected version of the weighted graph in which connections are not first restricted to local neighbourhoods and instead all samples are considered connected to all other samples and are only differentiated by the weights<sup>74</sup>. For reference, Haak and colleagues employed weighted radius-neighbourhood graphs<sup>22,28</sup>. Once the neighbourhood graph has been constructed, it is converted to a

graph Laplacian, which is then decomposed into its eigenvectors and eigenvalues. Selecting the eigenvectors corresponding to the smallest  $m$  non-zero eigenvalues yields the connectopic maps (here  $m=2$ ).

**Linear embeddings.** We employed two linear embeddings: principal and independent components analysis (PCA and ICA). Neither of these algorithms required any parameter selection.

**Kernel PCA.** Kernel PCA provides a non-linear extension of linear PCA using the kernel trick<sup>75</sup>. Many kernels are available, but we employed two of the more commonly used ones. Firstly, we used a second-order polynomial kernel:  $k(x, y) = (x^T y + 1)^2$ , where  $x$  and  $y$  are a given pair of feature vectors; this kernel had no free parameters. Secondly, we used a radial basis function (RBF) kernel:  $k(x, y) = \exp(-\gamma \|x - y\|^2)$ , where  $\gamma$  is a free parameter that determines the spread of the kernel.

**Locally linear embedding.** Locally linear embedding (LLE) derives a series of linear decompositions within local neighbourhoods of the feature space<sup>76</sup>. Provided the manifold is approximately linear within local neighbourhoods, the technique may recover a non-linear embedding over the manifold as a whole. We used a modified variant of the algorithm (MLLE) that is proposed to produce solutions which are more stable and less prone to distortions<sup>77</sup>. The algorithm takes a single free parameter ( $k$ ) determining the neighbourhood size.

**Local tangent space alignment.** Local tangent space alignment (LTSA) is closely related to LLE, and follows the intuition that local hyperplanes drawn at a tangent to the manifold should become aligned relative to one another if the manifold is appropriately unfolded<sup>78</sup>. As with LLE, it requires a free parameter ( $k$ ) determining the neighbourhood size.

**Isomap.** Isomap is a form of multidimensional scaling using a geodesic distance metric<sup>79</sup>. This respects the topology of the manifold and may allow more accurate reconstructions than multidimensional scaling using a standard Euclidean distance metric. First, a nearest neighbour graph was constructed (taking a free parameter  $k$ ) and weighted according to the Euclidean distances between samples. Pairwise geodesic distances between all samples were then calculated from the local Euclidean distances via Dijkstra's algorithm<sup>80</sup>. Finally, the geodesic distance matrix was submitted to multidimensional scaling.

**Diffusion maps.** Diffusion maps consider the diffusion distances between samples, defined by the transition probabilities between samples on a random walk over the manifold. This algorithm has previously been applied to connectopic mapping both within local regions and over the whole brain<sup>16,27</sup>. First, a nearest neighbour graph is constructed (taking a free parameter  $k$ ). This graph is then weighted by a Gaussian heat kernel:  $k(x, y) = \exp\left(-\frac{\|x-y\|^2}{\epsilon}\right)$ , where  $x$  and  $y$  are a given pair of samples, and  $\epsilon$  is a free parameter determining the spread of the kernel (not to be confused with the radius parameter of the same name in spectral embedding). From this, a transition matrix can be defined representing the probability of encountering one sample from another on a random walk over the manifold. In this stage, a free parameter ( $\alpha$ ) can be used to control the influence of the sampling density on the estimation, with values typically set between 0 (maximal influence) and 1 (no influence). Finally, diffusion distances between samples may be derived using the eigenvalues and associated eigenvectors of the transition matrix. Taking the first  $m$  eigenvectors (here,  $m=2$ ) yields the embedded space. Euclidean distances in this space approximate the diffusion distances in the original feature space. The diffusion distance reflects the connectivity of samples over the manifold, such that two samples have a small diffusion distance if they are strongly connected. We optimised all three parameters ( $k$ ,  $\epsilon$ , and  $\alpha$ ) via cross-validation.

