Temporally-independent functional modes of spontaneous brain activity

Supporting Information

1 Methods in Detail

We acquired 42 10-minute FMRI datasets in 7 sessions from 5 subjects (one subject being scanned on 3 separate occasions). 6 runs were acquired per session. The subjects were healthy adults (ages 18-25, 3 males and 2 females). The imaging protocol used for human studies was approved by the institutional review board at the University of Minnesota. Each of the subjects were scanned using this protocol and provided informed written and verbal consent prior to participating in the research. For two of the acquisition sessions (i.e., two subjects), 3 of the 6 runs were acquired with the subject performing a silent backwards-counting task; those runs were not used in the main TFM analyses reported here, leaving 36 runs of resting FMRI data. For two of the seven sessions, the subjects had their eyes open, and in the other five, the subjects had their eyes closed; we did not find any significant differences in the amplitude of the TFM fluctuations between the eyes-open and eyes-closed datasets.

Data were acquired on a standard 3T Siemens Trio, with 40 mT/m maximum gradients and a slew rate of 200 mT/m/s. A 32 channel receive-only radio-frequency (RF) head coil array was used along with a body RF transmitter. The FMRI time series data were acquired using multiband (MB) accelerated [1, 2, 3] 2D echo planar imaging (EPI) [4] with controlled aliasing [5] to reduce g-factors [6] associated with slice unaliasing. 3 slices were simultaneously excited (MB=3) while also being shifted by 1/3 of the field of view along the phase encode direction, with respect to each other. Each RF coil channel therefore receives a linear combination of the signals from each of the slices, weighted by the coil sensitivity profiles, which with a matrix inversion was used to reconstruct the signal for individual slices as described previously [2]. This slice unaliasing step was implemented as an ICE (Image Calculation Environment, Siemens, Germany) functor and incorporated into the Siemens EPI reconstruction pipeline. For in-plane spatial encoding, 2D single shot EPI was used with a 64x64 matrix (3 mm isotropic) with an echo time (TE) of 30 ms, echo spacing of 0.48 ms; the 3X slice-accelerated sequence achieved a whole brain (42 slices, including cerebellum) TR of 0.8 s. Ernst angles of 60 degrees were used for MB excitation pulses.

At the end of each 10-minute run, we acquired a single image without slice acceleration (MB=1) using a 90 degree flip angle with fully relaxed tissue magnetization; this EPI "reference" image was thus encoded identically to the rest of the FMRI data, (including geometric distortions), but had much better signal-to-noise and contrast-to-noise (between different tissue types). This image was used both as the reference image for head motion correction (aligning all of the FMRI images to this), and was also used to align the FMRI data accurately to the subject's structural image. A T1-weighted structural image was acquired at 1x1x1 mm^3 resolution, using a 3D MPRAGE [7] sequence, with a TR of 2.5 s and inversion time 1 s, resulting in an acquisition time of 6 minutes.

Data was analysed primarily using tools from FSL (FMRIB Software Libraty) [8, 9], as well as using FreeSurfer [10, 11] for cortical surface modelling and FastICA [12] for the temporal independent component analysis (ICA). For visualisation we used tools from FSL and Caret [13].

Each 10-minute FMRI dataset was corrected for head motion using FLIRT (FMRIB's Linear Image Registration Tool [14]), using the full-contrast FMRI image (described above) as the reference, and then high-pass filtered to remove drifts of characteristic timescales slower than approximately 60s. The data was not spatially smoothed, in order to retain maximal spatial detail.[1] The dataset was then fed into MELODIC (Multivariate Exploratory Linear Optimised Decomposition of Independent Components [15]) spatial-ICA with automatic dimensionality estimation (finding on average 95 components), in order to find and remove the majority of the artefacts. Each ICA component was automatically classified into "definitely-artefact" vs. "signal-or-artefact" using two simple heuristics that considered the fraction of timeseries power above 0.2Hz and the fraction of the component's spatial map that overlay low-intensity voxels in the original raw data. Upon comparing these automated classifications with manual inspection (on at least one dataset per subject), we confirmed that $< 1\%$ of non-artefact components were incorrectly labelled as artefact, and about $80 - 90\%$ of the artefactual components were correctly identified. We then regressed the (full space of all) artefact timeseries out of the 4D pre-processed data. Hence the *identification* of artefactual components at this stage was partially incomplete (in order to avoid risking removing valid resting-state neural-related signal), but the artefact *removal* from the data (achieved by regressing out the full space of the artefacts and not just their unique variance) was more aggressive than simply subtracting (the outer product of the timecourse and spatial map for) those components from the data. We aim to identify any remaining artefacts at the later stage of group-wise high-dimensional spatial-ICA.

We obtained an accurate registration between the full-contrast reference FMRI image and the subject's structural, initially with a 6 degree-of-freedom (rigid body) affine registration using FLIRT. This alignment was then refined using a version of the brain-boundary-registration (BBR [16]) algorithm that has recently been added into FLIRT. This optimises the functional-structural alignment by maximising the intensity difference in the FMRI image intensities on either side of the white-matter boundaries (derived from the structural using FAST (FMRIB's Automated Segmentation Tool [17]) tissue-type segmentation, and achieving greater accuracy and robustness than when using the whole original image.

We then mapped each 4D dataset into a standard co-ordinate system currently being developed for the NIH Human Connectome Project [18]; it is a combination of cortical surface *vertices* (from a FreeSurfer cortical surface model that was registered and resampled to a standard mesh representation of the 2D cortical surface) and 3D sub-cortical/cerebellar *voxels* (from a 3D nonlinear registration of the structural image to the MNI152 standard template image using FNIRT (FMRIB's Nonlinear Image Registration Tool [19, 20]) followed by resampling onto a standard set of subcortical grey matter voxels). This procedure attempts to optimise alignment across subjects on the surface [21], while at the same time ensuring that subcortical structures are well aligned. We term this co-ordinate system "grayordinates" as we are attempting to model all of the grey matter. The set of grayordinates can be arranged as a single long vector of vertices and voxels, and thus each FMRI timeseries dataset can be represented in this standard system as a 2D matrix of grayordinates X timepoints. Because both voxels and vertices are resampled onto a standard mesh or subcortical voxel set, these matrices have one-to-one correspondence across subjects (e.g., row 37 corresponds to the same spatial location in all subjects).

---

[1] Purely for display of the axial/coronal/sagittal slices, maps were generated by regressing the relevant TFM/RSN timeseries into the original 4D datasets after they had been warped into MNI152 2mm standard space using FNIRT, and smoothed by a Gaussian kernel of half-width 2mm, in order to ameliorate slightly the cross-subject spatial variability in 3D (non-surface) space.

We temporally concatenated all subjects' datasets in this format, to feed into the following group-wise ICA. We applied a high-dimensional spatial-ICA using MELODIC, in order to further identify artefactual components, as well as achieving a high-dimensional functional parcellation of the group data. From a 200-dimensional spatial-ICA, we manually identified 58 components as artefact on the basis of their spatial appearance (e.g., clearly outlining major arteries or veins) and temporal spectra (e.g., having a significant fraction of the power above 0.2Hz). These 58 components' timeseries were then regressed out of the remaining 142 timeseries, leaving 142 functional nodes' timecourses to feed into the temporal-ICA TFM network analysis.

The 142 timeseries were fed into temporal-ICA with an ICA dimensionality of 21, i.e., in order to find the 21 strongest TFMs present in the group data as a whole. We used the "symmetric" estimation approach, with the "tanh" nonlinearity, both options in the FastICA software. We also split the data into two groups of separate subjects (one group containing 21 runs from two subjects, one of whom had 3 sessions of acquisitions, and the other group containing 15 runs from 3 subjects). We carried out the temporal-ICA separately for the two groups (after first reducing the data dimensionality to 21 for the two groups together), in order to test the reproducibility of the components (TFMs) found.

The dimensionality of the temporal-ICA was chosen somewhat arbitrarily; there was no objective measure that would suggest this choice, just as has been the case in the literature using spatial-ICA for finding RSNs. However, we believe that a dimensionality of 21, given the current dataset, represents a good balance between richness/interpretability and robustness/estimability. Furthermore, as we raised the number of components estimated, we found that many of the most robust TFMs were still found (with almost no change spatially or temporally) up to temporal-ICA dimensionalities of at least 60.

2 The TFM model

We now present the mathematical model underlying the TFM analyses, including the initial high-dimensional spatial-ICA.

The spatial-ICA model is:

$$X_{(V \times T)} = S_{s(V \times I)} \times A_{s(I \times T)} + E_s, \tag{1}$$

where X is the data matrix of size *V*oxels (or grayordinates) × *T*imepoints, S_s are the estimated spatial maps (formally the ICA sources, also referred to as the "network nodes"), A_s are the associated timecourses (the ICA mixing matrix), I is the number of components estimated (in our case 200) and E_s is the residual noise remaining after the I-dimensional PCA-based dimensionality reduction carried out as the first stage of the spatial-ICA.

We now identify $J \leq I$ of the components as artefact, based on their spatial maps and/or timecourses, and regress these out of the remaining $K = I - J$ components to give K "cleaned" timeseries A_{sc}:

$$A^T_{sc(T \times K)} = A^T_{sK(T \times K)} - A^T_{sJ(T \times J)} \times ((A^T_{sJ(T \times J)})^{-1} \times A^T_{sK(T \times K)}), \tag{2}$$

where $(A^T_{sJ})^{-1}$ is the Moore-Penrose pseudo-inverse of the set of artefact timeseries.

We now ignore S_s (the nodes' spatial maps) and work only with the nodes' cleaned timeseries A_{sc}. These are fed into temporal-ICA:

$$A_{sc(K \times T)} = A_{t(K \times L)} \times S_{t(L \times T)} + E_t, \tag{3}$$

where A_t are the TFM node-weight vectors (the ICA mixing matrix), S_t are the TFM timecourses (the ICA sources), L is the number of components estimated (in our case 21, the number of TFMs), and E_t is the residual noise remaining after the L-dimensional PCA-based dimensionality reduction carried out as the first stage of the temporal-ICA.

Thus we can write the full analysis in a single equation:

$$X_{(V \times T)} = S_{s(V \times K)} \times A_{t(K \times L)} \times S_{t(L \times T)} + E, \tag{4}$$

where E combines all of the noise and artefact aspects of the data. The data is therefore represented as a product of three matrices: 142 nodes' spatial maps S_s, 21 TFMs' timecourses S_t and a central mixing matrix A_t that describes the node-membership for each TFM. The lth column in A_t contains the 142 node weights for TFM l; these weights can be positive or negative, and where we see significant positive *and* negative values in a single column, the analysis is telling us that different areas of the brain are anti-correlated with each other, with respect to the particular behaviour of that specific TFM. Figure 3B in the main paper is showing A_t.

If we combine S_s and A_t together into a new matrix S_m:

$$S_{m(V \times L)} = S_{s(V \times K)} \times A_{t(K \times L)}, \tag{5}$$

this gives us the (full voxelwise) spatial maps for each TFM. Alternatively, we can regress the TFM timecourses into the full original data, to find the TFM spatial maps, which gives similar results, but which also allows for a degree of consistency testing by generating separate spatial maps for each run/subject, and then feeding these into a t-test.

3 Relating the TFM model to LiNGAM network modelling

Another approach to network modelling, also utilising of temporal-ICA, is LiNGAM (Linear, Non-Gaussian, Acyclic causal Models [22]); this uses the ICA mixing matrix to estimate a *single* full network model, including (causal) directionalities. Our interpretation of temporal-ICA components as TFM timeseries and node-weight vectors is different from the interpretation made by LiNGAM, despite the fact that both are based on the same temporal-ICA decomposition. The “TFM interpretation” is a generalisation of the LiNGAM one, in the sense that it allows for a greater number of biological causes to drive the different components being temporally distinguished from each other. By making stronger assumptions about the form of the network (in particular, the assumptions that there is a single, stationary, acyclic, network, and that every node has its own unique external/hidden input[2]), LiNGAM is able to draw stronger and richer conclusions about the network, including a full effective connectivity model of directionalities, plus knowledge of which nodes are directly connected to which other nodes. On the other hand, the simpler TFM interpretation, while not making any claims about directionality, or whether certain connections are direct or indirect, is more “robust”, in the sense that each separate TFM can have been created by virtue of a single external input (in which case the node-weight vector tells us about the set of nodes that is driven, directly or indirectly, by this input), *or* by a hidden factor causing that TFM (i.e., both its activity and its connections; no distinction is made) to turn on and off as a whole. In either of these scenarios, if an external input feeds into multiple viewed nodes, the TFM interpretation is still valid, whereas the LiNGAM model is broken. As noted, however, this comes at the price of making a less rich interpretation of the *internal* network structure of each TFM.[3]

---

[2] In this context an “external input” is either a truly external input to the brain, such as sensory input, or is a part of the brain that is not being imaged and modelled.

[3] More explicitly: both the “TFM” interpretation and LiNGAM begin with the same decomposition of the (nodes X timecourses) data; X=AS, where S are the sources (timeseries) and A is the ICA mixing matrix. The TFM interpretation simply sees each column in A as a set of node weights describing the (weighted and signed) involvement of each node in a given TFM, whereas LiNGAM takes A and goes much further, turning A into $(I\text{-}B)^{-1}C$, where B is the exact causal (acyclic) network, and C (mapping of external inputs into viewed nodes) must be assumed to be identity.

4 Spatial-ICA group results

The initial groupwise spatial-ICA was run at a dimensionality of 200 using MELODIC. We identified 58 components as artefact and regressed these timecourses out of the remaining 142. The spatial maps for the 142 non-artefact components are shown in Figure 1, and in Figure 2 we show the left-most part in more detail, showing a 20-dimensional parcellation of visual areas. For each component we show up to 3 slices that best represent the component's spatial map, overlaid onto the average structural image from all subjects. At this dimensionality the components in general contain considerably fewer dis-contiguous regions than is the case at lower dimensionalities; the majority of components are either one contiguous region or a bilateral region-pair (though the actual maps contain a range of positive and negative values, so of course such statements depend on having applied a threshold first).

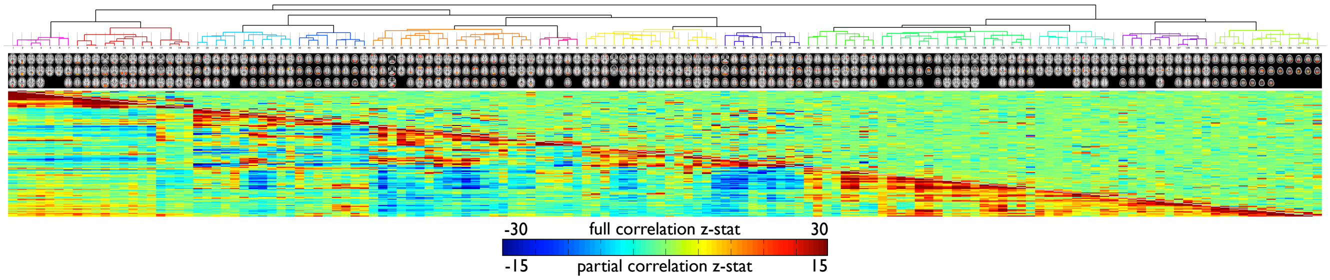

Figure 1: **Results from the initial groupwise spatial-ICA.** Each component ("network node") is shown with up to 3 slices that best represent the component's spatial map. The full and partial temporal correlation matrices are shown below and above the diagonal respectively. The full correlation matrix was used to reorder the nodes, bringing nodes with similar timecourses together, and causing the correlation matrix to have blocks of related nodes on the matrix diagonal.

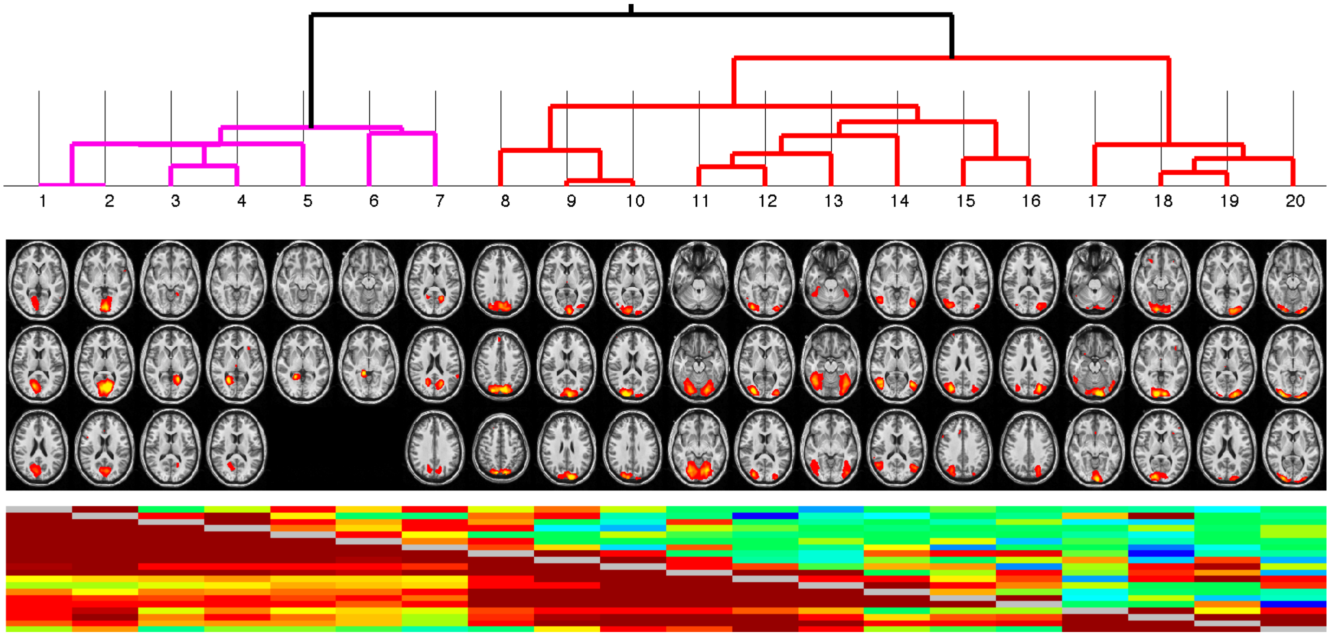

Figure 2: The same results as in Figure 1, but zooming in on the left-most part, showing visual areas.

