

Data-driven clustering reveals a fundamental subdivision of the human cortex into two global systems

Yulia Golland^a, Polina Golland^c, Shlomo Bentin^a, Rafael Malach^{b,*}

^a Department of Psychology, Hebrew University of Jerusalem, Jerusalem, Israel

^b Department of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel

^c Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, United States

Received 3 May 2007; received in revised form 16 August 2007; accepted 1 October 2007

Available online 13 October 2007

Abstract

Global organizational principles are critical for understanding cortical functionality. Recently, we proposed a global sub-division of the posterior cortex into two large-scale systems. One system, labeled *extrinsic*, comprises the sensory–motor cortex, and is associated with the external environment. The second system, labeled *intrinsic*, overlaps substantially with the previously described “default-mode” network, and is likely associated with inner-oriented processing. This global partition of the cerebral cortex emerged from hemodynamic imaging data the analysis of which was constrained by pre-determined hypotheses. Here we applied a hypothesis-free, unsupervised two-class clustering algorithm (*k*-means) to a large set of fMRI data. The two clusters delineated by this unsupervised hypothesis-free procedure showed high anatomical consistency across individuals, and their cortical topography coincided largely with the previously determined extrinsic and intrinsic systems. These new clustering-based results confirm that the intrinsic–extrinsic subdivision constitutes a fundamental cortical divide.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: fMRI; Clustering; Intrinsic activity; Default-mode; Hierarchy

1. Introduction

An appealing hypothesis in cognitive neuroscience is that the neuroanatomical organization bears direct relevance to the cognitive computations and the functional specialization implemented in the cerebral cortex. In line with this hypothesis, large-scale cortical partitions should reflect global functional principles. Such global-scale organizational principles have been proposed in cortical research both at the level of modality-specific systems (Grill-Spector & Malach, 2004; Levy, Hasson, Avidan, Hendler, & Malach, 2001; Ungerleider & Mishkin, 1982; Van Essen & Maunsell, 1983), and at the level of major cortical organization (Damasio & Damasio, 1994; Mesulam, 1998). As part of the effort to understand the functional organization of the human cortex in a recent study we proposed that the human cerebral cortex can be envisioned as a hierarchy of neuroanatomical sub-divisions, starting from large networks of

areas at the most global scale, and ending in columnar subdivisions within individual areas (Golland et al., 2007). In that study we found that the wide-spread cortical activity elicited during free-viewing a movie (Hasson, Nir, Levy, Fuhrmann, & Malach, 2004) subdivided the posterior cortical mantle into two major networks. The first network, which included major parts of sensory–motor cortex, was robustly activated by the natural movie stimuli. We labeled this network of cortical areas the “*extrinsic*” system, due to its evident association with the external input. Embedded among the externally activated areas of the extrinsic network, we found “islands” that formed a different, complementary network. While activity in these regions was not driven by external stimulation, it showed a high degree of inter-regional correlation, which suggested a common function. Since this activity was seemingly internally driven, we labeled this network the “*intrinsic*” system. The major components of the intrinsic system included medial prefrontal areas, the posterior cingulate and the precuneus, lateral inferior parietal cortex and the anterior aspect of infero-temporal cortex (Golland et al., 2007).

The anatomical extent of the intrinsic system largely corresponded to a network of areas inhibited by various cognitive

* Corresponding author. Tel.: +972 8 934 2758; fax: +972 8 934 4140.
E-mail address: rafi.malach@weizmann.ac.il (R. Malach).

tasks, termed the “default-mode”, or “task-negative” network (Fox et al., 2005; Gusnard & Raichle, 2001; Raichle et al., 2001). Indeed, many studies implicated the default network in a wide range of internally oriented activities (Binder et al., 1999; Goldberg, Harel, & Malach, 2006; Greicius, Krasnow, Reiss, & Menon, 2003; Greicius & Menon, 2004; Greicius, Srivastava, Reiss, & Menon, 2004; Gusnard, 2005; Gusnard, Akbudak, Shulman, & Raichle, 2001; Kelley et al., 2002; Kjaer, Nowak, & Lou, 2002; Lou et al., 2004; McKiernan, D’Angelo, Kaufman, & Binder, 2006; Northoff & Bermpohl, 2004; Vogeley & Fink, 2003).