**t-SNE.** t-distributed stochastic neighbourhood embedding (t-SNE) represents similarities between samples in terms of probabilities, and aims to derive a lower dimensional embedding in which the probabilities are optimally matched to those in the original feature space<sup>81</sup>. Euclidean distances are converted to probability scores via a Gaussian distribution in the original space and a Student t-distribution in the embedded space. t-SNE often provides good reconstructions of the local manifold structure, though may be less sensitive to the global structure. We optimised two parameters: the perplexity (which determines the number of effective nearest neighbours) and the learning rate (which controls the step size of the gradient descent). We initialised the t-SNE with the PCA embedding as this has been suggested to improve reconstruction quality<sup>82</sup>. Because the algorithm is stochastic, a different solution may be obtained each time the algorithm is run; we therefore fit the algorithm five times and selected the one with the smallest Kullback–Leibler divergence.

**Statistical analyses.** Prediction accuracy was assessed by taking the absolute correlation between corresponding retinotopic and connectopic maps. As a basic test of performance, we compared the correlations against zero via a series of one-sample t-tests; a Holm–Bonferroni correction<sup>83</sup> was applied over multiple comparisons within each algorithm (eccentricity and polar angle maps, variants of the algorithm where applicable, and resting-state and movie-watching tasks for Experiment 2).

Further statistical tests employed the same procedures in both experiments, with the exception that tests for Experiment 2 included an additional factor denoting the task (rest, movie-watching). First, we compared correlations within each algorithm. For the spectral embedding algorithm, correlations were entered into repeated-measures ANOVAs with factors for the map type (eccentricity, polar angle), graph type (weighted

fully-connected, unweighted radius, weighted radius, unweighted kNN, weighted kNN), and task (Experiment 2 only). For the remaining embeddings, correlations were compared between the map types with paired-samples t-tests in Experiment 1, and with two-way ANOVAs with factors of map type and task in Experiment 2. To compare between algorithms, we entered correlations from each algorithm into further repeated-measures ANOVAs with factors for the map type, algorithm type, and task (Experiment 2 only). For spectral embedding, only the unweighted nearest-neighbour variant was entered into this analysis as an example of a better performing variant.

Finally, the spectral embedding (unweighted kNN), diffusion map, and Isomap algorithms were applied to fingerprints derived from different connectivity sources. To compare between sources, these correlations were entered into repeated-measures ANOVAs with factors for the map type, algorithm type, connectivity source (cortex + subcortex, cortex, subcortex, noise, and within-V1), and task (Experiment 2 only).

In all ANOVAs, effect sizes are reported in units of partial and generalised eta-squared<sup>84,85</sup>. A Greenhouse-Geisser sphericity correction<sup>86</sup> was applied to all effects. Where necessary, post-hoc analyses were performed via Tukey contrasts<sup>87</sup>. All statistical tests employed an alpha criterion of 0.05 for determining significance.

## Data availability

All MRI data were obtained from already publicly available repositories: the StudyForrest project (<https://www.studyforrest.org/>) and Human Connectome Project (<https://db.humanconnectome.org/>).

## Code availability

All analysis code is available on the OSF (<https://osf.io/5vupk/>).