In order to help visualise the results of the initial high-dimensional spatial-ICA we estimated the temporal correlation matrix and used this to carry out simple hierarchical clustering; this does not affect the the temporal-ICA TFM analysis, although it is used to *re-order* the components (nodes) for all further analysis.

Below the diagonal in Figure 1 is shown the *full correlation* of all timecourses with each other. This is estimated by first concatenating all subjects' timecourses temporally, computing the full correlation between all pairs of timecourses, and then converting this to a z-statistic, after correcting the temporal degrees of freedom for the autocorrelation. Blue matrix elements imply negative average correlation, i.e., anti-correlation between network nodes.

Above the diagonal is shown the *partial correlation* [23, 24] matrix; this simple approach aims to emphasise the *direct* network connections, rather than showing all direct *and* indirect connections. To achieve this, before any two nodes' timeseries are correlated with each other, all other nodes' timeseries are first regressed out of both. This removed on average approximately one third of the temporal standard deviation from the nodes' timeseries.

5 Full set of 21 TFM results

We now show the full set of spatial maps for all 21 TFMs, both on the inflated cortical surface, and with axial/coronal/sagittal 2D slices, in Figures 3-23. All 2D slices are left-right reversed, and for all cortical surface renderings what appears to be the left (or right) hemisphere is the left (or right) hemisphere.

The TFM spatial maps are obtained by regressing the TFM timecourses into the original datasets, and averaging the resulting maps across all runs/subjects, with a within-run significance level of $p<0.05$, corrected for multiple comparisons across space (although, arguably, as these are data-driven decompositions, it does not make much sense to consider such p-values as very meaningful; what matters more is the reproducibility of the components, as described later). Carrying out a mixed-effects analysis across runs/subjects gives extremely similar results, as did also generating the spatial maps by multiplying the nodes' maps by the TFM node weight vectors as described above, so we only show the average maps here.

The sign of the spatial maps (and associated timecourses) is arbitrary; flipping the sign of one spatial map and its associated timecourse gives an equivalent model of the data. In the case of RSNs, which generally have one dominant sign in the spatial maps, it is conventional to force the dominant sign to be positive (and to adjust the sign of the timecourse accordingly), for convenience of interpretation. However, with TFMs, we generally see significant regions of both signs, and it should be remembered that the "positive" (red-yellow) and "negative" (blue) regions are anti-correlated with respect to each other for a given TFM, but both are equally associated with that TFM.

Power spectra for the TFM timeseries can be seen in Figure 24. These are calculated using Welch's method in Matlab, using non-overlapping windows of 222 timepoints to estimate temporal power spectra from all runs, averaging spectra across all runs, and then smoothing by fitting a local line to 10 neighbouring timepoints.

We now describe briefly the major components of all 21 TFMs, with the ordering somewhat "thematically" determined, starting with the three strongly visual-related TFMs, 2, 4 and 8.

TFM 2 contains all of the low-level visual areas, as well as lateral geniculate nucleus (LGN), with very little anti-correlation seen.

TFM 4 shows anti-correlation between foveal visual areas and high eccentricity (high visual angle away from the direct gaze direction). Correlations outside of the visual areas are bilateral and almost all of the same sign as the foveal region (as opposed to the high eccentricity). The anti-correlation seen is consistent with *de*-activation found in task studies [25, 26], the latter showing that the negative BOLD in the high vs. low eccentricity parts of V1 corresponds to a decrease in neural activity.

TFM 8 shows higher-level visual areas, positively correlated with the "dorsal visual stream" and including the frontal eye fields (FEF). This has also been referred to as the "task positive network", and can be seen clearly, for example, in an attention task in [27]. The anti-correlated regions include several parts of the default mode network, including (somewhat weakly compared with some other TFMs) posterior cingulate cortex (PCC), and, strongly, parts of primary auditory.

TFM 1 also contains much of the dorsal visual stream, similar to TFM 8. However, whereas in TFM 8 the anti-correlated regions include DMN/auditory areas, here the anti-correlation is strongly found in primary motor-somatosensory (mouth or face), along with SMA (supplementary motor area), posterior putamen/pallidum and pre-motor thalamus.

TFM 3 is dominated by primary and secondary sensori-motor areas (particularly at the level of the upper extremity representation), along with superior temporal sulcus and cerebellum VI. The anti-correlated regions are weaker and less widespread, possibly including parts of superior parietal lobule and anterior cingulate.

TFM 6 has strong anti-correlation between primary sensori-motor / SMA / pre-motor thalamus, SII, cingulate motor vs. superior temporal sulcus (primary auditory) / medial areas of the DMN / Brodmann 45.

TFM 10 contains widespread somatosensory and motor areas correlated with posterior superior temporal sulcus and areas 45 and 47l, anticorrelated with posterior IPS (intraparietal sulcus), angular gyrus and DLPFC (dorso-lateral pre-frontal cortex).

TFM 7 strongly involves medial parts of the DMN (PCC/precuneus and anterior cingulate), along with primary and early association auditory cortex, parts of the parietal-occipital fissure and frontal area 10. Anti-correlated regions include superior temporal sulcus and and Broca (area 45 more than 44).

TFM 9 includes dorso-lateral pre-frontal cortex / inferior parietal lobule / lateral occipital areas, anti-correlated with supramarginal/angular gyrus.

TFM 11 contains large areas of the DMN, including PCC/precuneus, anterior cingulate and widespread angular gyrus, along with Brodmann 22 and frontal areas 9/46. These areas are very similar to the "semantic network" found in a meta-analysis of 120 semantic system studies [28]. Anti-correlated regions are less prominant, including some primary auditory (Brodmann 42) and higher-level visual.

TFM 12 contains parts of the semantic network on the left anticorrelated with right somatosensory association areas.

TFM 13 is the only TFM to strongly cover Broca 44 (as opposed to 45). This forms part of a strongly lateralised "language" network, including bilateral supramarginal gyrus, which anti-correlates with much of the known DMN (though without including auditory), with the medial-frontal part of the DMN being strongly lateralised towards the right.

TFM 14 sees extrastriate visual correlating with parts of somatosensory cortex and retrosplenial cortex. Anti-correlations include large parts of poterior right frontal cortex, striate visual cortex and left cerebellum.

TFM 16 contains bilateral superior temporal sulcus correlated with parts of premotor bilaterally, retrosplenial cortex, DLPFC, and medial parietal cortex. Anticorrelated areas inclue pre-SMA and more anterior medial frontal cortex.

TFM 17 is dominated by a distinctive posterior-medial pattern containing two close (but not connected) areas. Because of the position in/near the PCC, this TFM is probably merged into many previous PCC-seeded DMN analyses. Anti-correlated regions include area 39 and Brodmann 6 and 8.

TFM 18 is dominated by somatosensory and motor regions, with focal anti-correlation in areas 39/40.

TFM 19 contains parts of the DMN and strong involvement of Crus I in the cerebellum. Anticorrelation includes primary and association auditory cortex.

TFM 20 contains some higher-level visual, superior parieto-occipital sulcus, anterior insula, and posterior IPS, and anti-correlated regions in retrosplenial cortex (areas 29 and 30).

TFM 21 sees DMN areas (including PCC) anticorrelated with lateral-frontal regions often seen as part of the "executive control" RSN.

TFMs 5 and 15 are unlike the other TFMs in that they show very widespread activity across much of the grey matter, with only very small areas of anti-correlation. Their temporal power spectra (shown in the SI), while being largely similar to the other TFMs in the frequency range below 0.2Hz, have raised power above 0.2Hz. There could well be some non-neural physiological contribution to these TFMs, although the temporal resolution is not quite high enough to avoid aliasing of cardiac pulsation, and hence it is not easy to differentiate (for example) cardiac and breathing effects in the spectra. It is interesting that TFM 15 includes anti-correlation in anterior insula (probably lateral agranular insula), as well as (more weakly) in the putamen. There is evidence for involvement of of the insula in heart rate control [29, 30, 31] and other autonomic activity [32]; it is possible that the signal seen here in the insula is neural in origin, and related to vascular (as opposed to neural) effects seen in grey matter in virtually all other regions. We found an almost identical TFM, including the anterior insula anti-correlation, in a separate study's dataset, as discussed below.