The network connectivity of the default mode/intrinsic regions was shown in a number of functional connectivity studies (Fox et al., 2005; Fransson, 2005; Golland et al., 2007; Greicius et al., 2003). In addition, rest-state connectivity studies revealed a second cortical network, including several sensory and attention related regions that were negatively correlated to the default-mode network (Fox et al., 2005; Fransson, 2005). The latter network showed partial similarity to the extrinsic regions as defined by Golland et al. (2007). Despite using different labels such as task-negative/task positive, extrospective/introspective, or extrinsic/intrinsic, all these studies suggest functionally similar large-scale partition of the cerebral cortex.

The networks identified in the majority of the above studies were revealed by ‘seed’ based analysis of temporal fluctuations in hemodynamic activity. While suggestive, this approach is subject to several limitations. First, it relies on the delineation of an initial, hypothesis driven, ‘seed’ region of interest. Second, the spatial expansion of the revealed network depends on an empirically chosen threshold at which the time course correlation is deemed significant. Third, the network structure and its interaction with other networks might be specific to the particular paradigm used in the study. For example, although the resting state is assumed to be task-free and, therefore, advantageous for revealing inherent connectivity, it is also considered to show higher activity levels in the ‘default-mode’ areas (Raichle et al., 2001). Therefore, inter-systems dynamics (such as negative correlation between the two systems) observed in several resting studies could be specific to the experimental condition rather than being a fundamental property of the cortical organization and functionality.

In order to examine the hypothesis of a fundamental *bipartite* division of the cortex into intrinsic and extrinsic systems while bypassing the above limitations, we used unsupervised clustering, which is free from a priori set hypotheses. Specifically we

applied the simple and widely used *k*-means clustering algorithm (Duda & Hart, 1973; MacKay, 2003) to a large data set of fMRI experiments of different functional characteristics (36–60 min in length for each subject, $N=7$). This clustering procedure assigns cortical voxels to a predefined number of subgroups based on the similarity of their time-courses. Unsupervised clustering eliminates the dependence on pre-determined selection of seed-regions and correlation value thresholds for defining cortical networks. Instead, the clustering algorithm automatically searches for the representative time courses that play the role of the regressor in the ‘seed’ based functional connectivity analysis. Furthermore, it adaptively adjusts the correlation threshold for each voxel, based on its similarity to the representative time course. Since this study focused on the hypothesis of a global, bi-partite organization, we limited the partition of the cortical space into two clusters ($k=2$) only. The hypothesis tested was whether forcing the cortical voxels to “choose” between just two types of clusters the algorithm would consistently highlight the intrinsic/extrinsic boundary. To further explore the cluster-based cortical topography we extended this analysis to $k=3$ and $k=4$. Moreover, the potential problem of biasing the resulting cortical subdivisions by particular experimental conditions has been reduced by performing the analysis on combined fMRI data from several experiments, each with different demand characteristics. These experiments included watching a continuous audio-visual movie, a block-design visual experiment, an experiment involving internal mental tasks in the absence of either sensory input or output, and a rest-state scan.

2. Methods

2.1. Participants

All participants to the fMRI experiments provided written informed consent to participate in the experiments. The Tel-Aviv Sourasky Medical Center ethics committee approved the experimental protocols. Our study included seven subjects who participated in a wide variety of experiments, as described below and summarized in Table 1. For technical reasons subjects NT and FM could not be brought for a complete list of experiments.

Overall the functional database included 372 min of EPI scan time. The details of the functional studies are described below.