Received: 23 February 2022; Accepted: 9 September 2022

Published online: 28 September 2022

## References

1. Bednar, J. A. & Wilson, S. P. Cortical maps. *Neuroscience* **22**, 604–617 (2016).
2. Hubel, D. H. & Wiesel, T. N. Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* **195**, 215–243 (1968).
3. Bonhoeffer, T. & Grinvald, A. Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* **353**, 429–431 (1991).
4. Wandell, B. A. & Winawer, J. Imaging retinotopic maps in the human brain. *Vision Res.* **51**, 718–737 (2011).
5. Penfield, W. & Boldrey, E. Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain* **60**, 389–443 (1937).
6. Sanchez-Panchuelo, R. M., Francis, S., Bowtell, R. & Schluppeck, D. Mapping human somatosensory cortex in individual subjects With 7T functional MRI. *J. Neurophysiol.* **103**, 2544–2556 (2010).
7. Merzenich, M. M. & Brugge, J. F. Representation of the cochlear partition on the superior temporal plane of the macaque monkey. *Brain Res.* **50**, 275–296 (1973).
8. Rauschecker, J. P., Tian, B. & Hauser, M. Processing of complex sounds in the macaque nonprimary auditory cortex. *Science* **268**, 111–114 (1995).
9. Formisano, E. *et al.* Mirror-symmetric tonotopic maps in human primary auditory cortex. *Neuron* **40**, 859–869 (2003).
10. Konkle, T. & Oliva, A. A real-world size organization of object responses in occipitotemporal cortex. *Neuron* **74**, 1114–1124 (2012).
11. Konkle, T. & Caramazza, A. Tripartite organization of the ventral stream by animacy and object size. *J. Neurosci.* **33**, 10235–10242 (2013).
12. Harvey, B. M. & Dumoulin, S. O. A network of topographic numerosity maps in human association cortex. *Nat. Hum. Behav.* **1**, 0036 (2017).
13. Sereno, M. I. & Huang, R.-S. Multisensory maps in parietal cortex. *Curr. Opin. Neurobiol.* **24**, 39–46 (2014).
14. Op de Beeck, H. P., Haushofer, J. & Kanwisher, N. G. Interpreting fMRI data: Maps, modules and dimensions. *Nat. Rev. Neurosci.* **9**, 123–135 (2008).
15. Goldberg, E. Gradiational approach to neocortical functional organization. *J. Clin. Exp. Neuropsychol.* **11**, 489–517 (1989).
16. Margulies, D. S. *et al.* Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc. Natl. Acad. Sci.* **113**, 12574–12579 (2016).
17. Engel, S., Glover, G. H. & Wandell, B. A. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex* **7**, 181–192 (1997).
18. Dumoulin, S. O. & Wandell, B. A. Population receptive field estimates in human visual cortex. *Neuroimage* **39**, 647–660 (2008).
19. Thomas Yeo, B. T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).
20. Power, J. D. *et al.* Functional network organization of the human brain. *Neuron* **72**, 665–678 (2011).
21. Jbabdi, S., Sotiropoulos, S. N. & Behrens, T. E. The topographic connectome. *Curr. Opin. Neurobiol.* **23**, 207–215 (2013).
22. Haak, K. V., Marquand, A. F. & Beckmann, C. F. Connectopic mapping with resting-state fMRI. *Neuroimage* **170**, 83–94 (2018).
23. Haak, K. V. *et al.* Connective field modeling. *Neuroimage* **66**, 376–384 (2013).
24. Bajada, C. J. *et al.* A tutorial and tool for exploring feature similarity gradients with MRI data. *Neuroimage* **221**, 117140 (2020).
25. Huentburg, J. M., Bazin, P.-L. & Margulies, D. S. Large-scale gradients in human cortical organization. *Trends Cogn. Sci.* **22**, 21–31 (2018).
26. Marquand, A. F., Haak, K. V. & Beckmann, C. F. Functional corticostriatal connection topographies predict goal-directed behaviour in humans. *Nat. Hum. Behav.* **1**, 0146 (2017).
27. Vos de Wael, R. *et al.* Anatomical and microstructural determinants of hippocampal subfield functional connectome embedding. *Proc. Natl. Acad. Sci.* **115**, 10154–10159 (2018).
28. Ngo, G. N., Haak, K. V., Beckmann, C. F. & Menon, R. S. Mesoscale hierarchical organization of primary somatosensory cortex captured by resting-state fMRI in humans. *Neuroimage* **235**, 118031 (2021).
29. Blazquez Freches, G., Haak, K. V., Beckmann, C. F. & Mars, R. B. Connectivity gradients on tractography data: Pipeline and example applications. *Hum. Brain Mapp.* <https://doi.org/10.1002/hbm.25623> (2021).
30. Finn, E. S. Is it time to put rest to rest?. *Trends Cogn. Sci.* **25**, 1021–1032 (2021).
31. Jääskeläinen, I. P., Sams, M., Glerean, E. & Ahveninen, J. Movies and narratives as naturalistic stimuli in neuroimaging. *Neuroimage* **224**, 117445 (2021).
32. Knapen, T. Topographic connectivity reveals task-dependent retinotopic processing throughout the human brain. *Proc. Natl. Acad. Sci.* **118**, e2017032118 (2021).



