

Temporally-independent functional modes of spontaneous brain activity

Stephen M. Smith^{a,1}, Karla L. Miller^a, Steen Moeller^b, Junqian Xu^b, Edward J. Auerbach^b, Mark W. Woolrich^{a,c}, Christian F. Beckmann^{a,d,e}, Mark Jenkinson^a, Jesper Andersson^a, Matthew F. Glasser^f, David C. Van Essen^f, David A. Feinberg^{g,h}, Essa S. Yacoub^b, and Kamil Ugurbil^b

^aOxford Centre for Functional Magnetic Resonance Imaging of the Brain, University of Oxford, Oxford OX3 9DU, United Kingdom; ^bCenter for Magnetic Resonance Research, University of Minnesota, Minneapolis, MN 55455; ^cOxford Centre for Human Brain Activity, University of Oxford, Oxford OX3 7JX, United Kingdom; ^dDonders Institute for Brain, Cognition and Behaviour, Radboud University, 6500 HC, Nijmegen, The Netherlands; ^eMIRA Institute for Biomedical Technology and Technical Medicine, University of Twente, 7500 AE, Enschede, The Netherlands; ^fDepartment of Anatomy and Neurobiology, School of Medicine, Washington University in St. Louis, St. Louis, MO 63110; ^gAdvanced MRI Technologies, Sebastopol, CA 95472; and ^hHelen Wills Institute for Neuroscience, University of California, Berkeley, CA 94720

Edited by Marcus E. Raichle, Washington University in St. Louis, St. Louis, MO, and approved January 17, 2012 (received for review December 23, 2011)

Resting-state functional magnetic resonance imaging has become a powerful tool for the study of functional networks in the brain. Even “at rest,” the brain’s different functional networks spontaneously fluctuate in their activity level; each network’s spatial extent can therefore be mapped by finding temporal correlations between its different subregions. Current correlation-based approaches measure the average functional connectivity between regions, but this average is less meaningful for regions that are part of multiple networks; one ideally wants a network model that explicitly allows overlap, for example, allowing a region’s activity pattern to reflect one network’s activity some of the time, and another network’s activity at other times. However, even those approaches that do allow overlap have often maximized mutual spatial independence, which may be suboptimal if distinct networks have significant overlap. In this work, we identify functionally distinct networks by virtue of their temporal independence, taking advantage of the additional temporal richness available via improvements in functional magnetic resonance imaging sampling rate. We identify multiple “temporal functional modes,” including several that subdivide the default-mode network (and the regions anticorrelated with it) into several functionally distinct, spatially overlapping, networks, each with its own pattern of correlations and anticorrelations. These functionally distinct modes of spontaneous brain activity are, in general, quite different from resting-state networks previously reported, and may have greater biological interpretability.

Functional connectivity in the brain can be observed using resting-state functional magnetic resonance imaging (fMRI), because the spontaneous temporal fluctuations from functionally related regions show temporal correlation with each other (1). However, brain function is mediated by many functionally distinct networks. These networks may overlap each other, either because a given region contains distinct functional units that cannot be separated with the limited spatial resolution of the data, or because a region truly is a part of more than one distinct functional network. If two regions participate in multiple functional networks, their apparent temporal correlation will reflect the combined contribution from all networks, obscuring the true underlying functional organization. The correlation will not necessarily be meaningful, being some unknown combination of correlations caused by various distinct processes. In this paper, we describe an approach that allows for overlap between functional networks, but differentiates the networks from each other on the basis of their temporal, rather than spatial, independence. The rationale is that in a resting fMRI dataset of sufficient duration, functionally distinct networks should be largely temporally distinct, even if spatially overlapping.

Two main modeling approaches have been used to estimate functional connectivity from resting fMRI data. The simplest (and most widely used) approach is seed-based correlation (1), in which temporal correlation is treated as a single fixed quantity to

be estimated over the whole experiment’s duration, although some recent work has examined how the correlation varies over time (2). A popular alternative approach is to carry out a low-dimensional spatial independent component analysis (ICA) (3) to define “resting-state networks” (RSNs). Spatial ICA decomposes the data into components that are maximally spatially independent of each other. With a low-dimensional ICA, each component contains (in general) several noncontiguous regions that all fluctuate together with similar spontaneous time courses; different components have different time courses. A further variant on these approaches is to parcellate the brain into a large number of functionally distinct regions (e.g., based on an anatomical atlas, or via a high-dimensional spatial ICA). The parcels (also referred to as network “nodes”) are then clustered into extended functional networks on the basis of the temporal correlations between their associated time courses.