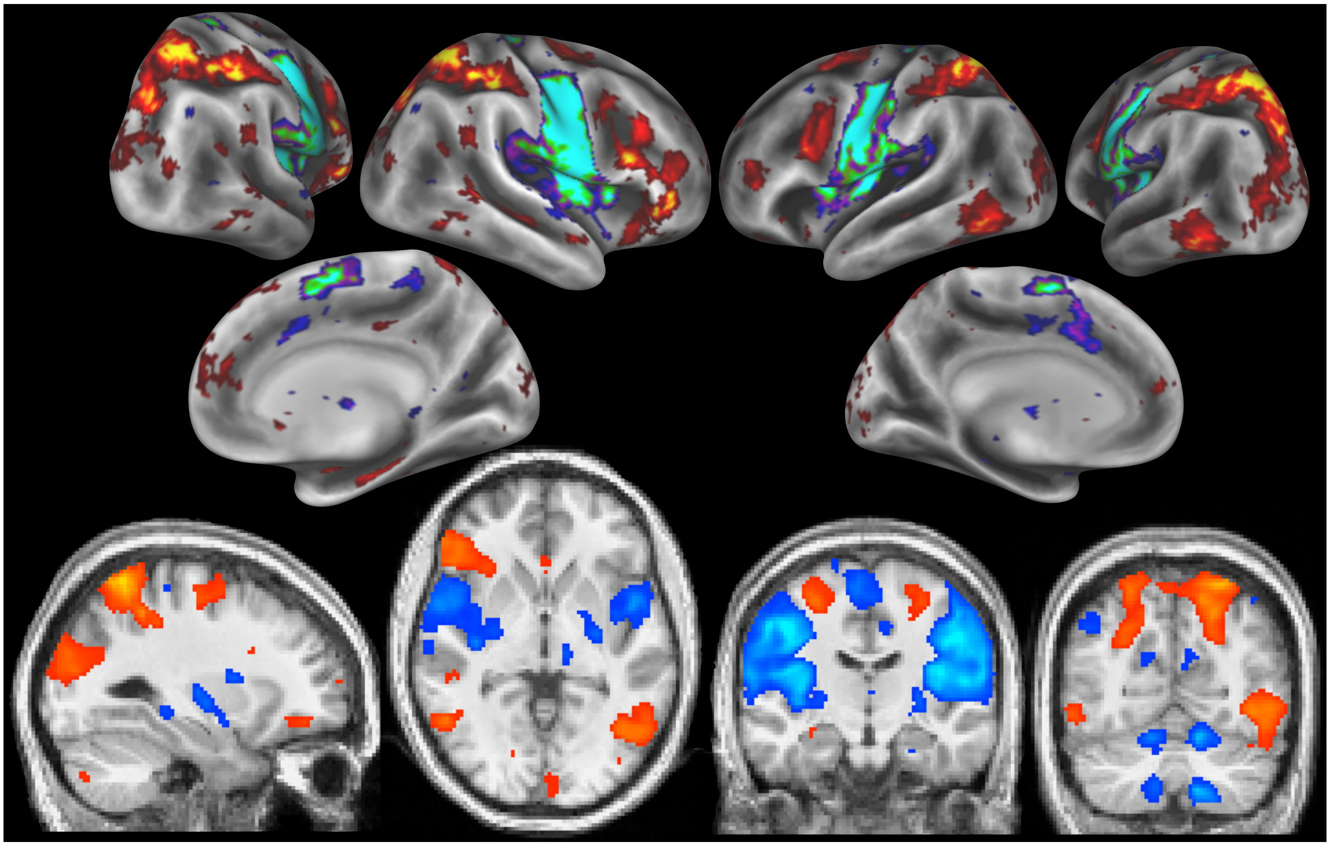

Figure 3: **Spatial map for TFM 1.**

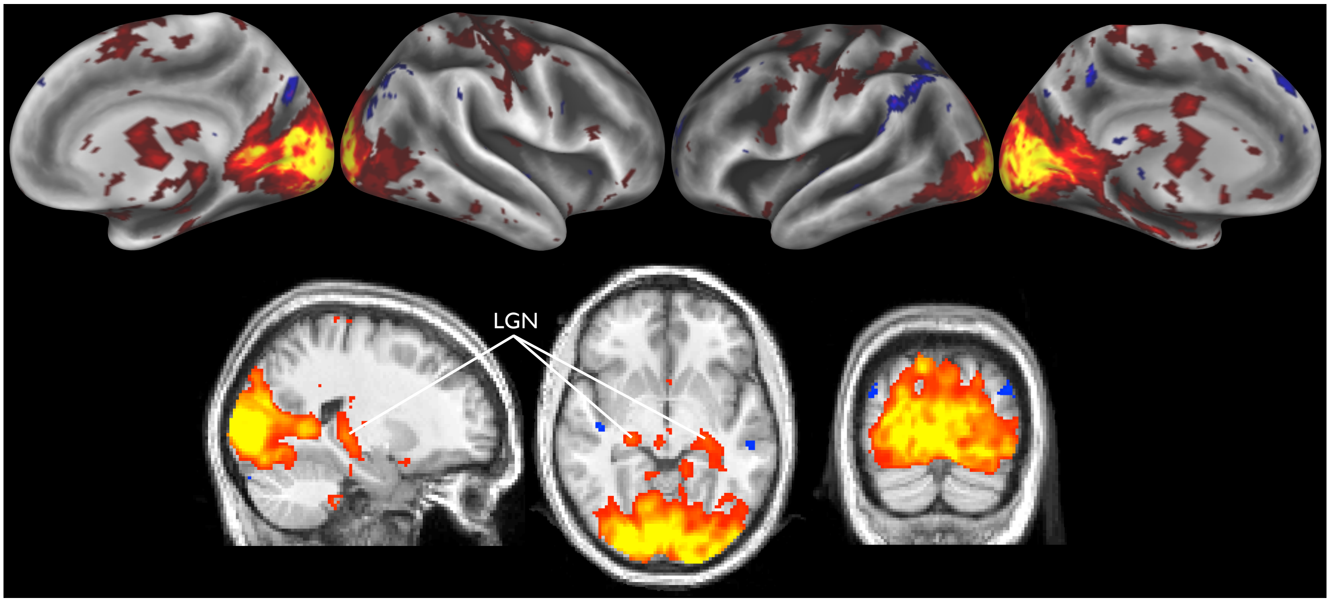

Figure 4: **Spatial map for TFM 2.**

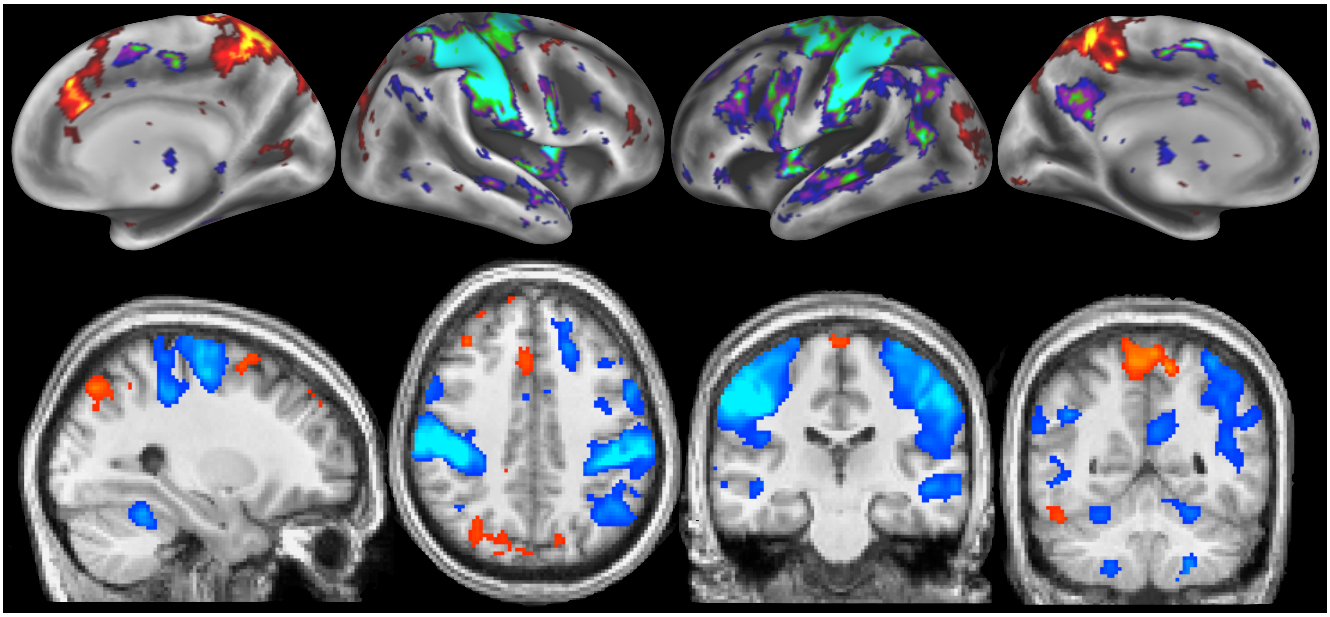

Figure 5: **Spatial map for TFM 3.**

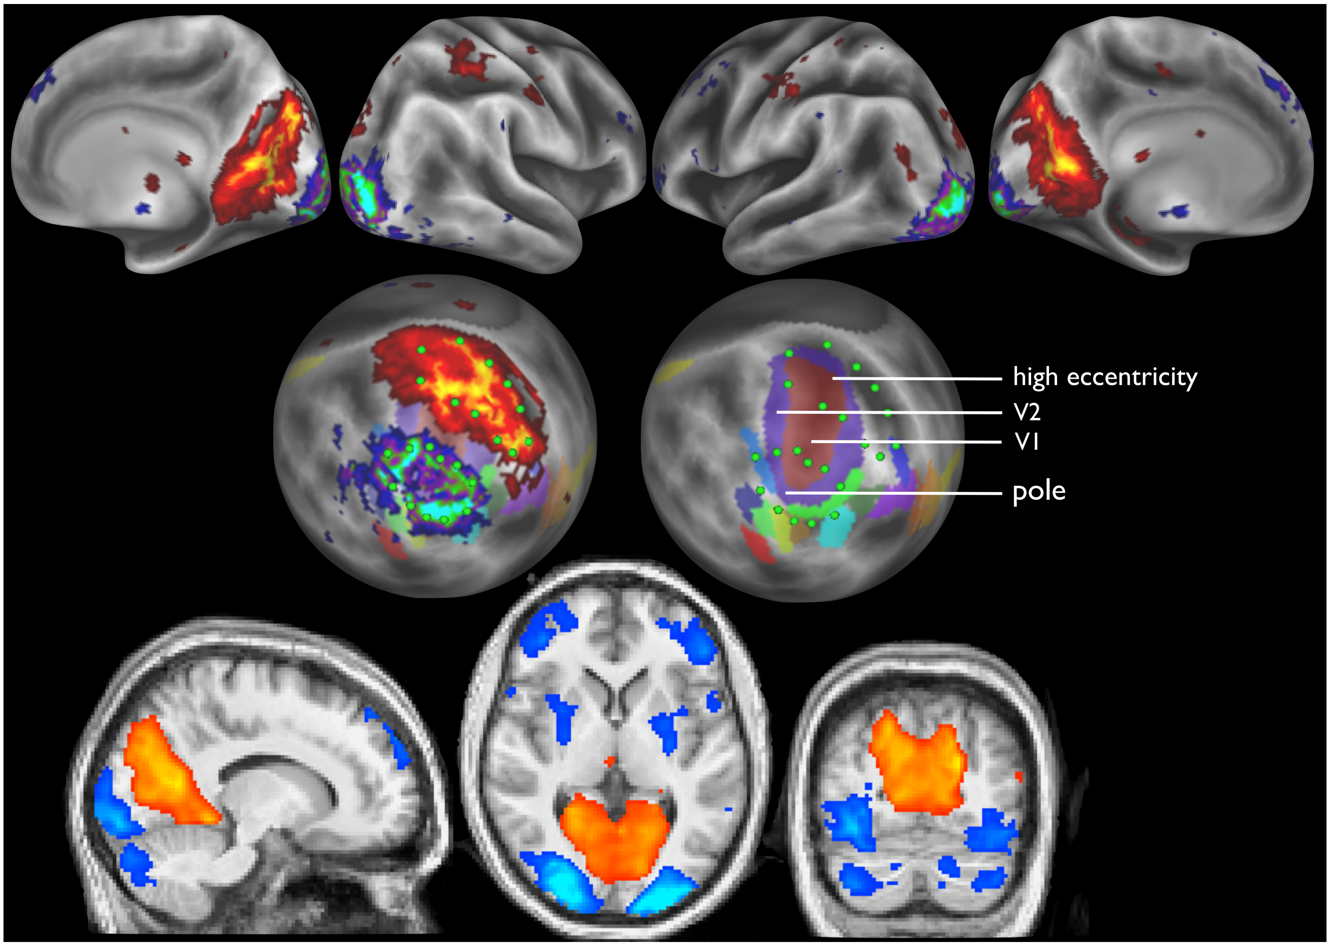

Figure 6: **Spatial map for TFM 4.**

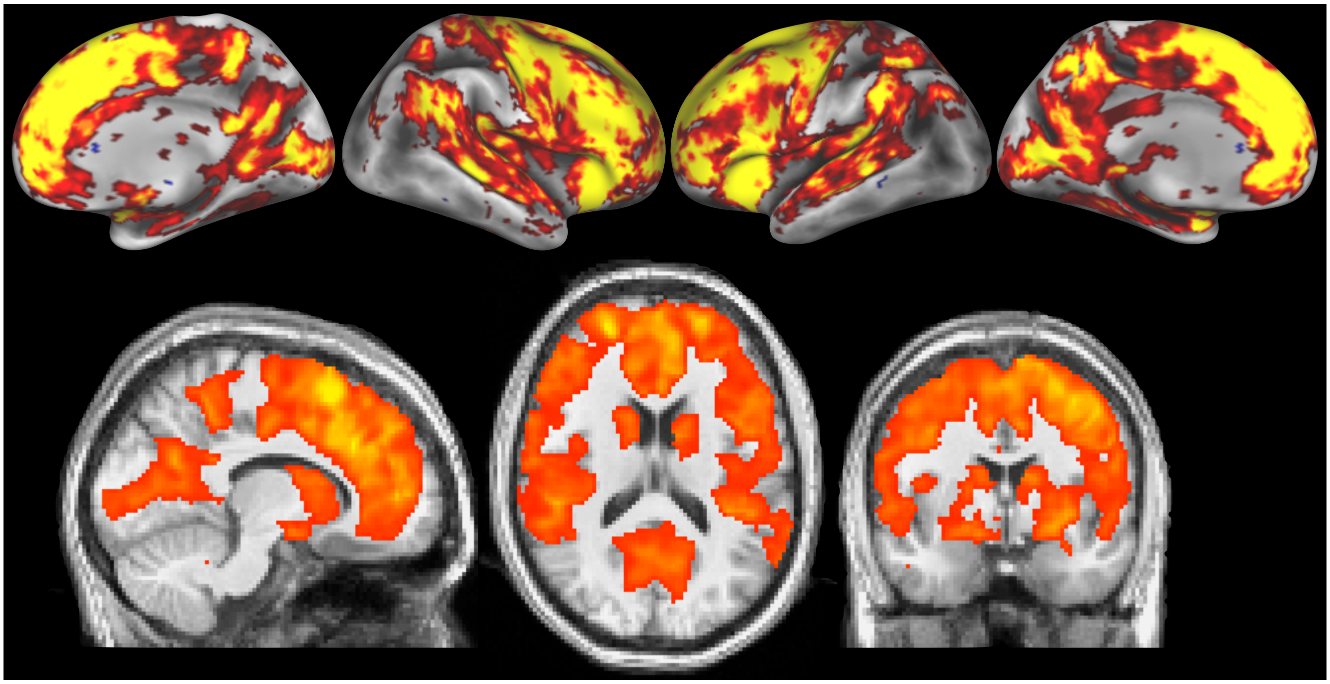

Figure 7: **Spatial map for TFM 5.**

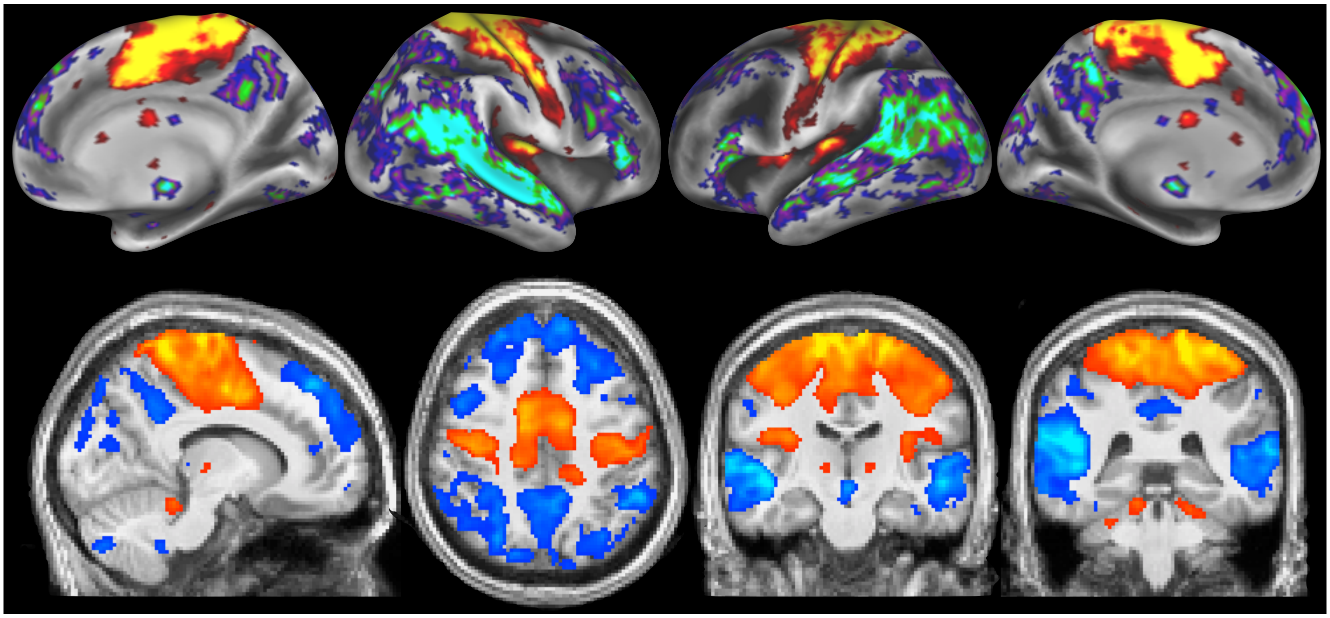

Figure 8: **Spatial map for TFM 6.**

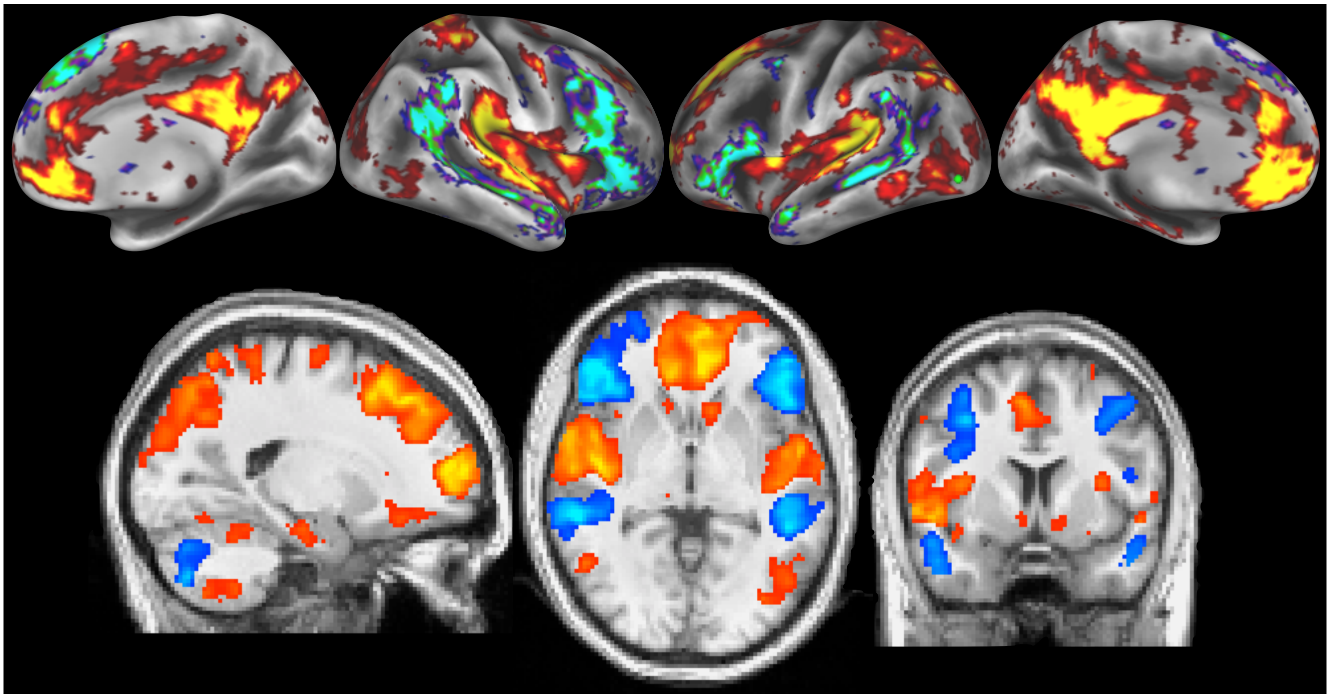

Figure 9: **Spatial map for TFM 7.**

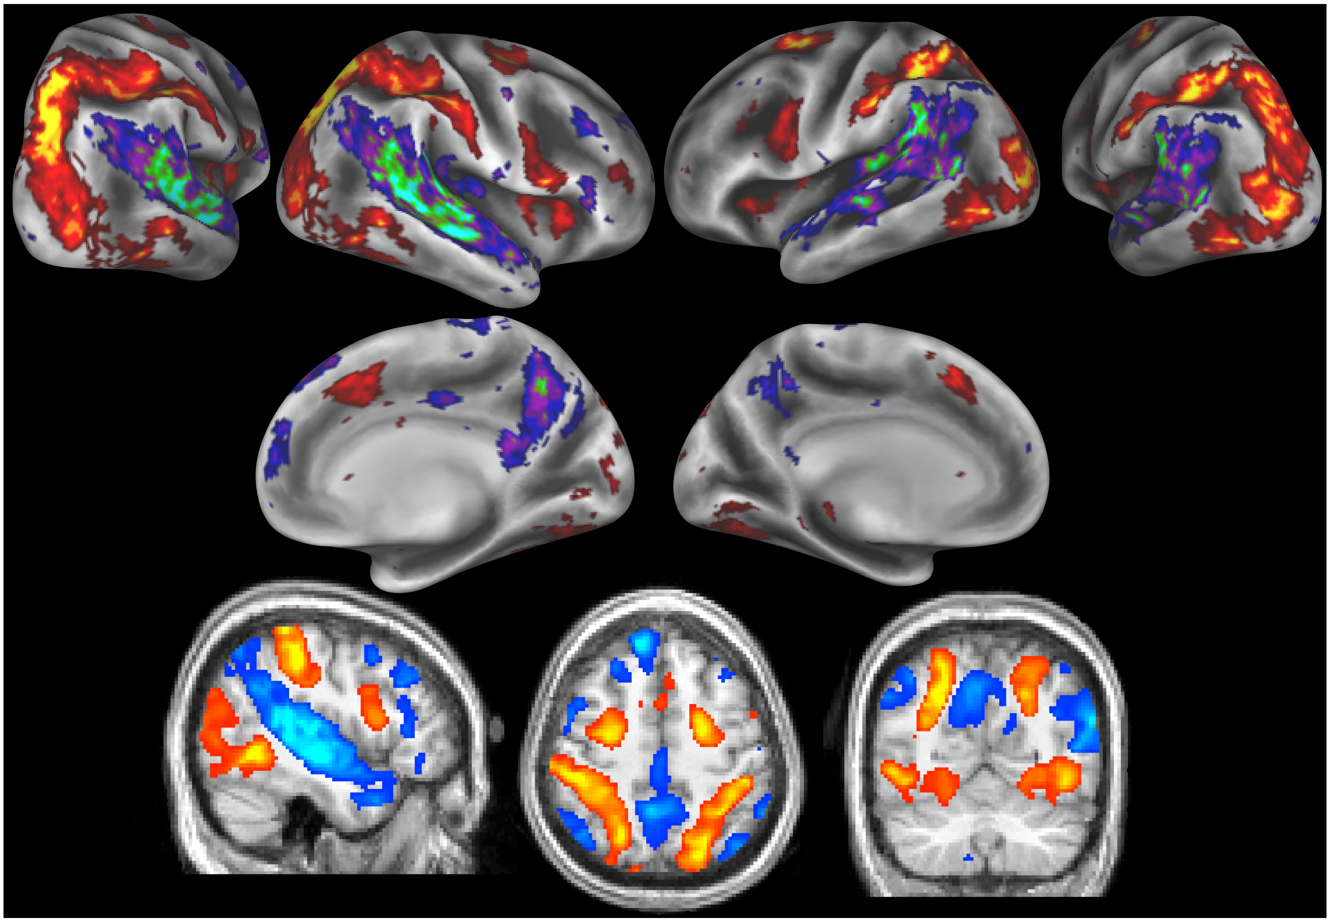

Figure 10: **Spatial map for TFM 8.**

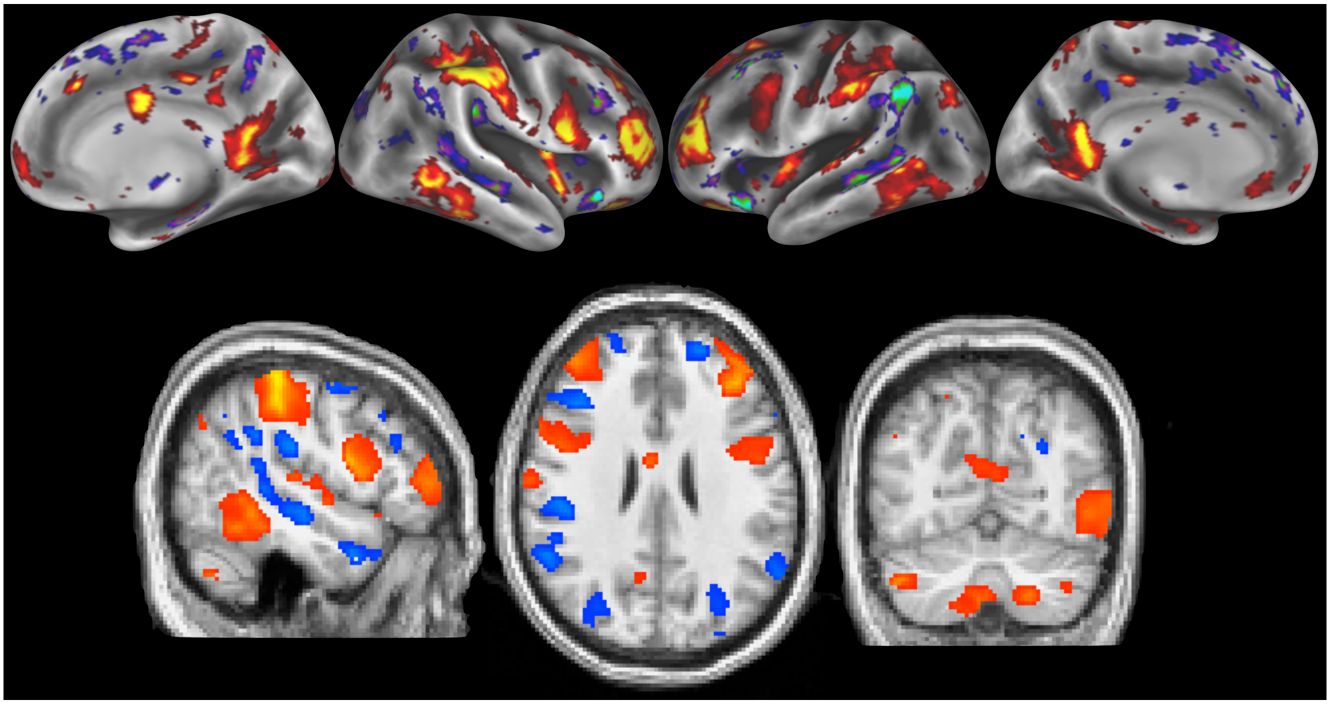

Figure 11: **Spatial map for TFM 9.**

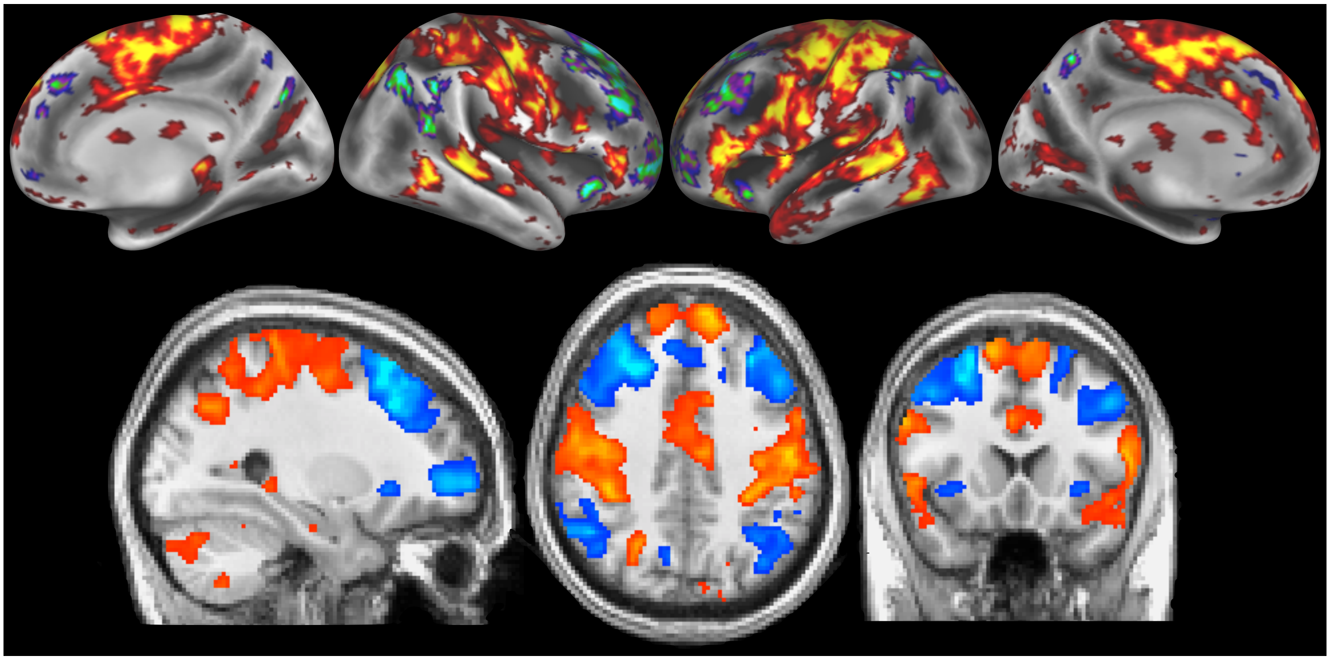

Figure 12: **Spatial map for TFM 10.**

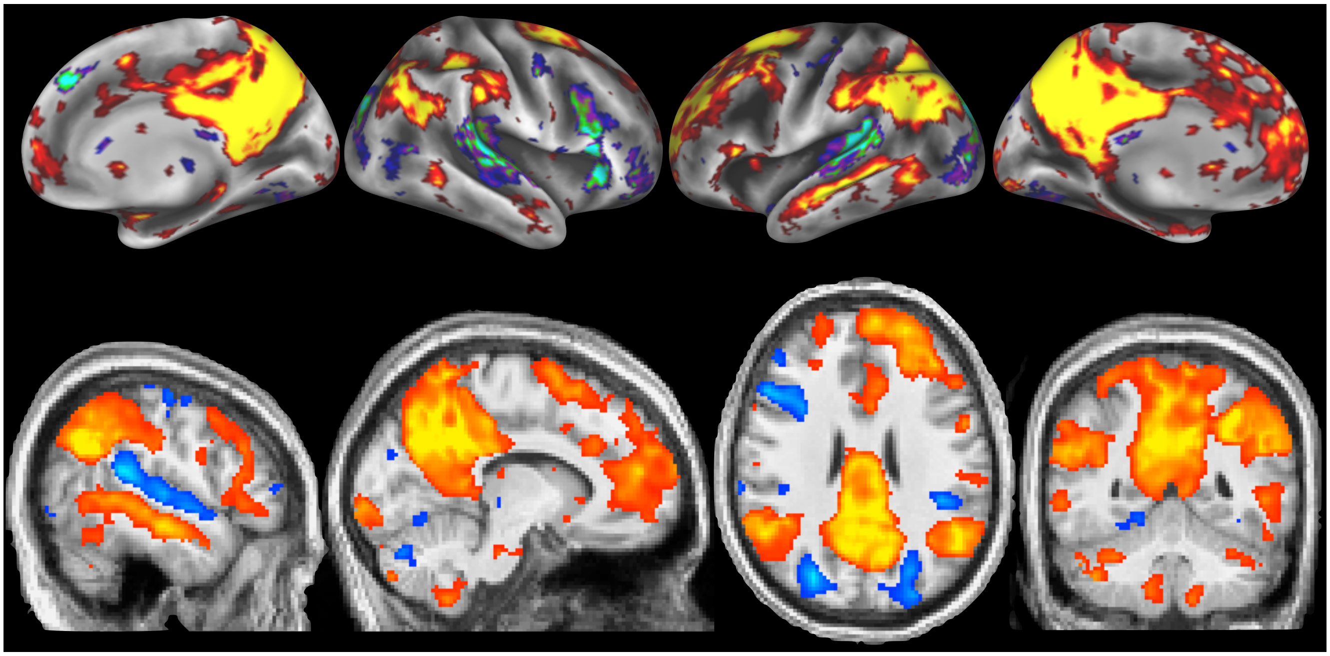

Figure 13: **Spatial map for TFM 11.**

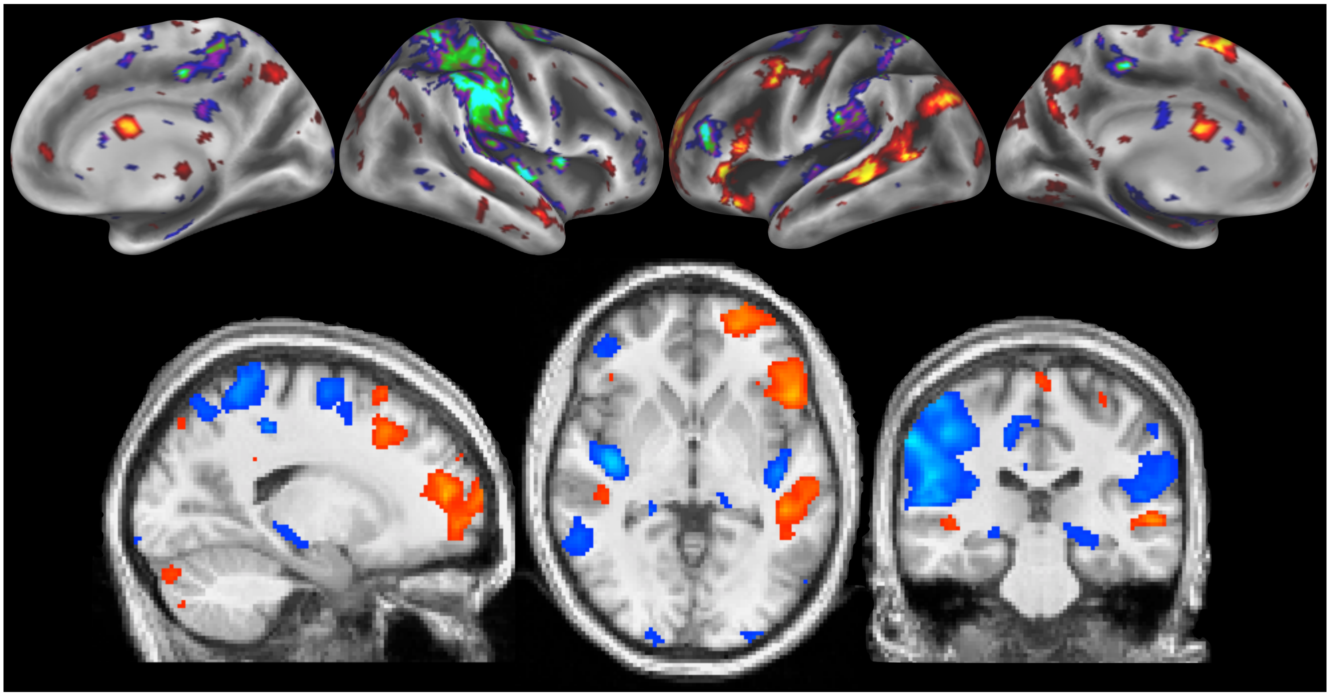

Figure 14: **Spatial map for TFM 12.**

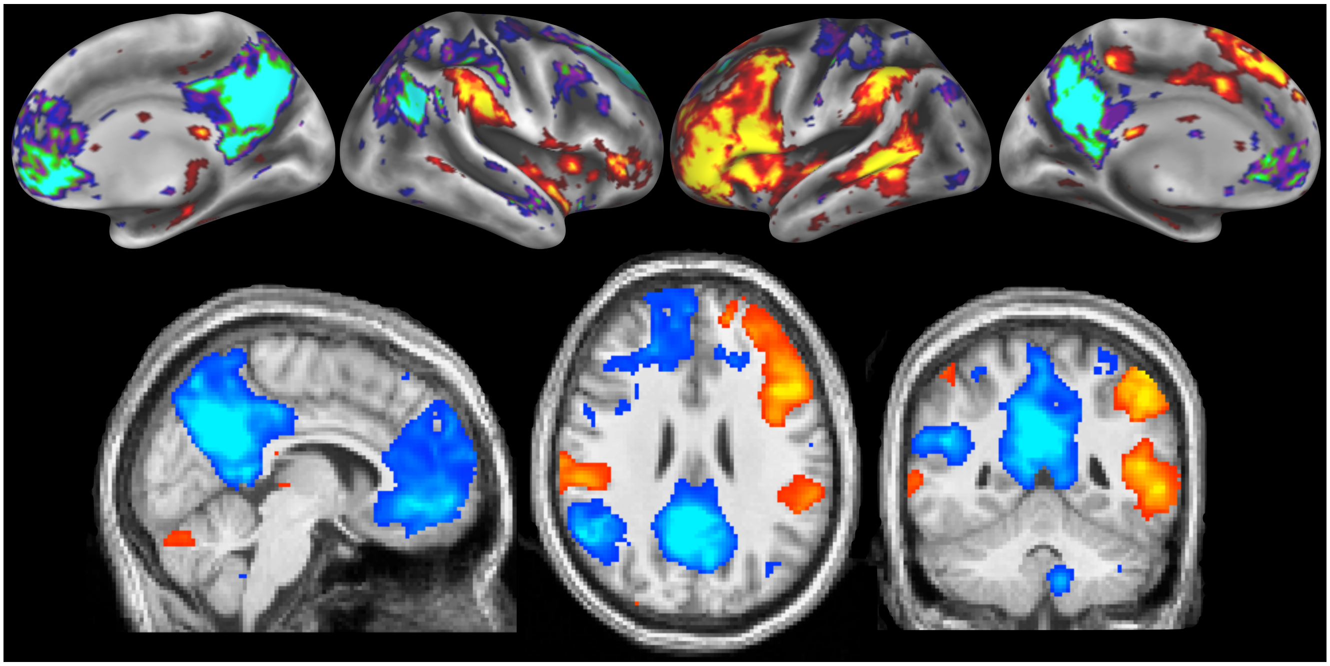

Figure 15: **Spatial map for TFM 13.**

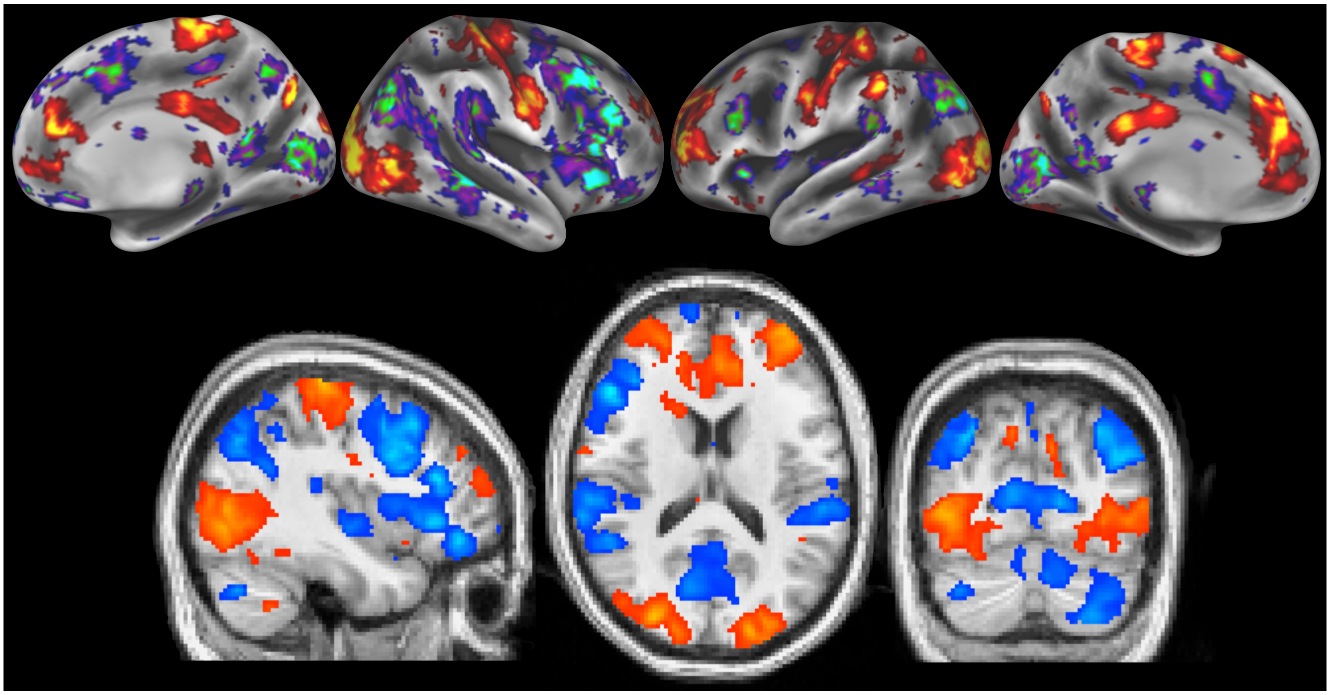

Figure 16: **Spatial map for TFM 14.**

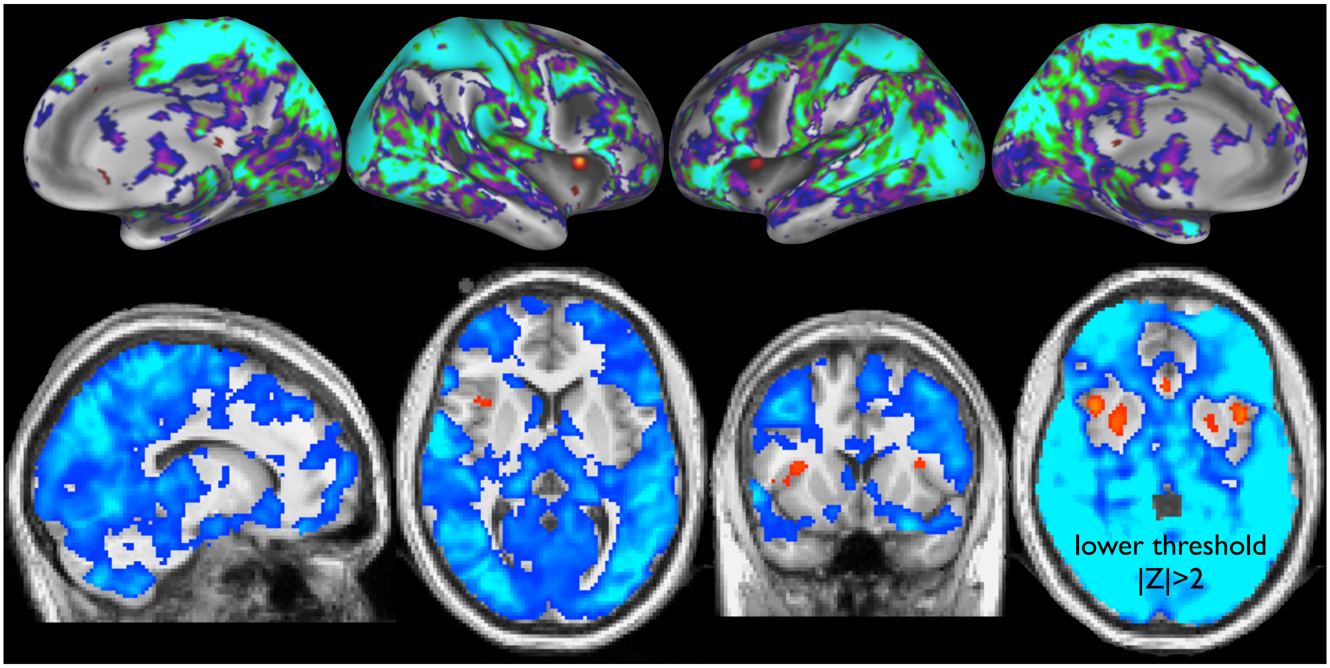

Figure 17: **Spatial map for TFM 15.**

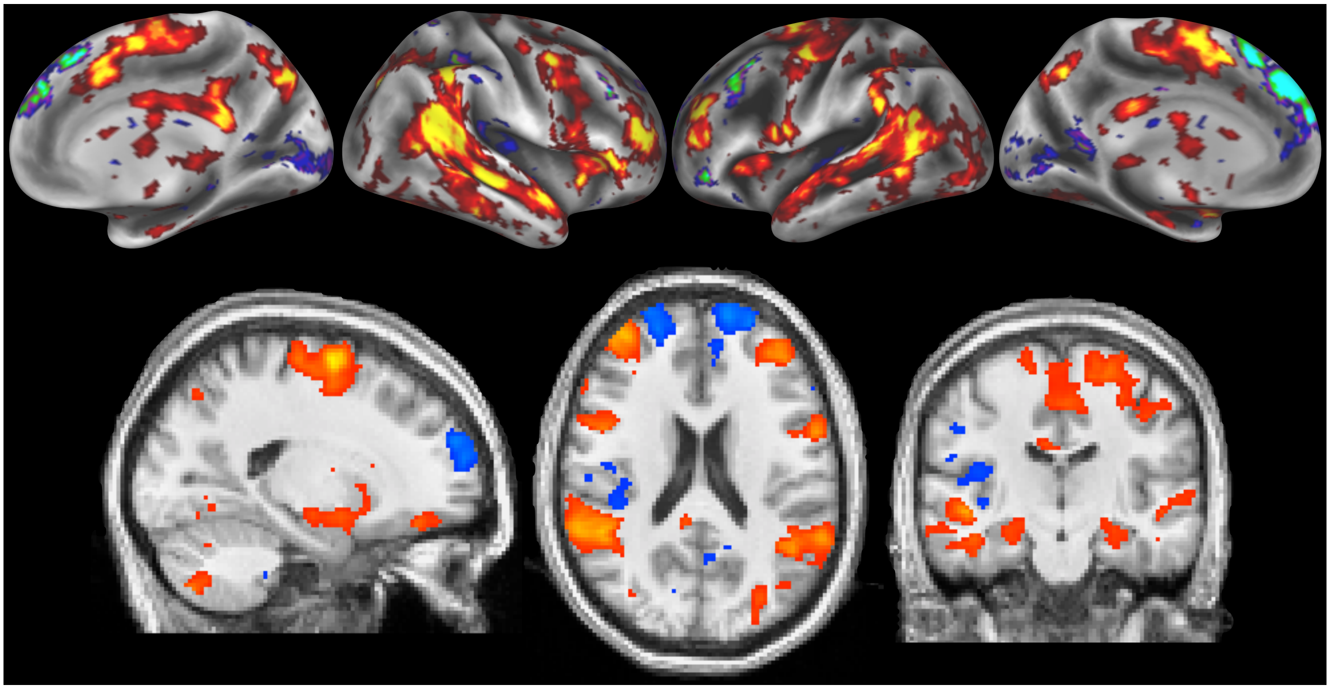

Figure 18: **Spatial map for TFM 16.**

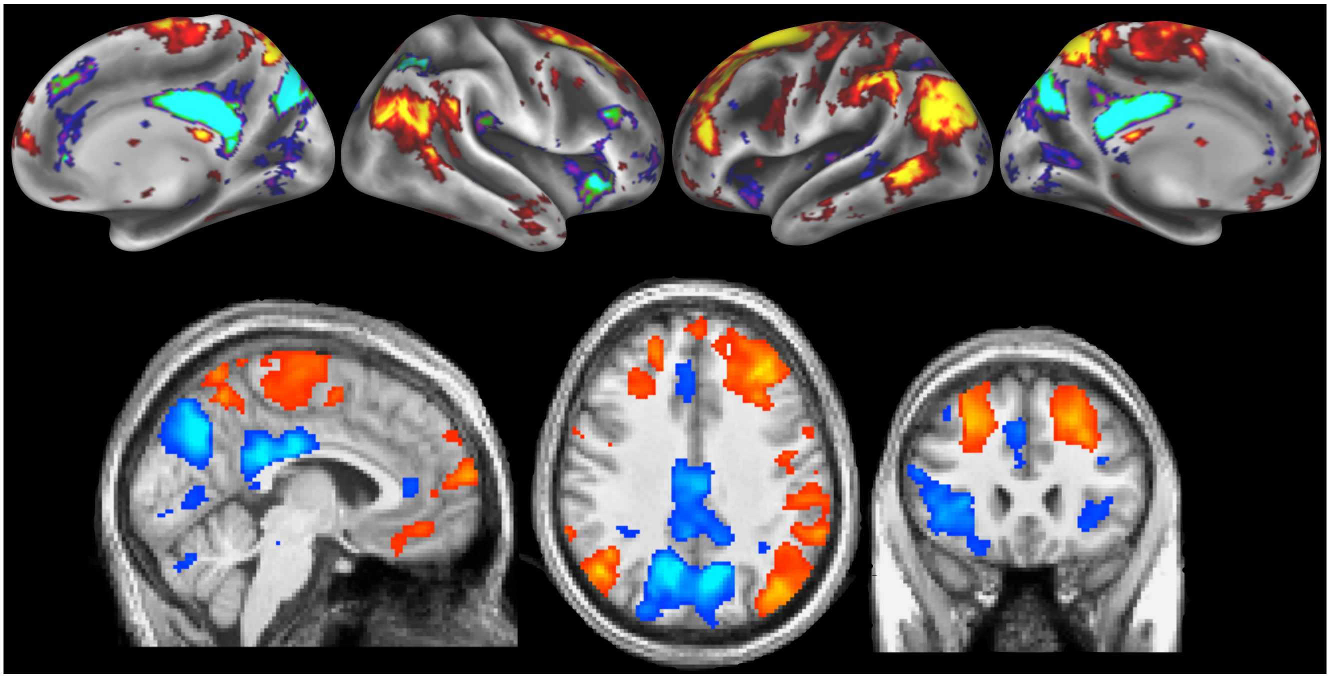

Figure 19: **Spatial map for TFM 17.**

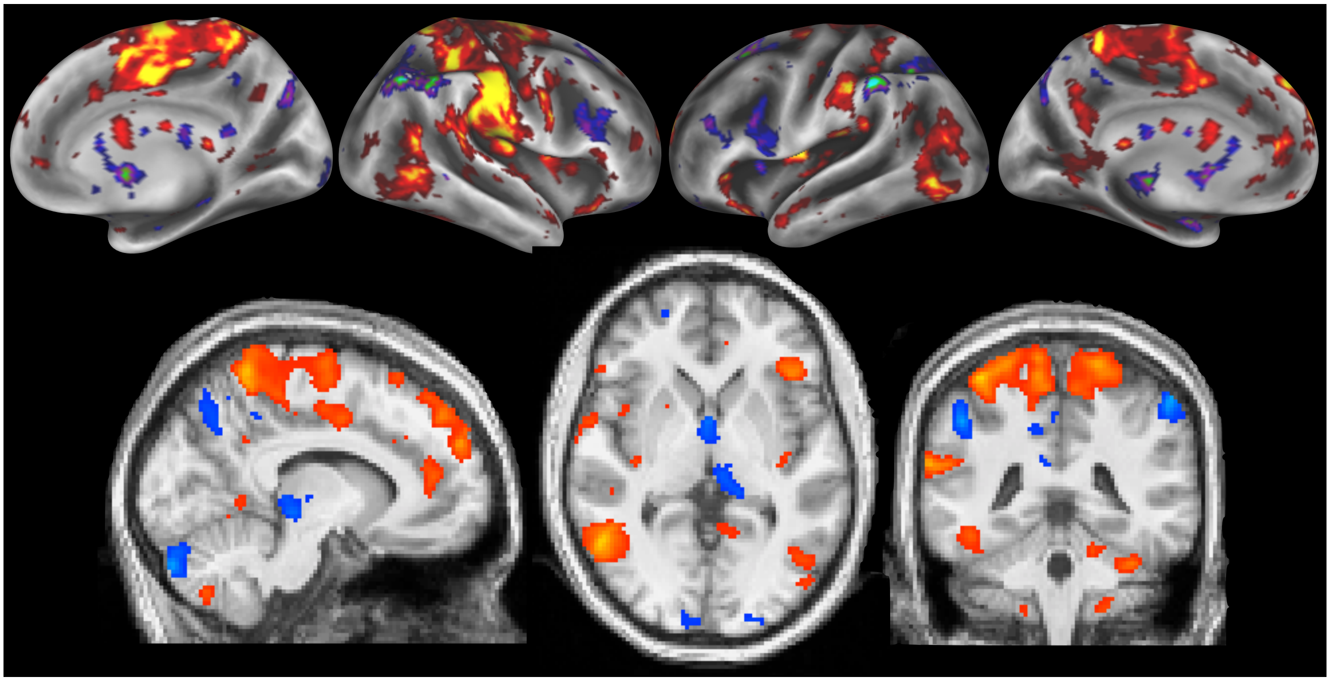

Figure 20: **Spatial map for TFM 18.**

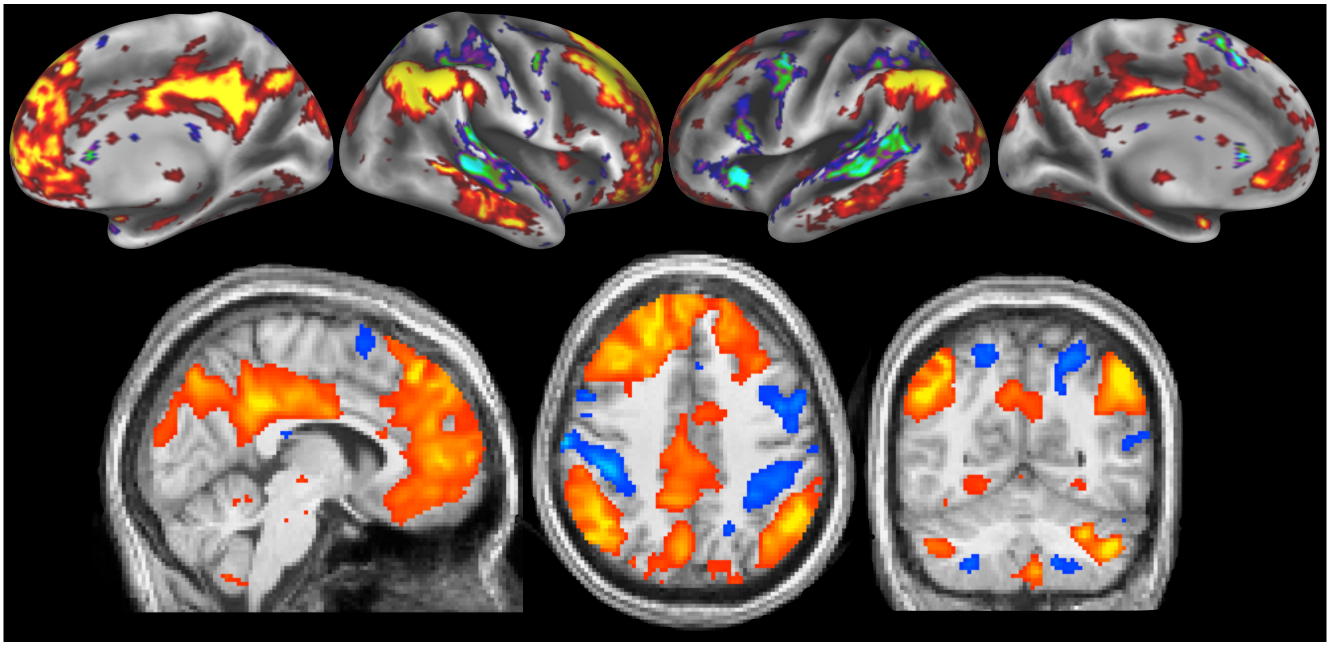

Figure 21: **Spatial map for TFM 19.**

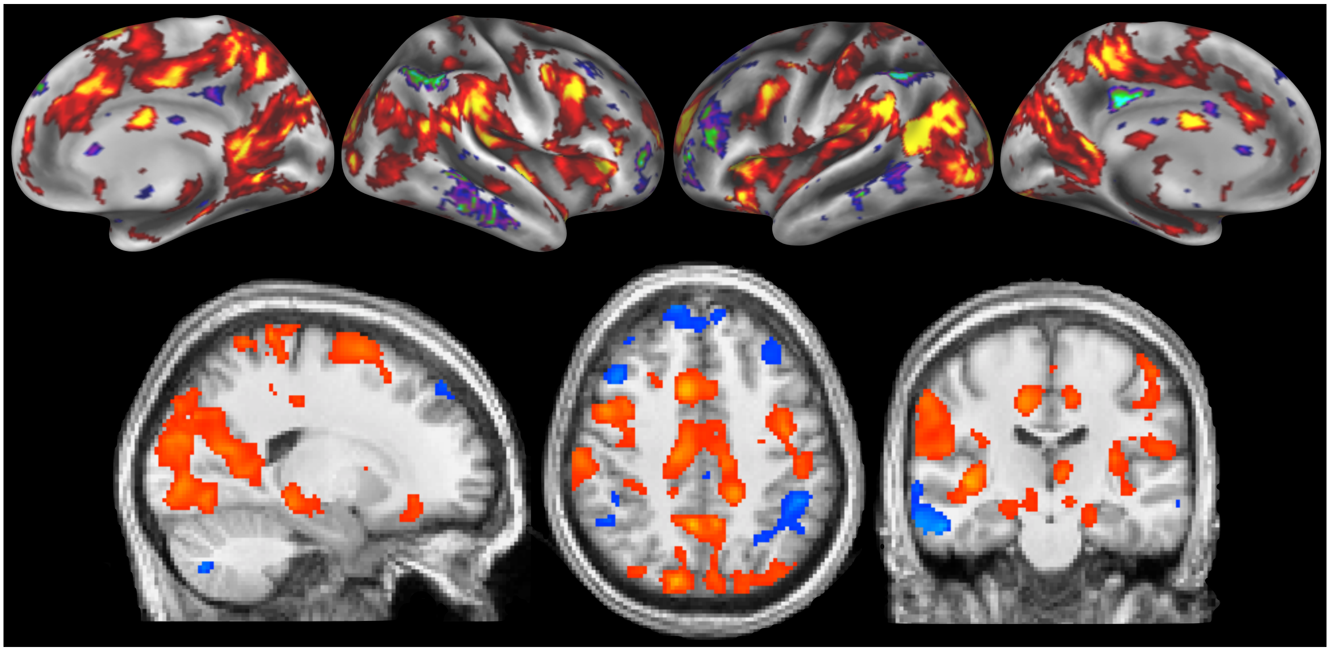

Figure 22: **Spatial map for TFM 20.**

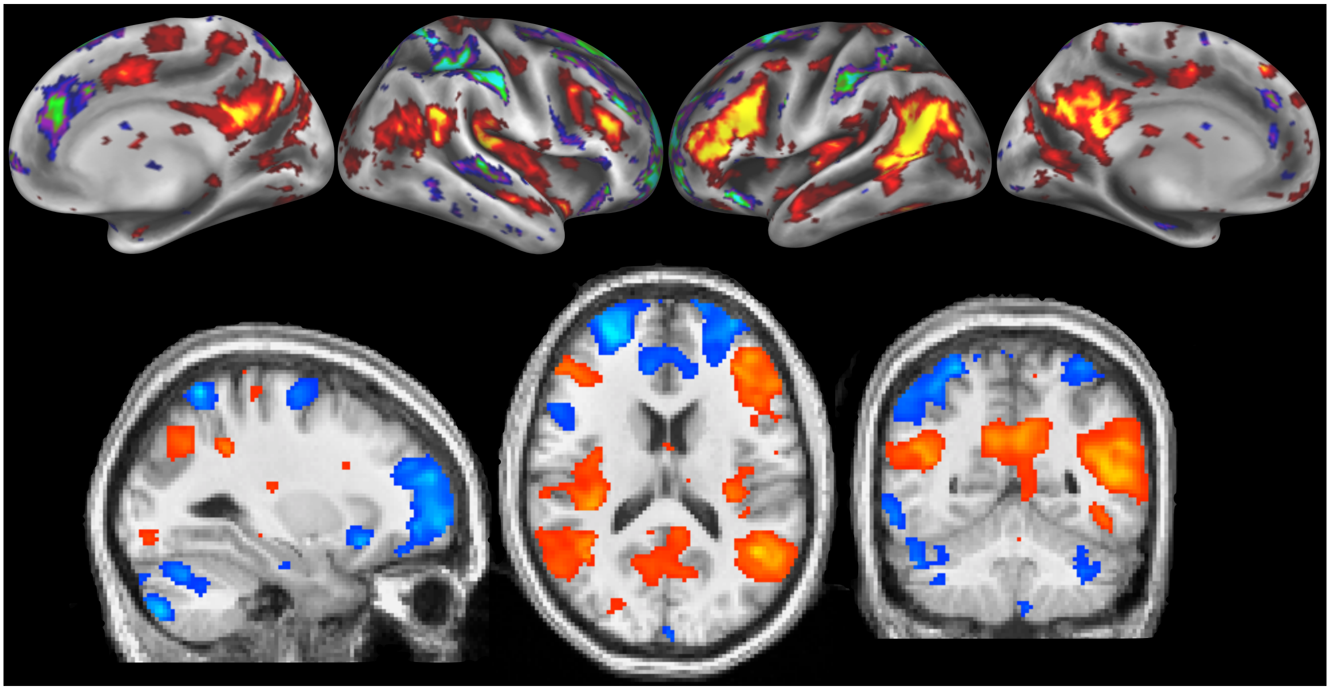

Figure 23: **Spatial map for TFM 21.**

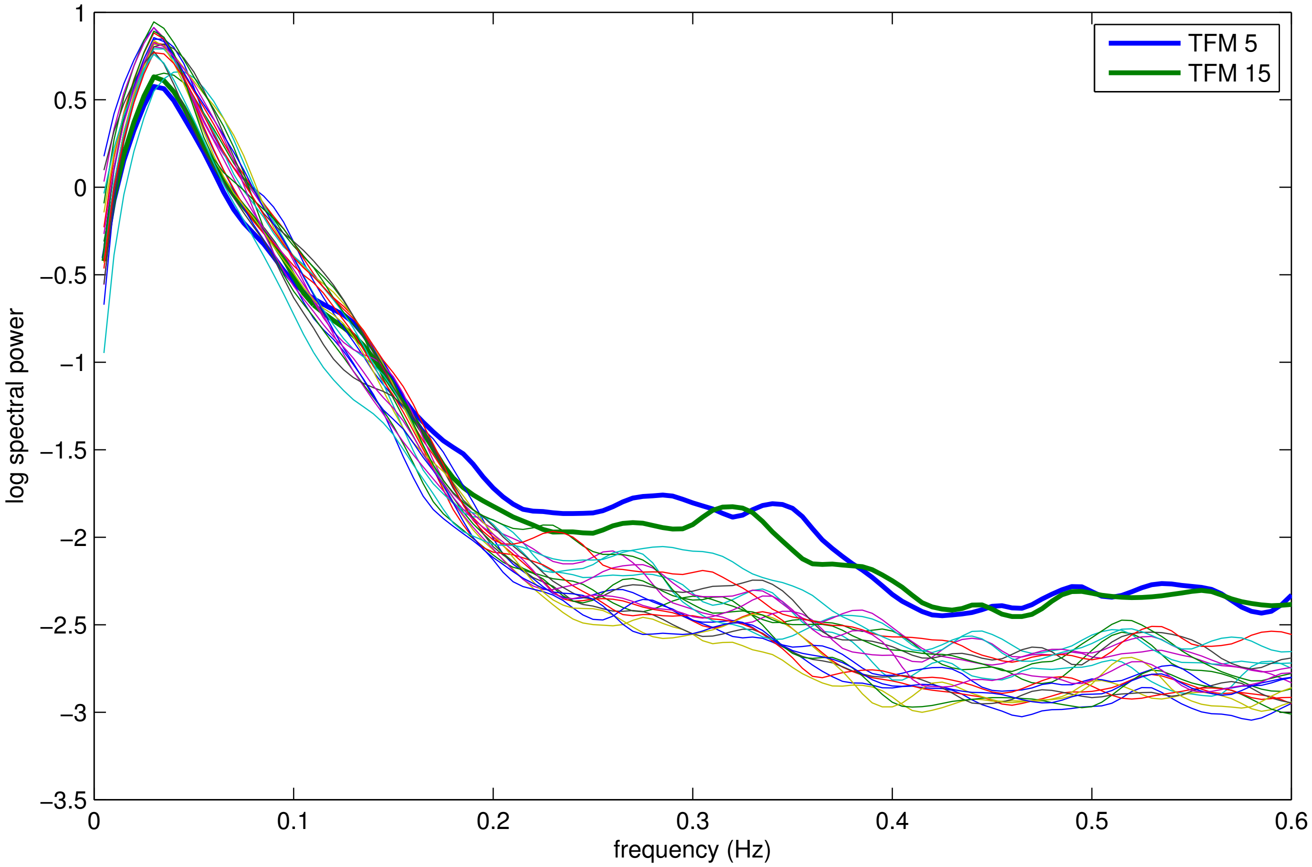

Figure 24: **Power spectra for all TFM timeseries.**

6 Extended correlation maps from visual RSNs

The 3 visual-related RSNs shown in Figure 1 in the main paper are almost entirely confined to visual areas. Because the spatial-ICA components' timecourses are not forced to be orthogonal to each other (in the way that temporal-ICA components are), the maps shown previously do not indicate the full (global) potential spatial involvement of these components - i.e., the extent to which these correlate with other RSNs. Thus, in order to allow a more complete comparison of the RSNs with the TFMs (in particular, TFM 8, which shows a very extended system including higher-level visual and the dorsal visual stream), we re-estimated the 3 visual RSNs, by regressing their 3 timeseries into the original data *without* also including the other 18 RSN timeseries in the multiple regression. This gives their extended areas of correlation in a manner that is more direct (and spatially detailed) than attempting to infer this from the 21x21 RSN timecourse correlation matrix.