2.2. Experiments used in the clustering analysis

2.2.1. Repeated movie experiment

This experiment included an uninterrupted, 969 s long segment of the classical western “The Good the Bad and the Ugly” directed by Sergio Leone. Each

Table 1
List of experiments in the clustering data set for each individual participant

	Movie experiment 1	Movie experiment 2	Rest	Internal tasks	Visual localizer	Face morphing	Total time (min)
EI	✓	✓	✓	✓	✓		60
EQ	✓	✓	✓	✓	✓		60
FM	✓	✓				✓	40.65
NT	✓	✓			✓		35.85
SZ	✓	✓	✓	✓		✓	55.8
SO	✓	✓	✓	✓	✓		60
ZH	✓	✓	✓	✓	✓		60

Table 2
List of protocol-driven visual experiments

Experiment	1	2	3	4	5	6	7	8	Overall
Epoch/blank (sec)	12/6	12/6	12/6	12/6	9/6	9/6	12/6	15/6	
Visual Stimuli	faces	faces	faces	faces	various	objects	faces	movie clips	
Subjects	EI	√	√	√	√	√	√	√	7
	EQ		√		√	√	√	√	6
	FM						√	√	3
	NT					√			1
	SZ			√		√		√	3
	SO	√	√	√		√		√	6
	ZH	√	√	√	√	√		√	7

subject was exposed to the same movie segment in two different sessions (for details see Golland et al., 2007).

2.2.2. Rest experiment

In this experiment, participants laid in the scanner with their eyes closed. They were instructed to pay close attention to any “visual-like” percept that might occur during the scan (e.g., dots), and to report it following the scan. The experiment was 615 s long (Nir, Hasson, Levy, Yeshurun, & Malach, 2006). There were five participants in this experiment.

2.2.3. Visual experiments

2.2.3.1. Visual localizer. Visual stimulation was composed of 15 s colour movie clips from four categories (faces, buildings, navigation, and objects), followed by 6 s blank screen. The experiment was 732 s long. No task was required. Five subjects in our set participated in this study (Levy et al., 2001).

2.2.3.2. Face morphing. Each epoch was 12 s long and contained visual facial stimuli with various degrees of morphing variability followed by a 6 s of blank screen. During the experiment, one or two consecutive repetitions of the same image occurred in each epoch. Subjects performed a “one-back” task. The experiment was 480 s long (Gilaie-Dotan & Malach, 2006). Data was included for two participants in our set who didn’t participate in the visual localizer experiment.

2.2.4. Internal tasks experiment

This experiment included three different tasks: in the first ‘mental calculation’ condition, participants consecutively subtracted 13 starting at 1000; in the second ‘internal speech’ condition, participants were instructed to silently speak in a continuous manner as similar as possible to a thought flow. In the third, ‘semantic memory’ condition, participants were instructed to retrieve words belonging to ‘animals’ category.

All the conditions were performed silently with no outward articulation. Participants kept their eyes closed throughout the recording. An interleaved block design was used in this experiment. Each epoch lasted 18 s, followed by a 6 s rest. Each condition was repeated eight times. Short (500 ms) “beep” sounds marked the transition between the experimental epoch and rest. There were five participants in this study.

2.3. Protocol-driven visual experiment used in the activation/deactivation analysis

To map general regions which undergo activation/deactivation comparing to baseline we combined a series of visual experiments available for each participant, in addition to the two visual experiments described above (labeled in

Table 2 as Experiments 7 and 8). A “one-back” task was used in a block-design structure.

2.4. Data acquisition

Participants were scanned in a 1.5 T Signa Horizon LX 8.25 GE MRI scanner equipped with a standard head coil. Blood oxygenation level-dependent (BOLD) contrast was obtained with gradient-echo echo-planar imaging (EPI) sequence (TR = 3000, TE = 55, flip angle = 90, field of view 24 cm × 24 cm, matrix size 80 × 80, functional voxel size: 3 × 3 × 4). The scanned volume included 24–27 nearly axial slices of 4 mm thickness and 1 mm gap, covering the whole cortical surface. T1-weighted high-resolution (1 mm × 1 mm × 1 mm) anatomical images and 3D-spoiled gradient echo (SPGR) sequence were acquired for each subject.