Soon after spatial ICA was first proposed for task fMRI (3), temporal ICA (optimizing for temporal independence between components) was also applied to task fMRI (4). It was shown that both spatial and temporal ICA could identify a small number of strongly distinct task activation and artifact processes (5). However, it quickly became apparent that (for fMRI) ICA performs much more robustly in the spatial than the temporal domain, partly because ICA requires a large number of samples to function well, and in fMRI there are orders of magnitude more voxels than time points. Also, the underlying neural processes in the data may well be more non-Gaussian in space than in time (particularly for resting-state data), adding to the greater robustness of spatial ICA. With respect to the separation of activation from artifacts, and of spatially distinct activations from each other, spatial independence may be the most appropriate model, for example, where stimulus-induced head motion would be temporally correlated to the signal of interest (3). For these reasons, nearly all applications of ICA to fMRI (including resting-state fMRI) have to date used spatial ICA.

However, for finding functionally independent, potentially spatially overlapping, functional networks, temporal ICA may provide a better (or at least complementary and informative) approach. Indeed, a relevant point was made in a commentary on the original use of spatial ICA for fMRI: “On the negative side, the very nature

Author contributions: S.M.S., K.L.M., S.M., J.X., E.J.A., M.W.W., C.F.B., M.F.G., D.C.V.E., D.A.F., E.S.Y., and K.U. designed research; S.M.S., K.L.M., S.M., J.X., E.J.A., M.W.W., C.F.B., M.F.G., D.C.V.E., D.A.F., E.S.Y., and K.U. performed research; S.M.S., K.L.M., S.M., J.X., E.J.A., M.W.W., C.F.B., M.J., J.A., M.F.G., D.C.V.E., D.A.F., E.S.Y., and K.U. contributed new reagents/analytic tools; S.M.S., K.L.M., M.W.W., C.F.B., M.F.G., D.C.V.E., and K.U. analyzed data; and S.M.S., K.L.M., M.W.W., C.F.B., M.F.G., D.C.V.E., and K.U. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: steve@fmrib.ox.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1121329109/-DCSupplemental.

of functional integration among brain areas means that [distinct] large scale neuronal dynamics can share a substantial anatomical infrastructure" (6). Here we propose to identify extended functional networks on the basis of their temporal independence. We show that these networks are in general different from the well-known RSNs identified by spatial ICA (7) and seed-based correlation; therefore we use a different terminology, referring to "temporal functional modes" (TFMs). We use the word "modes" rather than "networks" after observing many recurring spatial maps containing significant amounts of anticorrelated spontaneous activity; presumably, two functionally distinct networks may anticorrelate (e.g., if they represent alternative brain states that do not in general simultaneously coexist), and it would be inappropriate to consider the two networks as a single functional network—hence the more general term "modes."

Although we aim to identify temporally independent modes, temporal ICA for fMRI data remains considerably less robust than spatial ICA. This means that to be able to identify TFMs robustly and reproducibly, we need to improve data quality and increase the temporal sampling rate, and also to develop an efficient analysis strategy to allow the estimation of temporal independence. We address the former by using improved fMRI pulse sequences, developed initially to address the challenges of high-resolution fMRI at ultra-high fields and subsequently developed further as part of the National Institutes of Health (NIH) Human Connectome Project. We used accelerated imaging to acquire multiple slices of data simultaneously, covering the entire brain in 0.8 s. A total of 36 10-min runs of resting fMRI data

yielded 24,000 time points concatenated from five subjects. To maximize robustness of the final temporal ICA, we align multiple subjects' datasets to each other with high accuracy and then perform an initial decomposition of the whole-group data with spatial ICA at high dimensionality. After removing artifact functional components, the remaining components are the desired functional parcels, whose time series are then fed into temporal ICA. This identifies TFMs—each temporal ICA component is a distinct TFM, comprising a time series (guaranteed to be independent of all other TFM time series) as well as a vector that describes the (positive or negative) strength of involvement of each of the many network nodes (functional parcels derived by the spatial ICA) in the given TFM.