The results are shown on the right in Figure 25. While some areas of correlation outside the occipital visual areas are indeed now seen (for example, some primary sensory-motor areas correlating with the foveal visual areas seen in RSN 1), these extended maps are arguably less "clean" (certainly less *cleanly bilateral*) and less strong than was seen with the TFMs (particularly TFM 8).

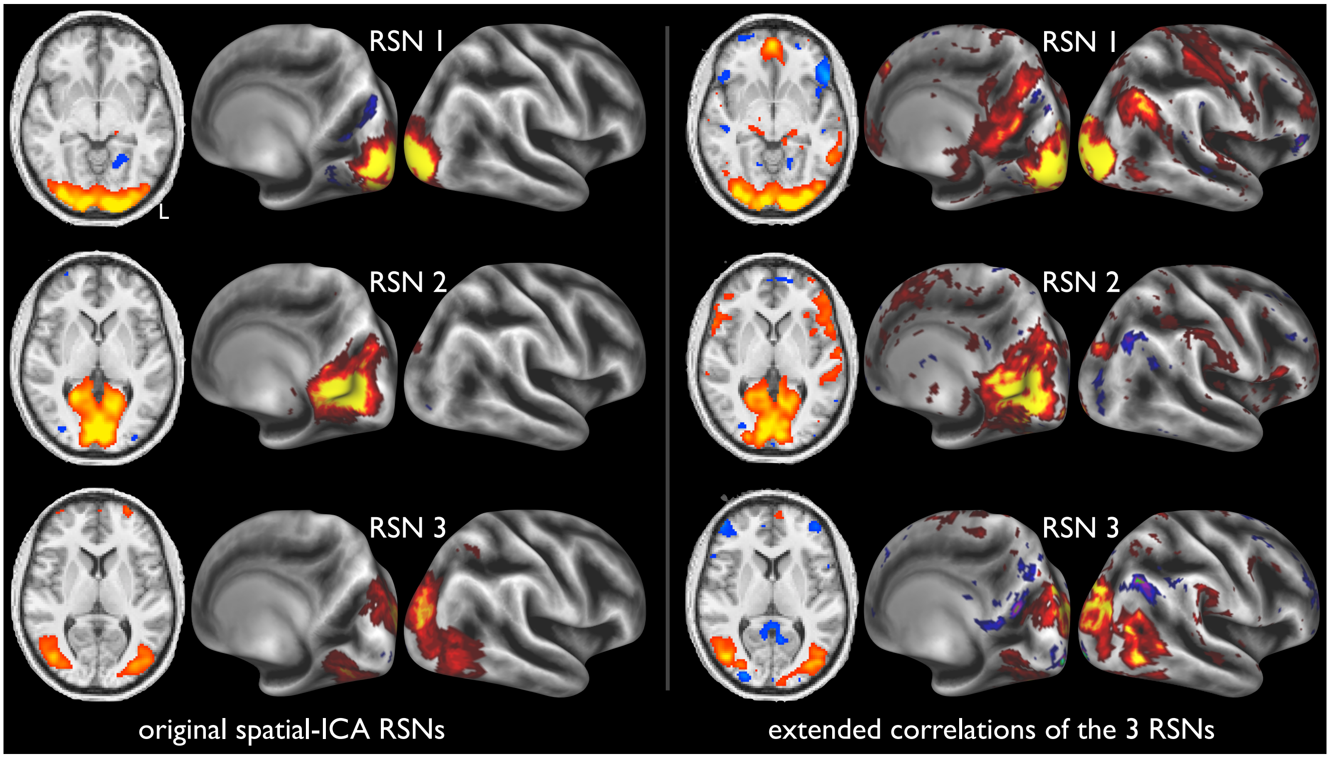

Figure 25: **3 visual components from a 21-dimensional spatial-ICA decomposition of the complete dataset.** On the left is shown the 3 visual-related RSNs from the original 21-dimensional spatial-ICA. On the right is shown the same components, re-derived by regressing the 3 RSNs' timecourses into the original data, *without* the other 18 RSNs' timeseries in the multiple regression, in order to show the extended areas of correlation for the 3 RSNs.

7 Amount of apparent nonstationarity explained by TFMs

Apparent "nonstationarity" (temporal variability) in the correlation between two regions is, in part, due to the temporal variation in activity of overlapping TFMs, despite the network structure *within* each TFM being static (according to our model). As the amplitude of each TFM varies over time (even if its internal network structure is not changing), any pair of nodes that are part of multiple TFMs will see the correlation between these nodes changing over time purely as a result of the independent TFMs' fluctuations.

By taking the outer product of the TFM node-weight vectors and timecourses, i.e., $A_{t(K \times L)} \times S_{t(L \times T)}$, we can reconstruct a version of the data that perfectly follows the TFM model. If we then take pairs of (reconstructed) node timeseries and look at apparent nonstationarity, we can estimate how much of that might be due purely to the overlapping TFMs, rather than other factors such as non-modelled noise or "deeper" nonstationarities (such as TFMs' internal network structure varying over time). We found that the reconstructed set of node timeseries contained 59% of the variance of the original 142 nodes' timeseries.

We measured nonstationarity via sliding window correlation. We used sliding windows (with smooth roll-off at the edges) of 30s. In order to ensure that correlation within the window was well conditioned, we applied a highpass filter to the nodes' timeseries, of approximately 0.15Hz. Without this highpass filter, lower frequencies dominate the correlation, and, because their temporal-period is similar to (or longer than) the timescale of the window, this would otherwise result in very noisy (and hence meaningless) estimates of correlation. We found that our results did not depend strongly on window length. With respect to changing the exact highpass cutoff, as the cutoff was raised or lowered from 0.15Hz, the general reliability/reproducibility of the sliding window correlation became progressively worse (keeping more