2.5. Data analysis

2.5.1. Preprocessing

fMRI data were analyzed using the *BrainVoyager* software package (R. Goebel, Brain Innovation, Maastricht, The Netherlands) and in-house, Matlab-based software. The functional images were superimposed on anatomical images and incorporated into the three-dimensional data sets through trilinear interpolation. The complete data set was transformed into Talairach space. The cortical surface was reconstructed from the three-dimensional SPGR scan, unfolded, cut along the calcarine sulcus, and flattened. The obtained activation maps were superimposed on the unfolded cortex. Pre-processing of functional scans included 3D motion correction, linear trend removal and slice scan time correction. Filtering out of low frequencies of up to 10 cycles per experiment was applied to the repeated movie experiment, of up to 6 cycles to the “rest” experiment; and of up to 4 cycles to the visual localizer, face morph and internal tasks experiments. The data were spatially smoothed with a Gaussian filter of full width half maximum value (FWHM) of 8 mm. Since this study focused on cortical structures only, the analysis was limited to the cortical gray matter, therefore only voxels in the cortical segmentation mask were included. To minimize noise artifacts, we excluded voxels with low mean activation value (below 20% of the median of all included voxels), and cropped the first five time points in each time course. To exclude the effects of signal magnitude and to focus the analysis on patterns of co-variation, we converted the values in each time course to z-scores, i.e., subtracted the mean value and divided by the standard deviation. We note that this normalization fits well with employing correlation as a similarity measure, as described below.

2.5.2. k-Means clustering

Formally, we seek a partition of a set of voxels into two or more subsets, i.e., a mapping from voxel indices $v \in \{1, \dots, V\}$ to cluster labels $k \in \{1, \dots, K\}$.

We employ the well-known k -means clustering algorithm (Duda & Hart, 1973) to perform this subdivision. The objective function optimized in the estimation is defined in terms of time course similarity in each subset:

$$J = \sum_{v=1}^V d^2(\mathbf{y}_v, \mathbf{m}_{k(v)}),$$

where \mathbf{y}_v is the time course of voxel v , $\mathbf{m}_{k(v)}$ is the mean time course of the voxel's assigned cluster, and $d^2(\cdot, \cdot)$ is the squared distance function. This formulation leads to a non-linear optimization problem that is solved iteratively. The algorithm alternates between estimating the cluster membership function $k(v)$, by assigning each voxel to the closest cluster center, and updating each cluster's mean time course.

The resulting partition crucially depends on the choice of the distance function. In this study, we adopt a commonly used approach of defining similarity between time courses as their correlation coefficient:

$$\rho(x, y) = \frac{\sum_{t=1}^T (x[t] - \bar{x})(y[t] - \bar{y})}{\sqrt{\sum_{t=1}^T (x[t] - \bar{x})^2 \sum_{t=1}^T (y[t] - \bar{y})^2}}$$

where \mathbf{x} and \mathbf{y} are the time courses of length T , and \bar{x} denotes the mean value of vector \mathbf{x} . The squared distance between two time courses is then naturally defined as $d^2(\mathbf{x}, \mathbf{y}) = 1 - \rho(\mathbf{x}, \mathbf{y})$. We treated multiple experiments as independent samples from the same generative process, using a natural extension of the distance function definition:

$$d^2(\mathbf{x}, \mathbf{y}) = 1 - \frac{\sum_{m=1}^M \sum_{t=1}^{T_m} (x_m[t] - \bar{x})(y_m[t] - \bar{y})}{\sqrt{\sum_{m=1}^M \sum_{t=1}^{T_m} (x_m[t] - \bar{x})^2 \sum_{m=1}^M \sum_{t=1}^{T_m} (y_m[t] - \bar{y})^2}}$$

where M is the number of experiments, and subscripts denote different experiments.

We employed a robust estimation procedure in determining the cluster mean time courses during clustering. At each iteration of the clustering algorithm, all voxels assigned to a particular system participate in the initial estimate of the system's mean time course. The final mean time course estimate is computed using only the voxels whose correlation with the initial estimate exceeds an empirically selected threshold (0.15).

To validate that the clustering results were strongly driven by the fMRI data rather than a regularization properties of a particular clustering method, we also evaluated a full Gaussian mixture modeling and the spectral clustering as alternative algorithms for determining the partition of voxels into functional systems. In our experience, using a full covariance matrix in the mixture model is impractical due to a large number of time points T ; using a diagonal covariance matrix with variable variance elements is more appropriate in this case. The partitions estimated through mixture modeling and spectral clustering were very similar to those obtained through k -means clustering, providing further evidence of the robustness of clustering results. The percentages of voxels per participant that were labeled consistently across all methods were 93%, 94%, 87%, 93%, 85%, 90%, and 87%.