We present results from a TFM analysis of a resting fMRI accelerated imaging dataset showing multiple functional networks that differ substantially from previously shown resting-state networks, and which are therefore of significant value in aiding our understanding of sensory, motor, and cognitive networks in the brain. We found both between-subject and between-study reproducibility.

Results

From 36 10-min sessions of resting fMRI data, concatenated temporally, we derived 200 spatial ICA components. We identified 58 of these as artifacts. The remaining 142 components have relatively compact spatial patterns, and we refer to these as "network nodes." The 142 nodes' time series were fed into 21-dimensional temporal ICA, identifying 21 TFMs. A TFM can be

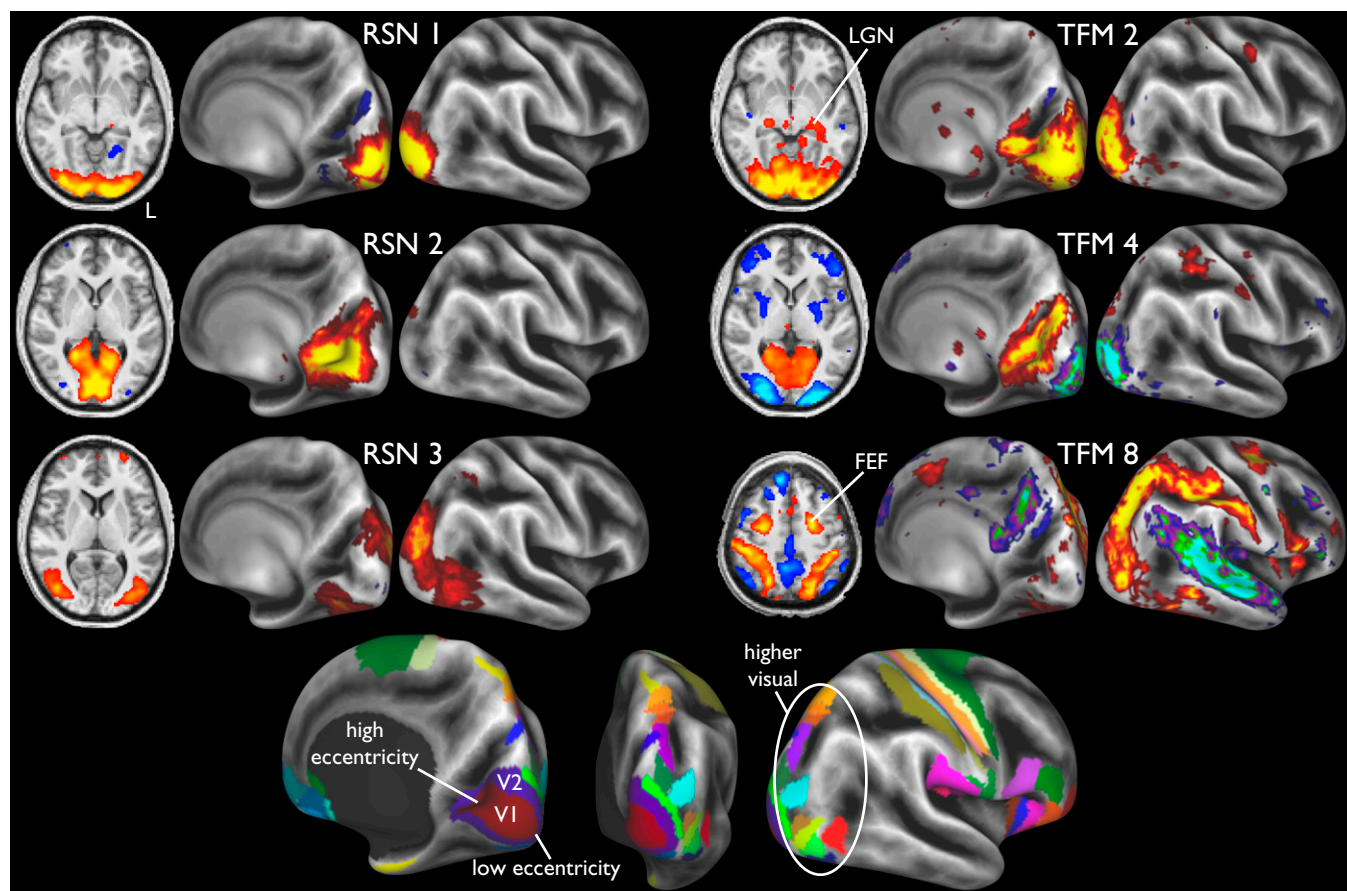


Fig. 1. Three visual components from a 21-dimensional spatial ICA decomposition of the complete dataset, as well as three components from the 21-dimensional TFM analysis. To help localize the maps structurally, they are shown on the partially inflated cortical surface, with sulci indicated by darker background intensity. To help localize the maps functionally, the bottom row shows several cytoarchitecturally based (V1 and V2) and retinotopically based (higher visual) areas from the "FS_LR" atlas (11). LGN, lateral geniculate nucleus. FEF, frontal eye fields.

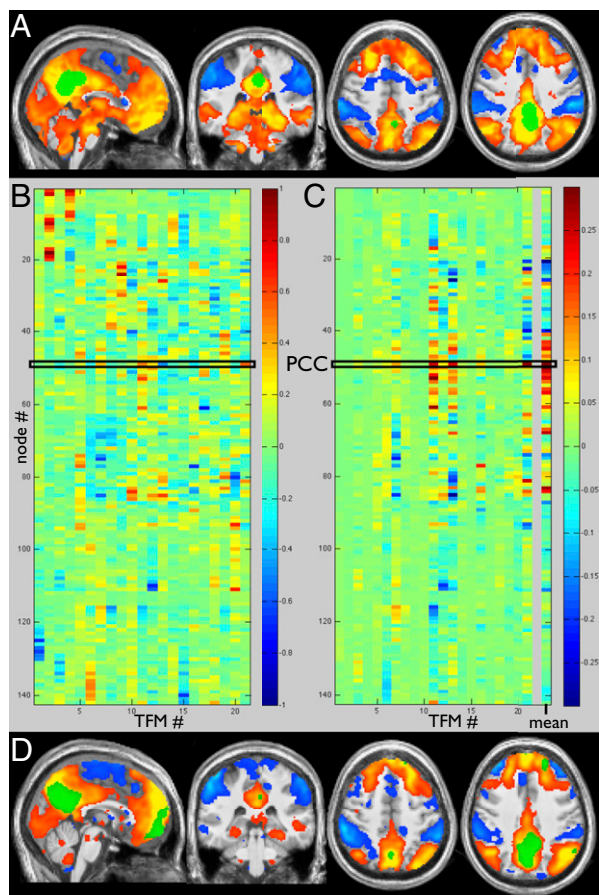


Fig. 3. Average correlations with PCC, estimated through seed-based correlation (A) and by averaging across all TFMs (B–D). (A) PCC seed region (green) is determined based on standard-space coordinates from the DMN literature: $(-5.1, -52.5, 40.7)$, $(-2, -39, 38.2)$, $(0, -54.9, 26.5)$, $(-6, -58, 28)$. The PCC's average time series is regressed against all voxels' time series, resulting in the average map of correlation (red-yellow) and