of the low frequencies damaged the estimate of correlation, as described above, and removing a larger range of higher-frequencies removed too much useful signal). Further, as the cutoff was changed (up or down), similarity of the overall correlation pattern to the full original data correlation matrix also became worse. The fact that the correlation structure (matching closely that found with average correlation applied to the full original data) can be found with such an aggressive highpass cutoff is an interesting result in its own right (and consistent with the results presented in [33]), but further discussion of this is not within the scope of this work.

Hence we estimated correlation between pairs of nodes at multiple window positions and investigated the mean and standard deviation of the correlation across all window positions, for each pair of nodes. The standard deviation was slightly reduced when applied to the original node timeseries compared with the reconstructed timeseries, and the mean correlation increased by a factor of about 2; hence, the (median across node pairs) coefficient of variation (variation across window positions) was 3.2 times lower with the reconstructed data. One might conclude therefore that approximately 25% of the apparent nonstationarity in raw data correlations is purely due to the fact that multiple functional networks (each potentially stationary, but with time-varying amplitude) are spatially overlapping.

8 TFM split-half reproducibility results

In addition to applying temporal-ICA to the full concatenated set of 36 runs, we also split the data into two groups of separate subjects (one group containing 21 runs from two subjects, one of whom had 3 sessions of acquisitions, and the other group containing 15 runs from 3 subjects). Splitting this way, grouping according to subjects, is more conservative and useful than splitting randomly across 10-minute runs. We carried out the temporal-ICA separately for the two groups (after first reducing the data dimensionality to 21 for the two groups together), in order to test the reproducibility of the components (TFMs) found. The reason for applying a whole-group dimensionality reduction is that the aspect of the TFM analysis that is of interest is the *reproducibility of the core ICA unmixing matrix*, i.e., the TFM spatial maps, as estimated through temporal non-Gaussianities in the data.

We used spatial similarity to identify the TFM from each of the split-half analyses that best matched a given TFM from the full dataset (i.e., based on the maximal correlation between the TFM node-weight vectors) This gave a pairing of TFMs between the two halves, with reproducibility then summarised via the correlation coefficient between the two vectors.