2.6. Group analysis

Clustering was performed for each participant separately. However, in order to evaluate the consistency of the resulting partitions, we compared the cluster maps across all participants. Transforming all the data into the Talairach space allowed us to directly compare cluster labels assigned to each voxel across different subjects. We note that anatomical differences among participants can cause misalignment of functional areas, contaminating the resulting group measurements. Similar to other methods in fMRI analysis, our group-wise results should directly benefit from improvements in spatial normalization methods that improve anatomical agreement across subjects.

In addition to estimating the optimal decomposition of the cortex into clusters, we also had to assign the cluster labels in a consistent way (i.e., the extrinsic system being labeled as cluster number 1 in all subjects). Cluster correspondence across subjects was established automatically by selecting the permutation of the cluster labels for each subject that maximized the agreement in the voxel labels across subjects. In order to quantify the repeatability of clustering across

participants, we computed group consistency maps for each cluster. In the group consistency map for a particular cluster, each voxel was assigned the proportion P of subjects whose individual clustering maps associated this voxel with the cluster of interest. The group consistency map quantifies agreement in clustering across participants with respect to the selected cluster.

2.7. Activation/deactivation maps

In order to assess the selective activations and de-activations in individual participants elicited by the combined protocol-driven visual experiments, we applied the standard General Linear Model analysis (Friston, 1995). A box-car predictor with hemodynamic delay of 3 s was constructed for all visual conditions and the model was independently fitted to the time course of each voxel. A regression coefficient was calculated for each predictor using the least-squares algorithm. After computing the coefficients for all regressors, we performed a two-tailed t -test of all visual conditions versus fixation on concatenated experiments, following the fixed-effect model (Fig. 3, $p < 0.001$, uncorrected).

3. Results

3.1. Two complementary networks

The clustering computation produced a spatially coherent grouping of cortical activity in two networks; this pattern was anatomically consistent across participants (Fig. 1). To estimate the anatomical consistency of this division we created a group-level consistency map, which summarized the individual maps normalized into Talairach space. The group-level map demonstrated the stability of the bi-partite patterns of individual participants (Fig. 2).

The anatomical distribution of one of the two systems (marked in blue in the figures) included auditory, visual and somato-sensory cortices, the motor cortex, and a fronto-parietal network associated with attention (Corbetta & Shulman, 2002). The second system (marked in green in the figures) included the posterior cingulate/precuneus ($P_{\text{Cing}}/P_{\text{cun}}$), lateral inferior parietal cortex (IPC), the anterior part of inferior temporal cortex, a small region in the posterior inferior frontal gyrus (pIFG), and major parts of medial and lateral prefrontal cortex.

In our previous study, searching for activation in response to natural audio-visual stimuli, we found a similar subdivision into two large-scale systems, which we labeled "extrinsic" and "intrinsic" systems (Golland et al., 2007). To compare the clustering results with our previous findings, we superimposed the borders of the intrinsic system identified through functional connectivity and reported in Golland et al. (2007) on the clustering patterns computed here. Since both types of analysis were performed on the same subjects, we could directly compare the individual results for each subject. Fig. 1 shows considerable agreement in the boundaries of the intrinsic system identified by the two methods. Critically, this agreement was consistent across different individuals demonstrating that the clustering result was not a consequence of random partition of the cortical surface. Given the close similarity of the clustering results to our previous findings, we maintained the labeling of the clustering-based partition of the cortex as extrinsic and intrinsic cortical systems.