anticorrelation (blue). (B) The 21 TFM node-weight vectors; each column shows the node weights that make up the "spatial map" for a given TFM. (C) Each vector was multiplied by its entry for node 49 (PCC), and the results were then averaged across all TFMs, giving the right-most column. This average depicts the node weights corresponding to the average correlation with PCC, and is shown as a voxelwise map (D) by multiplying the value for each node by the corresponding node's spatial map. The PCC node is in green.

weight was negative). This is the average PCC correlation information, calculated via the TFM decomposition (Fig. 3C, *Right*). Multiplying this by the 142 nodes' 3D voxelwise maps yields the average PCC correlation map (Fig. 3D).

Because the correlation with PCC, averaged across all TFMs, clearly shows the well-known DMN areas of correlation and anticorrelation, it is of interest to consider more carefully the message conveyed by the matrix in Fig. 3C. Several distinct TFMs contribute strongly to the average DMN appearance, in particular TFMs 7, 11, 12, 13, 16, and 21, and there is significant spatial overlap between these networks. Each of these TFMs has a distinct profile of node weights. For example, the areas showing anticorrelation with PCC (all blue elements in Fig. 3C) are quite different in the different TFMs. This indicates that separate temporal processes (that the time-averaged seed-based correlation does not separate from each other) give rise to the overall time-averaged pattern of correlations and anticorrelations traditionally associated with the DMN—it is not a single functional network. That distinct TFMs (e.g., 11, 13) relate to the DMN in quite different ways (positive correlation with semantics in TFM

11, and anticorrelation with language in TFM 13) provides insight into their functional significance.