In addition, random data of the same dimensions as the true data was generated, with matching covariance structure. This was fed through the same analysis pipeline as the real data (dimensionality reduction, split-half temporal-ICA, matching of paired components and ordering of matched components into decreasing paired similarity). Thus from 1000 different random datasets we built up the null distribution of the split-half reproducibility, as judged by the similarity of the paired split-half results (the TFM node-weight vectors). The reproducibility results are summarised in Figure 26. The split-half reproducibility for the true data lies above the 95th percentile of the null distribution, suggesting strong reproducibility of the TFM node-weight vectors across the two groups of subjects.

As a simpler test of reproducibility, from the same null dataset, we calculated the mean correlation coefficient over all matched pairs of TFMs. The mean from the true data was r=0.49, which corresponds to a significance of $P < 0.01$ (no correction for multiple comparisons needed as we reduced the entire split-half analysis to a single number), as judged from the null distribution of the mean correlation coefficient, again suggesting strong reproducibility between subjects.

This reproducibility of the TFMs means that the temporal-ICA, fed by the 21-dimensional initial PCA dimensionality reduction, is finding the same mixing matrix in the two groups of subjects; the non-Gaussianities in the data are causing a consistent decomposition into temporally-independent functional modes.

Finally, we show the 7 best-matching pairs of TFMs' node-weight vectors, in Figure 27. Many of the plots appear relatively smooth purely because of the re-ordering of the nodes by the initial hierarchical clustering of the correlation matrix after the initial spatial-ICA. This ordering does not affect the temporal-ICA (including its reproducibility); indeed, by definition, this correlation amongst nodes is removed by the PCA stage. Although the full number of nodes is 142, the "real" number of nodes seen by the core temporal-ICA unmixing is only 21 (because of the initial PCA dimensionality reduction), and so a more mathematically meaningful plot of the reproducibility (in the sense that the true degrees of freedom are more directly apparent) is shown in Figure 28; here we see the matched node-weight vectors as seen by the ICA.

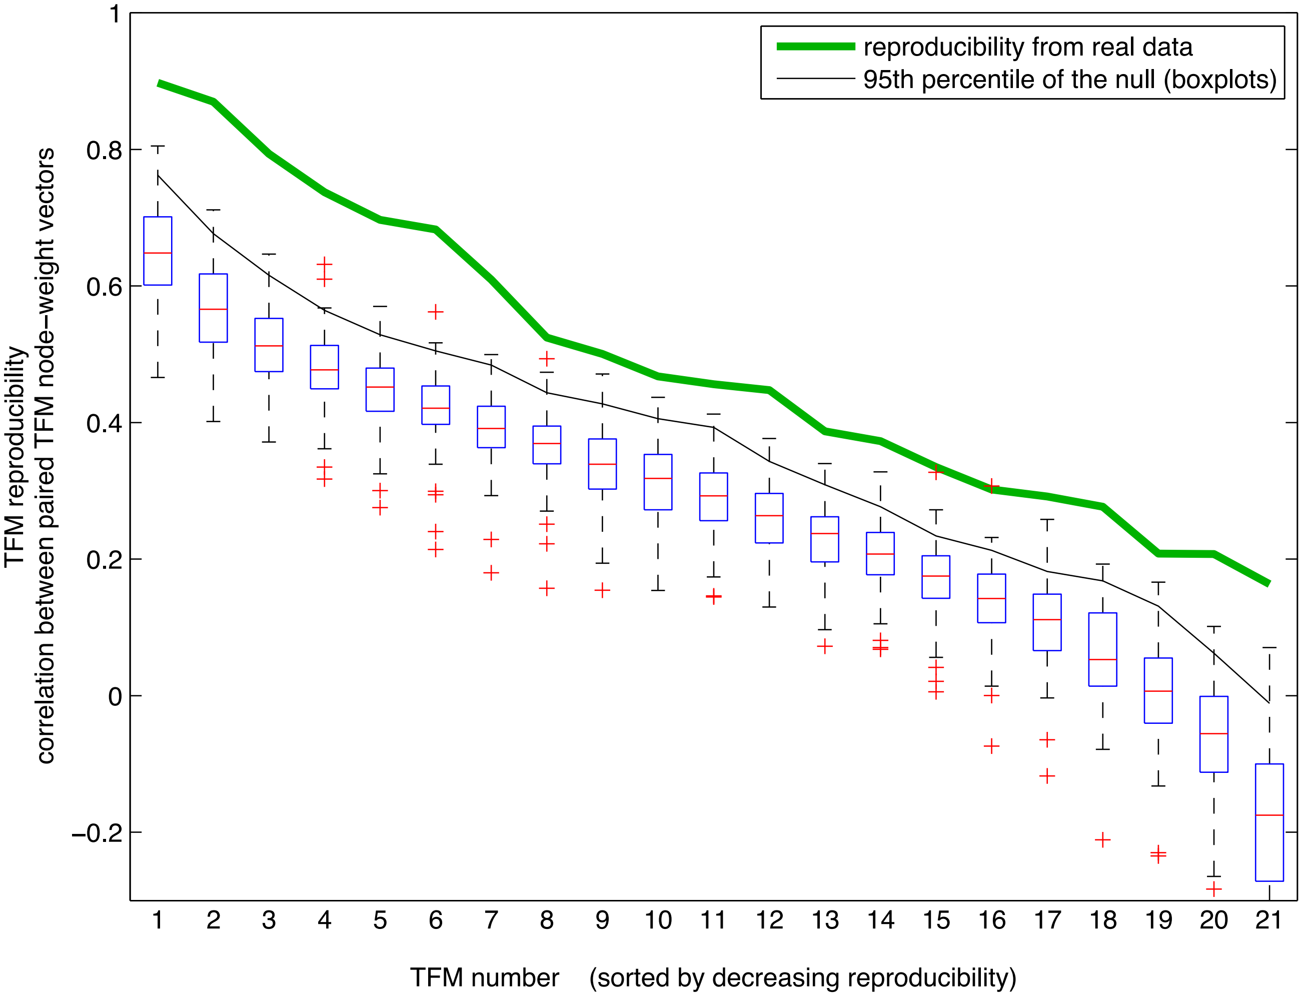

Figure 26: **Split-half reproducibility results.** Random data was fed through the same analysis pipeline as the real data, including split-half ICA and ordering of matched components into decreasing paired similarity. From 1000 different random datasets we built up the null distribution of the split-half reproducibility. The reproducibility for the true data lies above the 95th percentile of the null distribution, suggesting strong reproducibility of the TFM node-weight vectors across the two groups of subjects.