33. Finn, E. S. & Bandettini, P. A. Movie-watching outperforms rest for functional connectivity-based prediction of behavior. *Neuroimage* **235**, 117963 (2021).
34. Hanke, M. *et al.* A high-resolution 7-Tesla fMRI dataset from complex natural stimulation with an audio movie. *Sci. Data* **1**, 1–18 (2014).
35. Van Essen, D. C. *et al.* The Human Connectome Project: A data acquisition perspective. *Neuroimage* **62**, 2222–2231 (2012).
36. Hanke, M. *et al.* A studyforrest extension, simultaneous fMRI and eye gaze recordings during prolonged natural stimulation. *Sci. Data* **3**, 160092 (2016).
37. Sengupta, A. *et al.* A studyforrest extension, retinotopic mapping and localization of higher visual areas. *Sci. Data* **3**, 160093 (2016).
38. Betti, V. *et al.* Natural scenes viewing alters the dynamics of functional connectivity in the human brain. *Neuron* **79**, 782–797 (2013).
39. Wang, J. *et al.* Test-retest reliability of functional connectivity networks during naturalistic fMRI paradigms. *Hum. Brain Mapp.* **38**, 2226–2241 (2017).
40. Lynch, L. K. *et al.* Task-evoked functional connectivity does not explain functional connectivity differences between rest and task conditions. *Hum. Brain Mapp.* **39**, 4939–4948 (2018).
41. Gal, S., Coldham, Y., Tik, N., Bernstein-Eliav, M. & Tavor, I. Act natural: Functional connectivity from naturalistic stimuli fMRI outperforms resting-state in predicting brain activity. *Neuroimage* **258**, 119359 (2022).
42. Vanderwal, T. *et al.* Individual differences in functional connectivity during naturalistic viewing conditions. *Neuroimage* **157**, 521–530 (2017).
43. Byrge, L. *et al.* Video-evoked fMRI BOLD responses are highly consistent across different data acquisition sites. *Hum. Brain Mapp.* **43**, 2972–2991 (2022).
44. Kamps, F. S., Richardson, H., Murty, N. A. R., Kanwisher, N. & Saxe, R. Using child-friendly movie stimuli to study the development of face, place, and object regions from age 3 to 12 years. *Hum. Brain Mapp.* **43**, 2782–2800 (2022).
45. Vanderwal, T., Eilbott, J. & Castellanos, F. X. Movies in the magnet: Naturalistic paradigms in developmental functional neuroimaging. *Dev. Cogn. Neurosci.* **36**, 100600 (2019).
46. Kim, D., Kay, K., Shulman, G. L. & Corbetta, M. A new modular brain organization of the BOLD signal during natural vision. *Cereb. Cortex* **28**, 3065–3081 (2018).
47. Gruskin, D. C. & Patel, G. H. Brain connectivity at rest predicts individual differences in normative activity during movie watching. *Neuroimage* **253**, 119100 (2022).
48. Samko, O., Marshall, A. D. & Rosin, P. L. Selection of the optimal parameter value for the Isomap algorithm. *Pattern Recognit. Lett.* **27**, 968–979 (2006).
49. Vos de Wael, R. *et al.* BrainSpace: A toolbox for the analysis of macroscale gradients in neuroimaging and connectomics datasets. *Commun. Biol.* **3**, 103 (2020).
50. Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**, 825–841 (2002).
51. Smith, S. M. *et al.* Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* **23**, S208–S219 (2004).
52. Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W. & Smith, S. M. FSL. *Neuroimage* **62**, 782–790 (2012).
53. Smith, S. M. Fast robust automated brain extraction. *Hum. Brain Mapp.* **17**, 143–155 (2002).
54. Dale, A. M., Fischl, B. & Sereno, M. I. Cortical Surface-Based Analysis I. Segmentation and Surface Reconstruction. *Neuroimage* **9**, 179–194 (1999).
55. Greve, D. N. & Fischl, B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* **48**, 63–72 (2009).
56. Fischl, B., Sereno, M. I., Tootell, R. B. H. & Dale, A. M. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Hum. Brain Mapp.* **8**, 272–284 (1999).
57. Fischl, B., Sereno, M. I. & Dale, A. M. Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage* **9**, 195–207 (1999).
58. Andersson, J. L. R., Jenkinson, M. & Smith, S. *Non-linear registration, aka spatial normalization (FMRIB technical report TR07JA2)*. (2010).
59. Fischl, B. *et al.* Whole brain segmentation. *Neuron* **33**, 341–355 (2002).
60. Saad, Z. S., Ropella, K. M., Cox, R. W. & DeYoe, E. A. Analysis and use of fMRI response delays. *Hum. Brain Mapp.* **13**, 74–93 (2001).
61. Cox, R. W. AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* **29**, 162–173 (1996).
62. Kriegeskorte, N. *et al.* Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* **60**, 1126–1141 (2008).
63. Cutting, J. E., Brunick, K. L. & Candan, A. Perceiving event dynamics and parsing Hollywood films. *J. Exp. Psychol. Hum. Percept. Perform.* **38**, 1476–1490 (2012).
64. Faul, F., Erdfelder, E., Lang, A.-G. & Buchner, A. G\*Power: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behav. Res. Methods* **39**, 175–191 (2007).
65. Van Essen, D. C., Glasser, M. F., Dierker, D. L., Harwell, J. & Coalson, T. Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cereb. Cortex* **22**, 2241–2262 (2012).
66. Robinson, E. C. *et al.* Multimodal surface matching with higher-order smoothness constraints. *Neuroimage* **167**, 453–465 (2018).
67. Robinson, E. C. *et al.* MSM: A new flexible framework for Multimodal Surface Matching. *Neuroimage* **100**, 414–426 (2014).
68. Glasser, M. F. *et al.* The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* **80**, 105–124 (2013).
69. Smith, S. M. *et al.* Resting-state fMRI in the Human Connectome Project. *Neuroimage* **80**, 144–168 (2013).
70. Benson, N. C. *et al.* The Human Connectome Project 7 Tesla retinotopy dataset: Description and population receptive field analysis. *J. Vis.* **18**, 23 (2018).
71. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
72. Belkin, M. & Niyogi, P. Laplacian Eigenmaps and Spectral Techniques for Embedding and Clustering. in *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic* 585–591 (MIT Press, 2001).
73. von Luxburg, U. A tutorial on spectral clustering. *Stat. Comput.* **17**, 395–416 (2007).
74. Cerliani, L. *et al.* Probabilistic tractography recovers a rostrocaudal trajectory of connectivity variability in the human insular cortex. *Hum. Brain Mapp.* **33**, 2005–2034 (2012).
75. Schölkopf, B., Smola, A., Müller, K.-R. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Comput.* **10**, 1299–1319 (1998).
76. Roweis, S. T. & Saul, L. K. Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**, 2323–2326 (2000).
77. Zhang, Z. & Wang, J. MLE: Modified locally linear embedding using multiple weights. In *Advances in Neural Information Processing Systems* Vol. 19 (eds Schölkopf, B. *et al.*) (MIT Press, 2007).
78. Zhang, Z. & Zha, H. Principal manifolds and nonlinear dimension reduction via local tangent space alignment. *SIAM J. Sci. Comput.* **26**, 313–338 (2002).