Further Sensory-Motor and Cognitive TFMs. We now discuss the remaining "neuronal" TFMs, which can be seen in Fig. 4 (in addition to TFM 5; see below). TFM 1 contains much of the dorsal visual stream, similar to TFM 8. However, whereas in 8 the anticorrelated regions include the DMN and auditory areas, here anticorrelation is strongly found in primary motor-somatosensory (mouth or face), along with supplementary motor area (SMA) and posterior putamen/pallidum, premotor thalamus, and the cerebellar "face area." TFM 3 is dominated by primary and secondary sensory-motor areas (particularly at the level of the upper extremity representation), along with superior temporal sulcus and cerebellum VI. The anticorrelated regions are weaker, including parts of superior parietal lobule and ACC. TFM 6 has strong anticorrelation between primary sensory-motor, SMA, premotor thalamus, SII, and cingulate motor, versus primary auditory, medial DMN, and Brodmann 45. TFM 10 contains widespread motor-somatosensory, posterior superior temporal sulcus, and areas 45 and 47I, anticorrelated with posterior intraparietal sulcus (IPS), angular gyrus, and dorso-lateral prefrontal cortex (DLPFC). TFM 14 sees extrastriate visual correlating with parts of somatosensory and retrosplenial cortex; anticorrelations include large parts of posterior right frontal cortex, striate visual cortex, and left cerebellum. TFM 18 is dominated by motor-somatosensory regions, with focal anticorrelation in areas 39/40. TFM 20 contains some higher-level visual, superior parietal-occipital sulcus, anterior insula, and posterior IPS, and anticorrelated regions in retrosplenial cortex (areas 29 and 30).

TFM 7 strongly involves medial parts of the DMN (PCC/precuneus and ACC), along with primary and early association auditory cortex, parts of the parietal-occipital fissure, and frontal area 10; anticorrelated regions include superior temporal sulcus and Broca (area 45 more than 44). TFM 9 includes dorso-lateral prefrontal cortex, inferior parietal lobule, and lateral occipital areas, anticorrelated with supramarginal/angular gyrus. TFM 12 contains parts of the semantic network on the left, anticorrelated with right somatosensory association areas. TFM 16 contains bilateral superior temporal sulcus correlated with parts of premotor bilaterally, retrosplenial cortex, DLPFC, and medial-parietal cortex; anticorrelated areas include pre-SMA and more anterior medial-frontal cortex. TFM 21 sees DMN areas (including PCC) anticorrelated with lateral-frontal regions often seen as part of the "executive control" RSN. TFM 17 is dominated by a distinctive posterior-medial pattern containing two close (but not connected) areas. Because of the position in/near PCC, this TFM has likely merged into many previous PCC-seeded DMN analyses; anticorrelated regions include area 39 and Brodmann 6 and 8. TFM 19 contains parts of the DMN and strong involvement of Crus I in the cerebellum; anticorrelation includes primary and association auditory cortex.

"Global" TFM Spatial Maps. Finally, TFMs 15 (Fig. 5) and 5 (Fig. 4) are unlike the other TFMs in that they show widespread signal across much of the gray matter, with very small areas of anticorrelation. Their temporal power spectra (shown in *SI Appendix*) are largely similar to the other TFMs below 0.2 Hz, but have raised power above 0.2 Hz. There could well be some nonneural 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. TFM 15 includes anticorrelation in anterior insula (probably lateral agranular insula) and (more weakly) in the putamen. The insula is implicated in heart rate control and other autonomic functions (13–15); hence, the insular signal seen here might be neural in origin,




Fig. 4. A further 15 TFMs from the set of 21. The most informative view(s) for each is chosen from the left/right, medial/lateral surfaces.

and related to vascular (as distinct from neural) effects seen in gray matter in virtually all other regions.