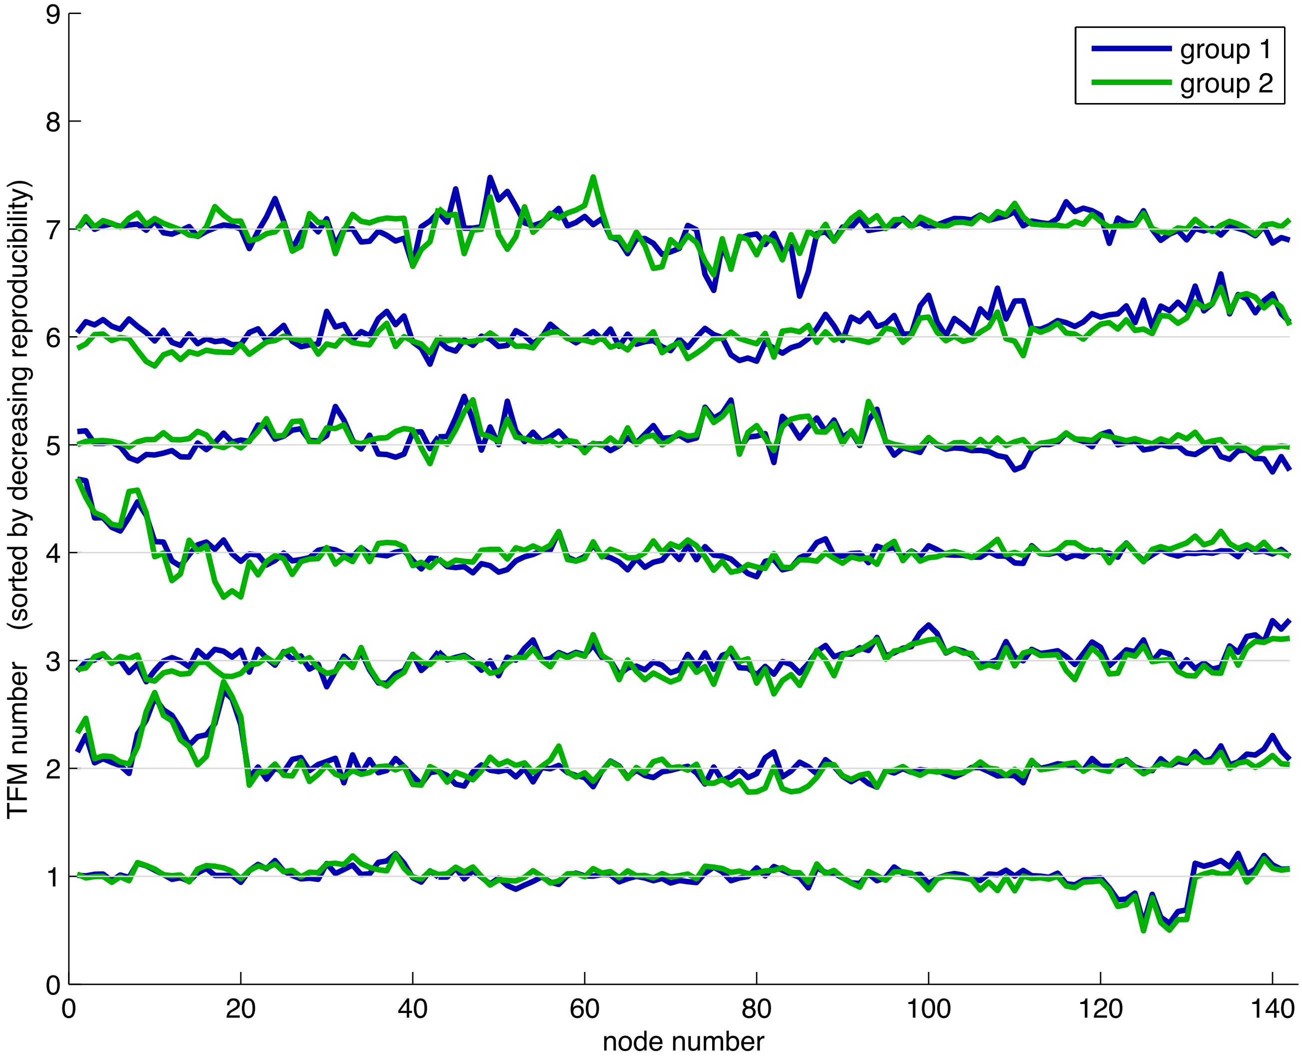

Figure 27: **Split-half reproducibility; 7 best-matched TFM node-weight vectors.** Here we see the full 142-element node-weight vectors as projected out of the 21-dimensional reduced space of the data seen by the temporal-ICA. Many of the plots appear relatively smooth purely because of the re-ordering of the nodes by the initial hierarchical clustering of the correlation matrix after the initial spatial-ICA.

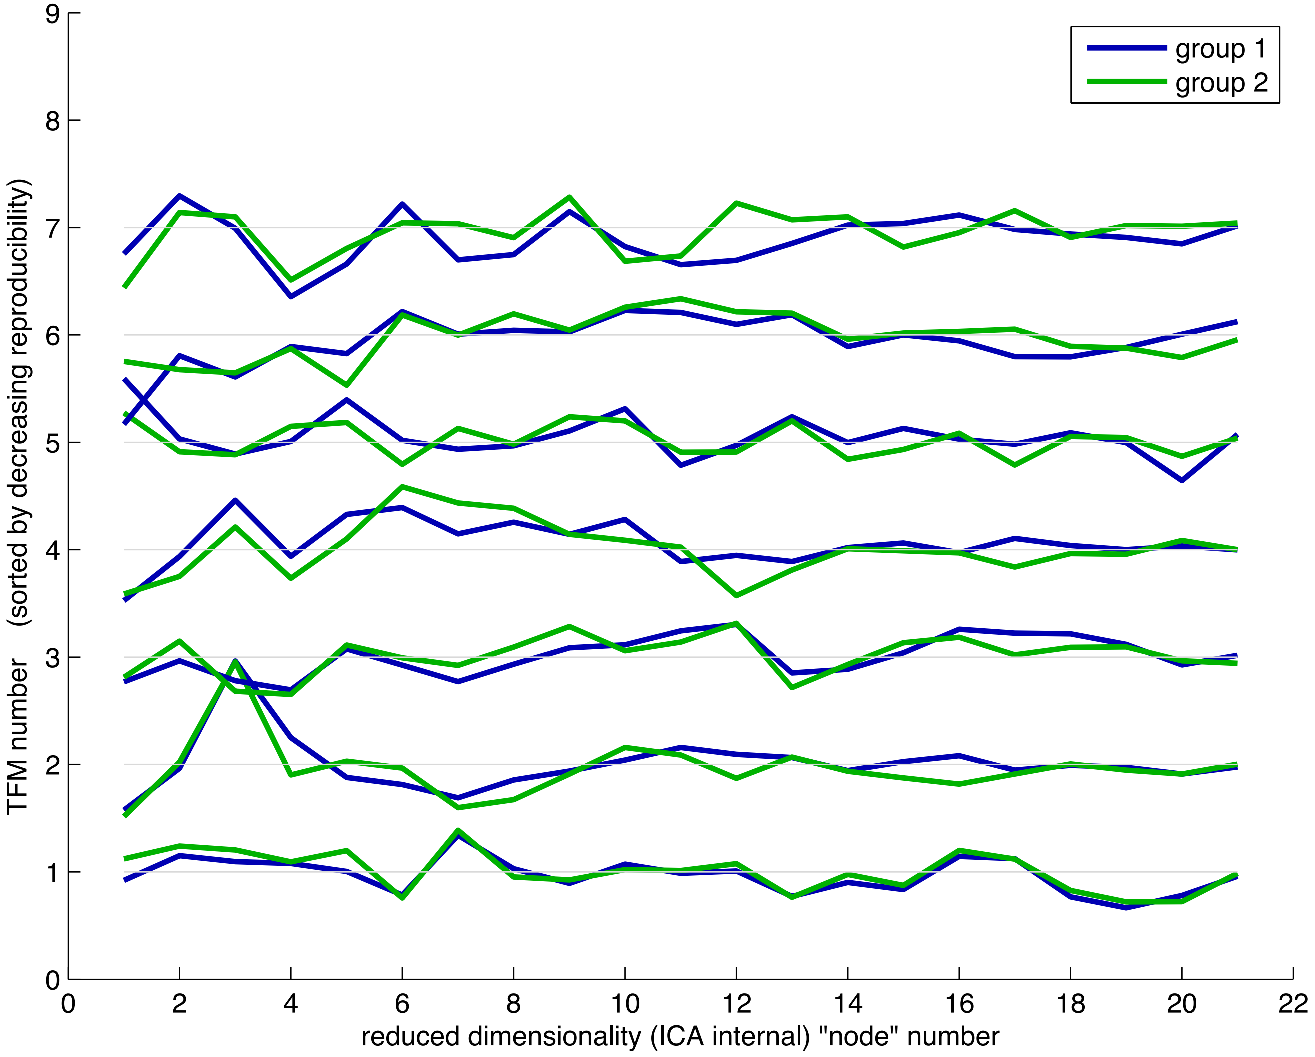

Figure 28: **Split-half reproducibility; 7 best-matched TFM node-weight vectors.** Here we see the 21-dimensional node-weight vectors (the internal mixing matrix) seen by the temporal-ICA.

9 Comparing TFMs against analyses from a separate dataset

In order to gain further confidence in the robustness and reproducibility of the TFMs that we found in the short-TR dataset, we carried out some analyses using a more "standard" dataset, with 36 healthy adult subjects' 6-minute resting data with a TR of 2.04s acquired at 3T on a Siemens Trio, with the subjects instructed to lie with eyes open, with ambient light minimised. This is the dataset used in [34] and [35], i.e., the same dataset that we used previously to compare RSN maps against spatial-ICA components from the BrainMap task-FMRI database. Because these datasets were of shorter duration, lower quality and longer TR than the short-TR dataset used for the majority of our TFM work, the robustness of TFMs estimated from this data set was worse than from the short-TR data. Nevertheless, data quality was sufficient to produce several very encouraging results.

The initial group-wise spatial-ICA was carried out with a dimensionality of 150, from which 25 artefactual components were excluded, leaving 125 functional nodes with associated timeseries to feed into the temporal-ICA.

Our first investigations concentrated on the visual area nodes. On the basis of an initial hierarchical clustering of the raw correlation matrix (similar to that shown in SI Figure 1), and on inspection of the nodes' spatial maps, we chose 20 nodes that covered the visual areas, and fed the timeseries from just these nodes into temporal-ICA, with a dimensionality of 3, to extract the 3 main visual-related TFMs. The results are shown in Figure 29. We used just this subset of the data, and a low number of components, in order to increase robustness of the temporal-ICA, given the limited suitability of this standard dataset for the TFM analysis. Hence the spatial maps from the standard dataset are mostly limited in extent to the visual areas primarily covered by the 20 visual nodes used for the analysis. Nevertheless, the results are compelling, and spatial correlations between matched pairs of maps from the two datasets are significant with $P < 10^{-5}$.

When running a 21-dimensional temporal-ICA on all 125 nodes' timeseries from the standard dataset, we found reasonable reproducibility on a split-half analysis (18 subjects vs. the other, distinct, 18 subjects), with at least 10 TFMs showing significant reproducibility. 3 TFMs of particular note, in terms of their reproducibility, spatial similarity to 3 TFMs found from the short-TR data, and non-similarity to known RSN spatial maps, are shown in Figures 30-32. These show reproducibility in interesting aspects of each of the 3 TFMs: TFM8 (showing anti-correlation between DMN areas and higher-level visual), TFM13 (showing anti-correlation between DMN areas and language areas) and TFM15 (showing a nearly-global signal anti-correlated with insula).

Finally, we regressed the 21 short-TR TFM spatial maps onto the standard data in order to obtain the characterisitc TFM timecourses from the standard dataset without needing to run the full TFM analysis (as discussed above, because of the relative non-suitability of the standard data for TFM analysis). We then estimated the mutual-information between all pairs of TFM timecourses, in order to test the hypothesis that these should be more statistically independent than timecourses derived from spatial-ICA. For the latter, we ran a group-wise 21-dimensional spatial-ICA on the standard dataset, and extracted the RSN timecourses. The results are shown in Figure 33.

The main result is that the TFM timecourses are indeed significantly more statistically independent of each other than the RSN timecourses (i.e., significantly lower mutual information $t = 15.4, p < 10^{-5}$). We also evaluated the effect of random mixing and random rotation (i.e., applying an orthogonal random mixing matrix) of the RSN timeseries; this either increased

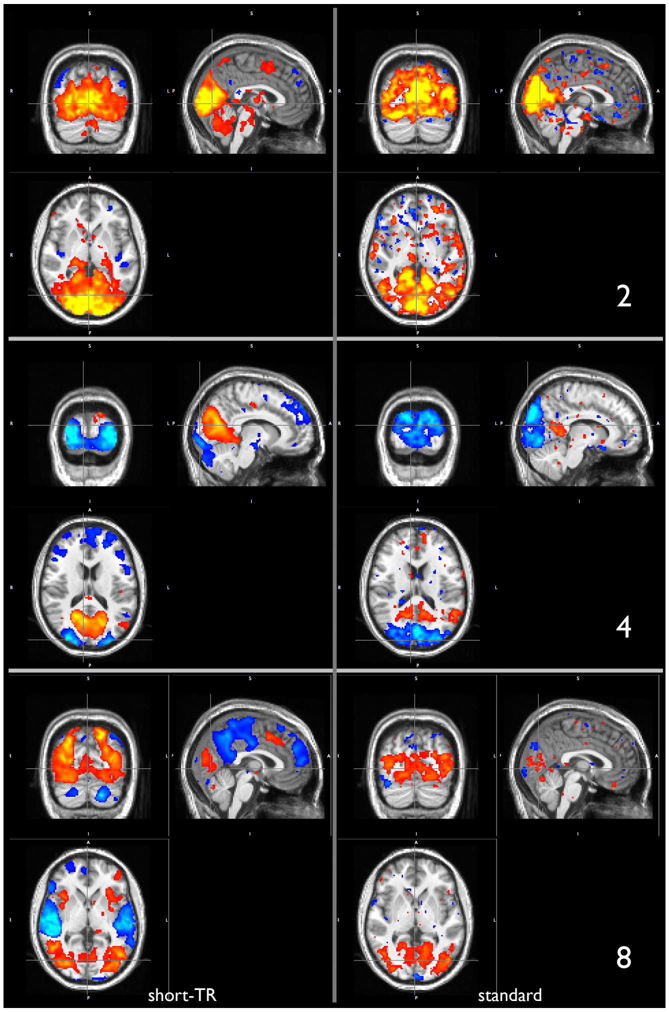

Figure 29: **3 visual TFMs from a TFM analysis of 20 visual nodes (right part of each sub-image) from a standard resting-FMRI dataset, compared against the short-TR visual TFMs 2, 4 and 8 (left part of each sub-image).**

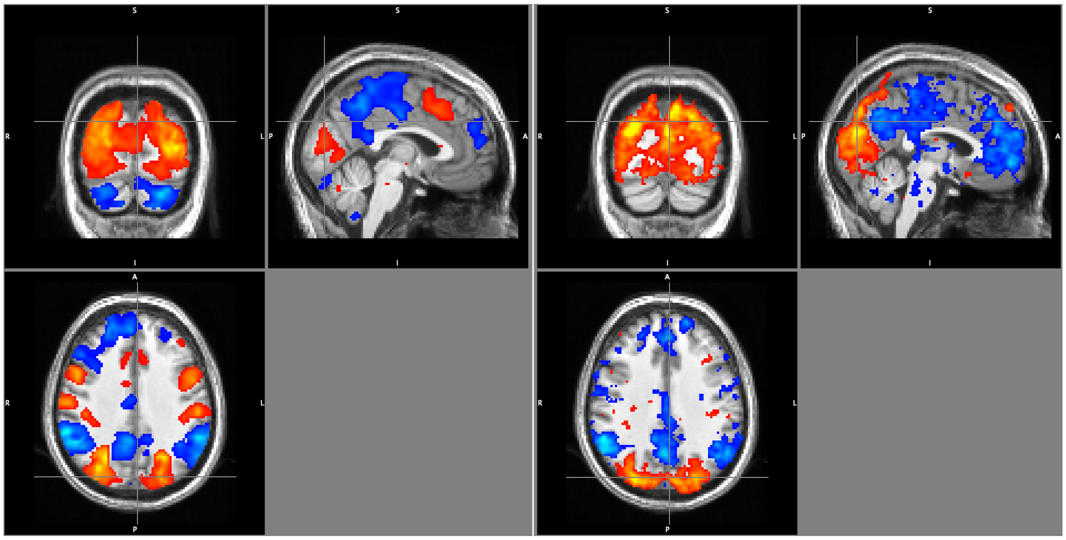

Figure 30: **Left: TFM 8 from the short-TR data. Right: the matching TFM from the standard dataset.**