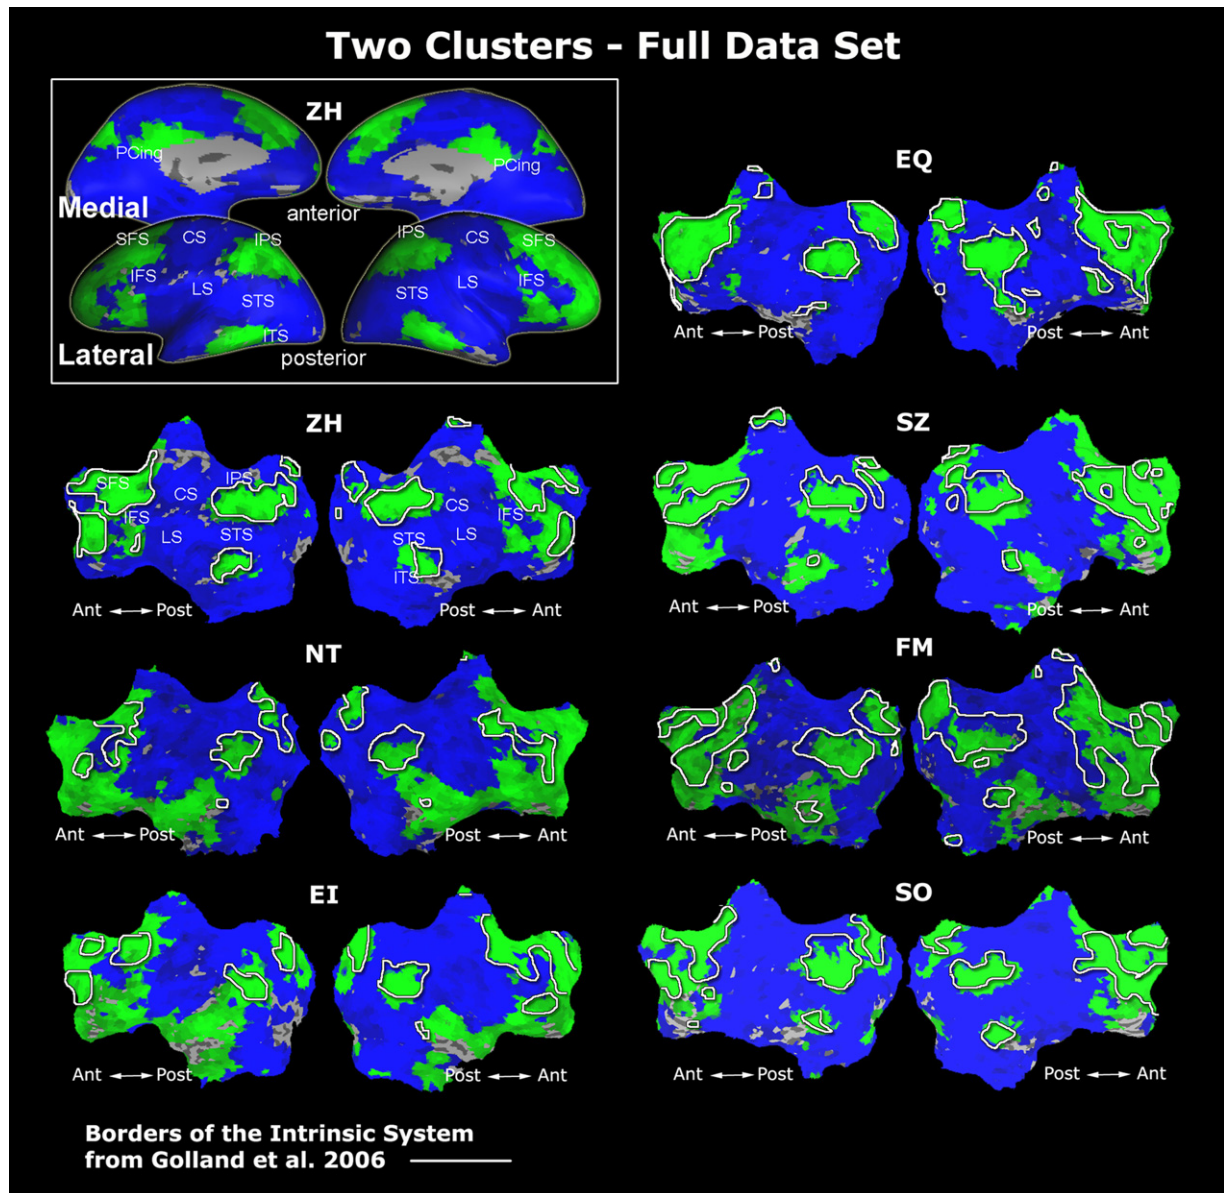


Fig. 1. The topography of clustering-based two-systems partition of cortical space. Individual clustering into two global systems, presented on unfolded cortical hemispheres for each subject. Clustering results for subject ZH are presented on an inflated and unfolded cortex (left upper corner). Borders of the intrinsic system from Golland et al. (2007) are marked by white lines for comparison. Note the strong anatomical agreement between the clustering results of the current analysis and the results from the previous study, despite the differences in the experimental setup and methodology. SFS: superior frontal sulcus; IFS: inferior frontal sulcus; CS: central sulcus; LS: lateral sulcus; IPS: intra-parietal sulcus; STS: superior temporal sulcus; ITS: inferior temporal sulcus; PCing: posterior cingulate.