TFM Reproducibility. To gain further confidence in our results, we confirmed significant reproducibility in tests between separate groups of subjects. We also found encouraging results (several TFMs matching what we have shown above, including TFMs 8, 13, and 15) from a separate study utilizing a more traditional acquisition. See *SI Appendix* for details.

Discussion

We have demonstrated spontaneously varying modes of brain function, identified on the basis of their temporal independence. This approach presents technical challenges, as it is fundamentally less robust than the seed-based correlation and spatial ICA methods in common use. However, by using accelerated MRI acquisition that generates much more densely sampled time series, along with improved cross-subject alignment and two stages of artifact removal, we identified a set of temporal functional modes that are reproducible and robust. Several of these TFMs seem to be functionally interpretable, although much remains to be done on this front.

The well-known patterns of correlation and anticorrelation that make up the default mode network and the task-positive network appear to reflect the summation of multiple functionally (temporally) distinct processes, which overlap spatially, and hence become merged together in averaged seed-based correlation analyses. The notion that the DMN comprises multiple functional subnetworks has previously been suggested, based on changing patterns of correlation when a seed region is moved across distinct neighboring regions (e.g., 16). However, such analyses have generally been based on simple correlation, and so are fundamentally limited in their ability to differentiate spatially overlapping networks from each other.

Of relevance to this, an area of increasing interest is the study of temporal nonstationarities in resting-state patterns. The term “nonstationarity” implies that a statistic of interest is non-constant, and it is used in a variety of ways; for example, nonstationarity might (*i*) refer to an apparent change over time in correlation between two regions (*2*); (*ii*) refer to changes in the mean and/or variance in the time course of a functional network (or TFM); or (*iii*) refer to the underlying network structure changing over time. If multiple functional networks are static over time in terms of each network’s “internal” connection strengths, but each has its own fluctuating overall activity level (case *ii*), the network structure may be stationary, but apparent nonstationarity in correlation will occur (e.g., as measured via sliding-window correlation, or apparent as session variability). Indeed, many other factors (that may not be functionally interesting) can also cause apparent changes in correlation, including changing noise level or poor sampling of the correlation (e.g., allowing the signal to remain dominated by wavelengths longer than the scale of the window in sliding-window correlation). Given that the TFM model allows brain regions to be part of multiple networks, this model may account for some of the apparent nonstationarity seen in simple correlation (case *i*). By reconstructing the data from the TFM model (i.e., assuming each TFM’s set of connections, and their strengths, is fixed), we can estimate the remaining apparent nonstationarity in correlation

and hence estimate how much apparent correlation nonstationarity is due purely to the overlapping networks rather than other factors such as nonmodeled noise or “deeper” nonstationarities (case *iii*). Using sliding-window correlation (see *SI Appendix* for more details), we found that at least 25% of the apparent nonstationarity in raw data correlations is attributable to the fact that multiple (internally fixed) functional networks with time-varying amplitude are spatially overlapping.

Spatially, many (although not all) of the TFMs contain significant spatial areas of both strongly positive and negative values, meaning that, with respect to a given TFM’s time course (and hence, we propose, function), there is considerable deactivation synchronized with activation. This is much less apparent in the spatial maps of resting-state networks identified from spatial ICA. Anticorrelations involving the DMN have been reported using seed-based correlation (9), and although there has been some debate over certain preprocessing enhancing (or even creating) this anticorrelation, it does now seem clear that there is indeed real underlying anticorrelation (12). The anticorrelations reported here support that conclusion, and are not subject to the same (preprocessing-related) methodological concerns. The most obvious interpretations of anticorrelations relate to mutual functional incompatibility (particularly in cognitive networks) and suppression (more likely in sensory areas). The former scenario might be most strongly argued for functional modes that require conscious involvement; if it is the case that conscious thought can only strongly be involved in one kind of cognitive task at any one moment (semantic, decision making, memory recall, etc.), this might lead to anticorrelation between the areas involved in those different tasks. If indeed there are multiple networks that are all mutually incompatible, this suggests that our focus on temporal independence, although being more appropriate than spatial independence, may still be suboptimal; one may need to relax constraints of total mutual independence (and orthogonality) in time (temporal ICA) and space (spatial ICA). One possible strategy would be to maximize non-Gaussianity in space and time, while not fully enforcing independence/orthogonality. Along these lines, although enforcing spatial sparsity may be a practical way to optimize non-Gaussianity, this is unlikely to be a sensible approach temporally (we see no evidence that the TFM time courses are sparse). Although the spatiotemporal approach of ref. 17 might be useful in this context, if it is the case that anticorrelation really does