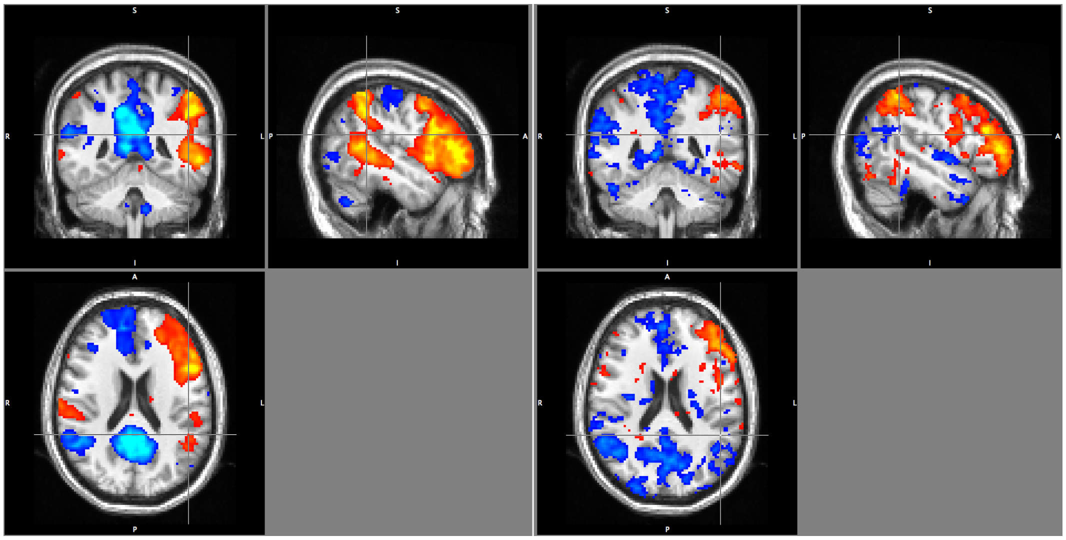

Figure 31: **Left: TFM 13 from the short-TR data. Right: the matching TFM from the standard dataset.**

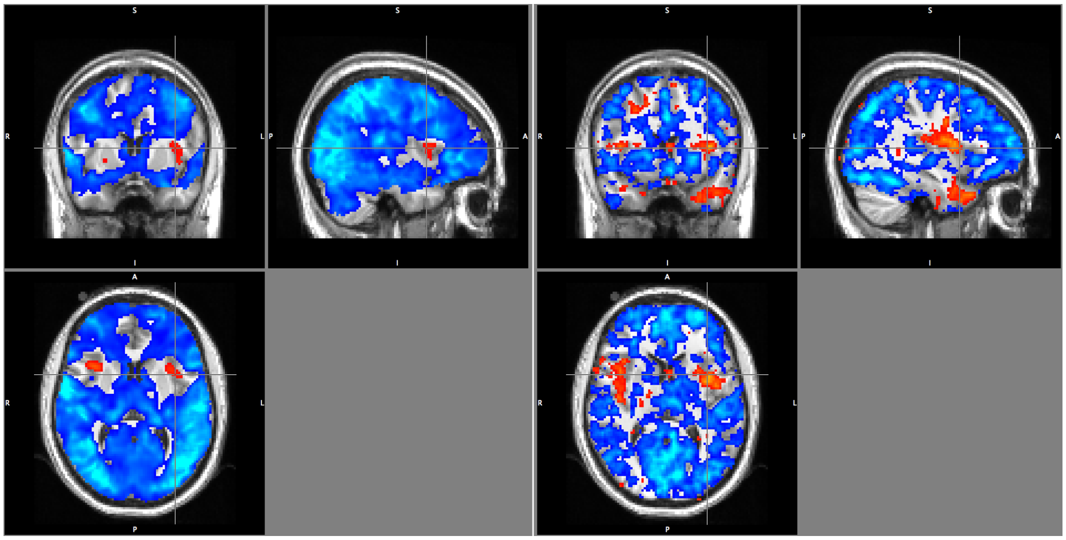

Figure 32: **Left: TFM 15 from the short-TR data. Right: the matching TFM from the standard dataset.**

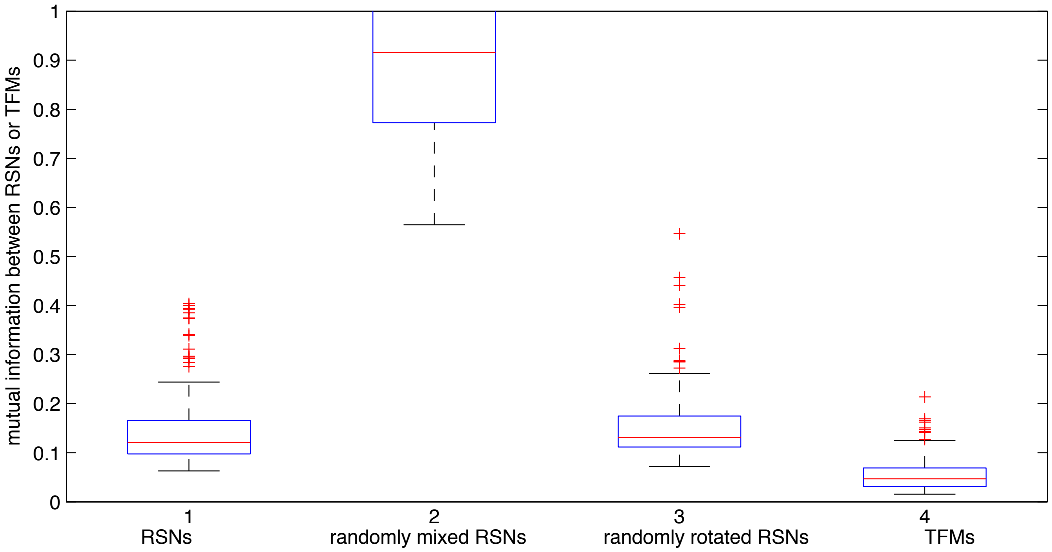

Figure 33: **Mutual information between RSN and TFM timeseries from the standard dataset, with TFMs derived by regressing short-TR TFM spatial maps into the standard-TR FMRI data.** Each boxplot shows the distribution of mutual information values across all possible pairs of RSN or TFM timeseries; the higher the mutual information, the lower the statistical independence. The standard dataset TFM timeseries are significantly more independent of each other, despite being derived via a separate dataset. Random mixing of the RSN timeseries did not result in greater independence of the timeseries.

the mutual information between timeseries, or left it unchanged - further evidence that the TFM timeseries are indeed more independent of each other than the RSN timeseries, in a meaningful way.

10 Consideration of TFM analysis vs. spatial-ICA and temporal-ICA on original data

An implementational question of interest is whether it is necessary to carry out the initial spatial-ICA (to generate a large number of nodes), given that the dimensionality is then reduced to a comparatively small number for the final temporal-ICA. Mathematically, feeding the original data directly into temporal-ICA (with the dimensionality reduced by an initial PCA to the number of temporal-ICA components to be found) is very similar to what we have done. However, being able to identify artefactual components from the high-dimensional spatial-ICA, and regress these out of the non-artefactual components, is one practical advantage of the approach that we have taken.[4] Nevertheless, we have also carried out a direct 21-dimensional temporal-ICA on the original group-wise (concatenated) data, and indeed achieved similar results to the TFMs shown above.

A second difference between the two approaches is that the PCA dimensionality reduction will act differently. If PCA is applied to the original data, the number of voxels acting with a similar timecourse will affect the reduction (the more voxels that are synchronised with respect to a particular underlying process, the more likely that process is to be kept by the dimensionality reduction). However, if PCA is applied to the set of node timeseries, as we have done in our main TFM analyses, the number of voxels that effectively fed into a given node is not considered; all nodes potentially carry equal weight with respect to the dimensionality reduction. This may be an advantage (e.g., in terms of being able to identify important, smaller, nodes) or may be a disadvantage (e.g., in terms of making less robust use of the available signal in the data), and is something we will investigate further in the future.

Finally, it is interesting to consider the recent results reported in [36], where *temporal-ICA* was applied to MEG timeseries (more specifically, source-reconstructed Hilbert envelope timeseries), and at least 8 of the resulting spatial maps were found to be very similar to the RSN spatial maps derived from a low-dimensional *spatial-ICA* from resting FMRI data (in fact, the same "standard" dataset that we have used above). The question arises: if the arguments and results that we have presented in this paper are correct, surely the MEG temporal-ICA results should not look so similar to the spatial-ICA results from resting FMRI (in the sense that the TFM spatial maps presented above are in many aspects quite distinct from known RSN spatial maps). We hypothesised that this might be due to the much greater spatial smoothness present in the MEG data, and possibly (related to this) the lower intrinsic dimensionality of the less spatially detailed data. We carried out a basic test of this hypothesis by applying a very large amount of smoothing (Gaussian smoothing of full-width half-maximum 20mm) to our 4D datasets before re-running groupwise temporal-ICA directly on these 4D datasets concatenated temporally across all runs/subjects (working in 4mm resolution MNI152 standard space for computational practicality). The ICA spatial maps that resulted from this of course had considerably less spatial detail in them, and, judging qualitatively, did indeed look much more like low-dimensional RSN spatial patterns than do the TFM results presented above. However these are just initial, qualitatively judged, results, and these issues deserve further study in the future.

---

[4] If the group-level ICA is to do a good job of estimating artefacts as well as signal, it is likely that we would want to estimate a large number of components (to best separate out all possible artefacts and signals). This argues for the use in the first instance of *spatial*-ICA (rather than moving straight into the temporal-ICA); because (even with our accelerated FMRI) it is still the case that spatial-ICA will be more robust than temporal, we can therefore do a better job of finding (and then regressing out) the (potentially large number of) artefacts at the group level using spatial-ICA *before* moving on to the final temporal-ICA decomposition.

References

[1] D. J. Larkman, J. V. Hajnal, A. H. Herlihy, G. A. Coutts, I. R. Young, and G. Ehnholm. Use of multicoil arrays for separation of signal from multiple slices simultaneously excited. *J Magn Reson Imaging*, 13(2):313–7, 2001.

[2] S. Moeller, E. Yacoub, C. A. Olman, E. Auerbach, J. Strupp, N. Harel, and K. Ugurbil. Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magn Reson Med*, 63(5):1144–53, 2010.

[3] D.A. Feinberg, S. Moeller, S.M. Smith, E. Auerbach, S. Ramanna, M.F. Glasser, K.L. Miller, K. Ugurbil, and E. Yacoub. Multiplexed echo planar imaging for sub-second whole brain FMRI and fast diffusion imaging. *PLoS ONE*, 5(12):e15710, 2010.

[4] P. Mansfield. Multi-planar image formation using NMR spin echoes. *J.phys. C: Solid State Phys.*, 10:L55–L58, 1977.

[5] F. A. Breuer, M. Blaimer, R. M. Heidemann, M. F. Mueller, M. A. Griswold, and P. M. Jakob. Controlled aliasing in parallel imaging results in higher acceleration (CAIPIRINHA) for multi-slice imaging. *Magn Reson Med*, 53(3):684–91, 2005.

[6] R.G. Nunes, J.V. Hajnal, X. Golay, and D.J. Larkman. Simultaneous slice excitation and reconstruction for single shot EPI. In *ISMRM 14th Annual Meeting*, page 293, 2006.

[7] J. P. Mugler III and J. R. Brookeman. Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP RAGE). *Magn Reson Med*, 15(1):152–7, 1990.

[8] S.M. Smith, M. Jenkinson, M.W. Woolrich, C.F. Beckmann, T.E.J. Behrens, H. Johansen-Berg, P.R. Bannister, M. De Luca, I. Drobnjak, D.E. Flitney, R. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J.M. Brady, and P.M. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23(S1):208–219, 2004.

[9] M.W. Woolrich, S. Jbabdi, B. Patenaude, M. Chappell, S. Makni, T. Behrens, C. Beckmann, M. Jenkinson, and S.M. Smith. Bayesian analysis of neuroimaging data in FSL. *NeuroImage*, 45:S173–186, 2009.

[10] A.M. Dale, B. Fischl, and M.I. Sereno. Cortical surface-based analysis I: Segmentation and surface reconstruction. *NeuroImage*, 9:179–194, 1999.

[11] B. Fischl, M.I. Sereno, and A.M. Dale. Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage*, 9:195–207, 1999.

[12] A. Hyvärinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, 10(3):626–634, 1999.

[13] D.C. Van Essen, J. Dickson, J. Harwell, D. Hanlon, C.H. Anderson, and H.A. Drury. An integrated software system for surface-based analyses of cerebral cortex. *Journal of American Medical Informatics Association*, 8(5):443–459, 2001.

[14] M. Jenkinson and S.M. Smith. A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2):143–156, 2001.

[15] C.F. Beckmann and S.M. Smith. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans. on Medical Imaging*, 23(2):137–152, 2004.

[16] D.N. Greve and B. Fischl. Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, 48:63–72, 2009.

[17] Y. Zhang, M. Brady, and S. Smith. Segmentation of brain MR images through a hidden Markov random field model and the expectation maximization algorithm. *IEEE Trans. on Medical Imaging*, 20(1):45–57, 2001.

[18] D.S. Marcus, J. Harwell, T. Olsen, M. Hodge, M.F. Glasser, F. Prior, M. Jenkinson, T. Laumann, S.W. Curtiss, and D.C. Van Essen. Informatics and data mining tools and strategies for the Human Connectome Project. *Frontiers in Neuroinformatics*, 5(4), 2011.

[19] J. Andersson, S. Smith, and M. Jenkinson. Non-linear optimisation. Internal Technical Report TR07JA1, Oxford Centre for Functional Magnetic Resonance Imaging of the Brain, Department of Clinical Neurology, Oxford University, Oxford, UK, 2007. Available at `www.fmrib.ox.ac.uk/analysis/techrep` for downloading.

[20] J. Andersson, M. Jenkinson, and S. Smith. Non-linear registration aka spatial normalisation. Internal Technical Report TR07JA2, Oxford Centre for Functional Magnetic Resonance Imaging of the Brain, Department of Clinical Neurology, Oxford University, Oxford, UK, 2007. Available at `www.fmrib.ox.ac.uk/analysis/techrep` for downloading.

[21] B. Fischl, M.I. Sereno, R.B.H. Tootell, and A.M. Dale. High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping*, 8(4):272–284, 1999.

[22] Shohei Shimizu, Patrik O. Hoyer, Aapo Hyvärinen, and Antti Kerminen. A linear non-Gaussian acyclic model for causal discovery. *J. Mach. Learn. Res.*, 7:2003–2030, 2006.

[23] G. Marrelec, A. Krainik, H. Duffau, M. Pélégrini-Issac, S. Lehéricy, J. Doyon, and H. Benali. Partial correlation for functional brain interactivity investigation in functional MRI. *NeuroImage*, 32:228–237, 2006.

[24] S.M. Smith, K.L. Miller, G. Salimi-Khorshidi, M. Webster, C.F. Beckmann, T.E. Nichols, J.D. Ramsey, and M.W. Woolrich. Network modelling methods for FMRI. *NeuroImage*, 54:875–891, 2011.

[25] A. Shmuel, E. Yacoub, J. Pfeuffer, P.F. Van de Moortele, G. Adriany, X. Hu, and K. Ugurbil. Sustained negative BOLD, blood flow and oxygen consumption response and its coupling to the positive response in the human brain. *Neuron*, 36(6):1195–1210, 2002.

[26] A. Shmuel, M. Augath, A. Oeltermann, and N.K. Logothetis. Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nature Neuroscience*, 9(4):569–577, 2006.

[27] R. Leech, S. Kamourieh, C.F. Beckmann, and D.J. Sharp. Fractionating the default mode network: Distinct contributions of the ventral and dorsal posterior cingulate cortex to cognitive control. *The Journal of Neuroscience*, 31(9):3217–3224, 2011.

[28] J.R. Binder, R.H. Desai, W.W. Graves, and L.L. Conant. Where is the semantic system? a critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, 19(12):2767–2796, 2009.

[29] S.M. Oppenheimer, A. Gelb, J.P. Girvin, and V.C. Hachinski. Cardiovascular effects of human insular cortex stimulation. *Neurology*, 42:1727–32, 1992.

[30] E.R. Locatelli, J.P. Varghese, A. Shuaib, and S.J. Potolicchio. Cardiac asystole and bradycardia as a manifestation of left temporal lobe complex partial seizure. *Annals of internal medicine*, 130(7):581–583, 1999.

[31] K. Lamb, K. Gallagher, R. McColl, D. Mathews, R. Querry, and J.W. Williamson. Exercise-induced decrease in insular cortex rCBF during postexercise hypotension. *Medicine & Science in Sports & Exercise*, 39(4):672–679, 2007.

[32] H.D. Critchley. Neural mechanisms of autonomic, affective, and cognitive integration. *The Journal of Comparative Neurology*, 493(1):154–166, 2005.

[33] R.K. Niazy, J. Xie, K. Miller, C.F. Beckmann, and S.M. Smith. Spectral characteristics of resting state networks. In E.J.W. Van Someren, editor, *Progress in Brain Research, Vol. 193*, pages 259–276. Elsevier, 2011.

[34] N. Filippini, B.J. MacIntosh, M.G. Hough, G.M. Goodwin, G.B. Frisoni, S.M. Smith, P.M. Matthews, C.F. Beckmann, and C.E. Mackay. Distinct patterns of brain activity in young carriers of the APOE-e4 allele. *Proc Natl Acad Sci USA (PNAS)*, 106:7209–7214, 2009.

[35] S.M. Smith, P.T. Fox, K.L. Miller, D.C. Glahn, P.M. Fox, C.E. Mackay, N. Filippini, K.E. Watkins, R. Toro, A.R. Laird, and C.F. Beckmann. Correspondence of the brain's functional architecture during activation and rest. *Proc Natl Acad Sci USA (PNAS)*, 106(31):13040–13045, 2009.

[36] M.J. Brookes, M.W. Woolrich, H. Luckhoo, D. Price, J.R. Hale, M.C. Stephenson, G.R. Barnes, S.M. Smith, and P.G. Morris. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proc Natl Acad Sci USA (PNAS)*, 108(40):16783–16788, 2011.