79. Tenenbaum, J. B., de Silva, V. & Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319–2323 (2000).
80. Dijkstra, E. W. A note on two problems in connexion with graphs. *Numer. Math.* **1**, 269–271 (1959).
81. van der Maaten, L. & Hinton, G. Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).
82. Kobak, D. & Linderman, G. C. Initialization is critical for preserving global data structure in both t-SNE and UMAP. *Nat. Biotechnol.* **39**, 156–157 (2021).
83. Holm, S. A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* **6**, 65–70 (1979).
84. Bakeman, R. Recommended effect size statistics for repeated measures designs. *Behav. Res. Methods* **37**, 379–384 (2005).
85. Olejnik, S. & Algina, J. Generalized eta and omega squared statistics: Measures of effect size for some common research designs. *Psychol. Methods* **8**, 434–447 (2003).
86. Greenhouse, S. W. & Geisser, S. On methods in the analysis of profile data. *Psychometrika* **24**, 95–112 (1959).
87. Tukey, J. W. Comparing individual means in the analysis of variance. *Biometrics* **5**, 99 (1949).

## Acknowledgements

Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

## Author contributions

Both authors conceived and developed the experiments. D.M.W. performed the data analysis under the supervision of T.J.A. Both authors contributed to the writing of the manuscript and approved the final version for submission.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-20322-4>.

**Correspondence** and requests for materials should be addressed to D.M.W.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022