Fig. 5. TFM 15 (global versus insula mode).

occur within a single biological component, it is unlikely to make sense to optimize for skew of the spatial maps. In the temporal domain, given that many object to the concept of *deactivation* as a meaningful biological concept (as opposed to *nonactivation*, including inhibition), it might make sense to constrain the component time courses to be only positive (or zero); this is achieved in a nonnegative matrix factorization (18). However, such methods assume that both domains (time and space) have to be fully nonnegative, so again this does not perfectly suit our needs. Instead, it would be valuable to develop methods for weakening the temporal independence constraint (and hence potentially improving interpretability of the components found), while also seeking to improve computational robustness.

We believe that the advances presented here make a contribution toward improved mapping of the brain's functional modes. Through fMRI data quality improvements such as the accelerated imaging used here, and with increased integration with other modalities such as magnetoencephalography, new modeling methods should be able to derive ever richer and more interpretable mapping of the brain's functions and networks.

Methods

For details, see *SI Appendix Methods*. We acquired 36 10-min resting fMRI datasets in seven sessions from five healthy adults (one subject being scanned on three separate occasions). The imaging was approved by the institutional review board at the University of Minnesota, and subjects provided informed consent. Data were acquired on a standard 3T Siemens Trio, with 40 mT·m⁻¹ maximum gradients and a slew rate of 200 mT·m⁻¹·s⁻¹. 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 (19) echo-planar imaging with controlled aliasing. Three slices were simultaneously excited (MB = 3), giving a whole-brain temporal resolution of 0.8 s.

Data were analyzed primarily using tools from FSL (20) as well as FreeSurfer (21) for cortical surface modeling, and FastICA (22) for temporal ICA. For visualization, we used tools from FSL and Caret (<http://brainvis.wustl.edu/wiki/index.php/Caret>About>). Each 10-min fMRI dataset was corrected for head motion and high-pass-filtered to remove drifts before being fed into spatial ICA to find and remove the majority of artifacts. We then mapped each 4D dataset into a standard representation of gray matter ("grayordinates") currently being developed for the NIH Human Connectome Project; this is a combination of cortical midgray surface vertices and 3D subcortical/cerebellar voxels (from a 3D nonlinear registration of the structural to the MNI152 standard template image). The set of grayordinates is arranged as a single long vector of vertices and voxels, and thus each fMRI time-series dataset can be represented in this standard system as a 2D matrix of grayordinates × time points.

We concatenated temporally all subjects' datasets in this format to feed into the following groupwise ICA. We applied a high-dimensional spatial ICA to identify further artifactual components as well as achieve a high-dimensional functional parcellation of the group data. From a 200-dimensional spatial ICA, we manually identified 58 components as artifacts. These were then regressed out of the remaining 142 time series, leaving 142 functional nodes—node time series and associated spatial maps. The 142 node time series were then fed into temporal ICA (the TFM analysis) with an ICA dimensionality of 21, that is, to find the 21 strongest TFMs present in the group data as a whole. Thus, the overall model (see *SI Appendix* for details) is

$$X_{(V \times T)} = S_{S(V \times K)} \times A_{t(K \times L)} \times S_{t(L \times T)} + E. \quad [1]$$

X is the data matrix of size Voxels (or grayordinates) × Timepoints, represented as a product of three matrices: S_s are the spatial maps (network nodes) estimated by the initial K -dimensional spatial ICA (after regressing out the artifact components), S_t are the TFM time courses from the L -dimensional temporal ICA, and A_t is a central mixing matrix that describes the node membership for each TFM. The l th 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 anticorrelated with each other with respect to the particular behavior of that specific TFM. Fig. 3B is showing A_t . E combines noise and artifacts.

All TFM spatial maps shown were obtained by regressing the TFM time courses 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. The sign of the spatial maps (and associated time courses) is arbitrary; inverting the sign of one spatial map and its associated time course gives an equivalent model of the data. TFM spatial maps are shown both on the inflated cortical surface (what looks like the right hemisphere is the right hemisphere) and with axial/coronal/sagittal 2D slices (shown left–right reversed).

All analysis software used is either already publicly available or will be released in upcoming versions of FSL, Caret, and Human Connectome Project toolkits. The multiband fMRI pulse sequence and reconstruction code is available from the Center for Magnetic Resonance Research, University of Minnesota, for the Siemens VB17 platform with a 32-channel receive coil, and is expected shortly to be available as a Siemens "Works in Progress" package. Collection of the full Human Connectome Project datasets is due to begin in 2012, once the imaging protocols have been finalized; these datasets will be made publicly available. Our TFM and RSN maps are available from SumsDB (http://sumsdb.wustl.edu/sums/directory.do?id=8288032&dir_name=TFM_PNAS).

ACKNOWLEDGMENTS. We are grateful to G. Douaud, A. Groves, T. Blumensath, and A. Hyvärinen for useful discussions, and to N. Filippini and C. Mackay for providing the "standard-TR" dataset. This work was funded in part by the NIH Human Connectome Project (1U54MH091657-01), as well as NIH Grants R01 EB000331, P30 NS057091, and P41 RR08079. K.L.M. was funded by the Wellcome Trust.

1. Biswal B, Yetkin FZ, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med* 34: 537–541.
2. Chang C, Glover GH (2010) Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage* 50(1):81–98.
3. McKeown MJ, et al. (1998) Analysis of fMRI data by blind separation into independent spatial components. *Hum Brain Mapp* 6(3):160–188.
4. Biswal BB, Ulmer JL (1999) Blind source separation of multiple signal sources of fMRI data sets using independent component analysis. *J Comput Assist Tomogr* 23: 265–271.
5. Calhoun VD, Adali T, Pearlson GD, Pekar JJ (2001) Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Hum Brain Mapp* 13(1):43–53.
6. Friston KJ (1998) Modes or models: A critique on independent component analysis for fMRI. *Trends Cogn Sci* 2:373–375.
7. Beckmann CF, DeLuca M, Devlin JT, Smith SM (2005) Investigations into resting-state connectivity using independent component analysis. *Philos Trans R Soc Lond B Biol Sci* 360:1001–1013.
8. Shmuel A, et al. (2002) Sustained negative BOLD, blood flow and oxygen consumption response and its coupling to the positive response in the human brain. *Neuron* 36:1195–1210.
9. Fox MD, et al. (2005) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci USA* 102:9673–9678.
10. Binder JR, Desai RH, Graves WW, Conant LL (2009) Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cereb Cortex* 19:2767–2796.
11. Van Essen DC, Glasser MF, Dierker DL, Harwell J, Coalson T (November 2, 2011) Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cereb Cortex*, 10.1093/cercor/bhr291.
12. Fox MD, Zhang D, Snyder AZ, Raichle ME (2009) The global signal and observed anticorrelated resting state brain networks. *J Neurophysiol* 101:3270–3283.
13. Oppenheimer SM, Gelb A, Girvin JP, Hachinski VC (1992) Cardiovascular effects of human insular cortex stimulation. *Neurology* 42:1727–1732.
14. Locatelli ER, Varghese JP, Shuaib A, Potolicchio SJ (1999) Cardiac asystole and bradycardia as a manifestation of left temporal lobe complex partial seizure. *Ann Intern Med* 130:581–583.
15. Lamb K, et al. (2007) Exercise-induced decrease in insular cortex rCBF during post-exercise hypotension. *Med Sci Sports Exerc* 39:672–679.
16. Margulies DS, et al. (2009) Precuneus shares intrinsic functional architecture in humans and monkeys. *Proc Natl Acad Sci USA* 106:20069–20074.
17. Stone JV, Porrill J, Porter NR, Wilkinson ID (2002) Spatiotemporal ICA of event-related fMRI data using skewed probability density functions. *Neuroimage* 15:407–421.
18. Paatero P, Tapper U (1994) Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics* 5(2):111–126.
19. Moeller S, et al. (2010) 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* 63:1144–1153.
20. Woolrich MW, et al. (2009) Bayesian analysis in FSL. *Neuroimage* 45(Suppl 1): S173–S186.
21. Fischl B, Sereno MI, Dale AM (1999) Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage* 9(2):195–207.
22. Hyvärinen A (1999) Fast and robust fixed-point algorithms for independent component analysis. *IEEE Trans Neural Netw* 10:626–634.