We also examined the within-system coherence of the fMRI signals by examining the average correlation of the voxels' time courses within each system with that system's mean time course. Both the intrinsic and the extrinsic systems showed a similar level of internal coherence: extrinsic = 0.497, S.E.M. = 0.017; intrinsic = 0.498, S.E.M. = 0.013.

3.2. Exploring the anatomy and the functionality of the intrinsic network

The neuroanatomical organization of the intrinsic system seems to coincide largely with the previously reported "default-mode" or "task negative" networks that were essentially defined

by deactivation during a large array of perceptual and cognitive tasks (Gusnard & Raichle, 2001; Mazoyer et al., 2001; Shulman, Fiez et al., 1997). In order to directly examine the similarity between the intrinsic network and the "task-negative" activation maps, we combined the data from all visual, protocol-driven experiments that were available for the subjects in our study and mapped the regions which showed activity above or below the baseline fixation condition across all these experiments (see Table 2 for the list of experiments for each participant). These activation and deactivation maps were compared with the intrinsic system identified through clustering in individual participants. As shown in Fig. 3A, a considerable proportion of regions which showed deactivation during visual stimulation

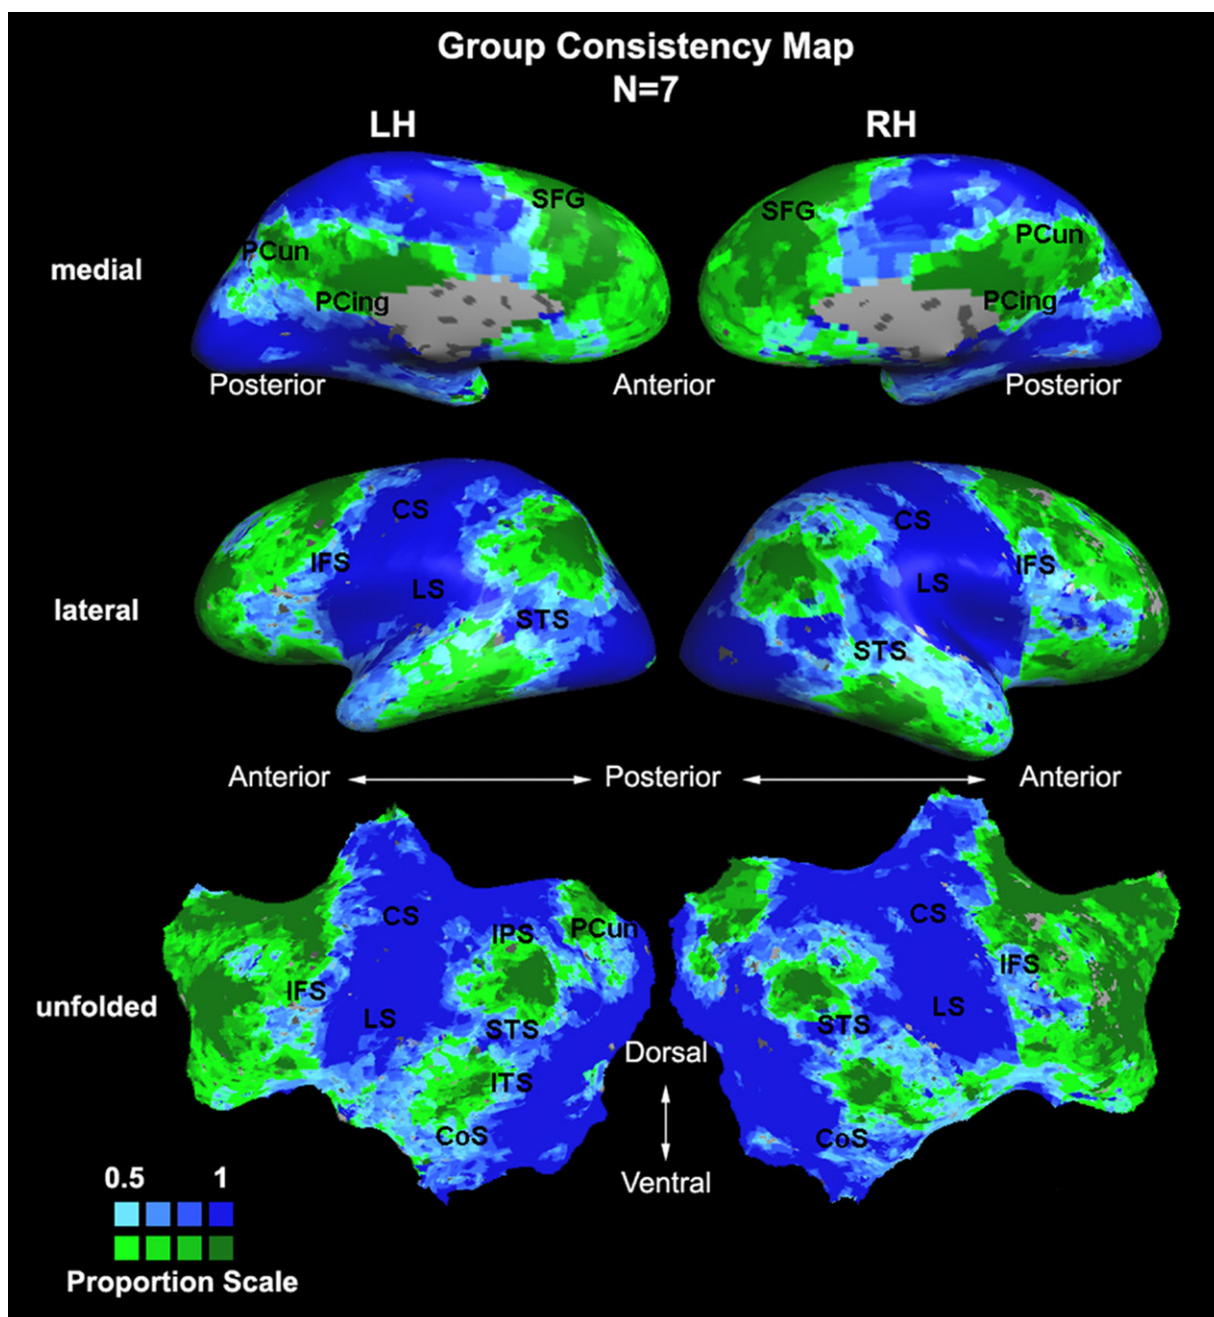


Fig. 2. Group-level consistency map. In the group consistency map for a particular network, each voxel was assigned the value of the proportion (P) of subjects whose individual clustering maps associated this voxel with the network of interest. The color depicts the strength of agreement (proportion) in clustering across subjects with respect to the selected network. Dark colors mark regions of full agreement ($P=1$) among individual results. Maps are presented on inflated (top) and unfolded left (LH) and right (RH) hemispheres.

matched the anatomical loci of the intrinsic network as defined by clustering. However, the deactivation foci appeared to be consistently smaller in extent than the clustering-defined intrinsic boundaries, and were confined primarily to the posterior cortex. This might suggest that the visually induced deactivation marks regions that constitute a regional sub-set within the intrinsic network (see Fransson, 2006). In addition, we found small patches of deactivation in auditory and language regions. However, as previously shown, these patches of deactivations are modality

specific and are not part of the default system (Shulman, Corbetta et al., 1997), nor of the intrinsic system defined in our study.

In order to directly evaluate the consistency of modulation produced by external input in the extrinsic and intrinsic systems, we compared the variability of the average systems' dynamics across participants in different experimental conditions. If an activity of a region is driven by external input, the dynamics of activity in this region should be similar across participants. We calculated the average signal from the entire intrinsic and

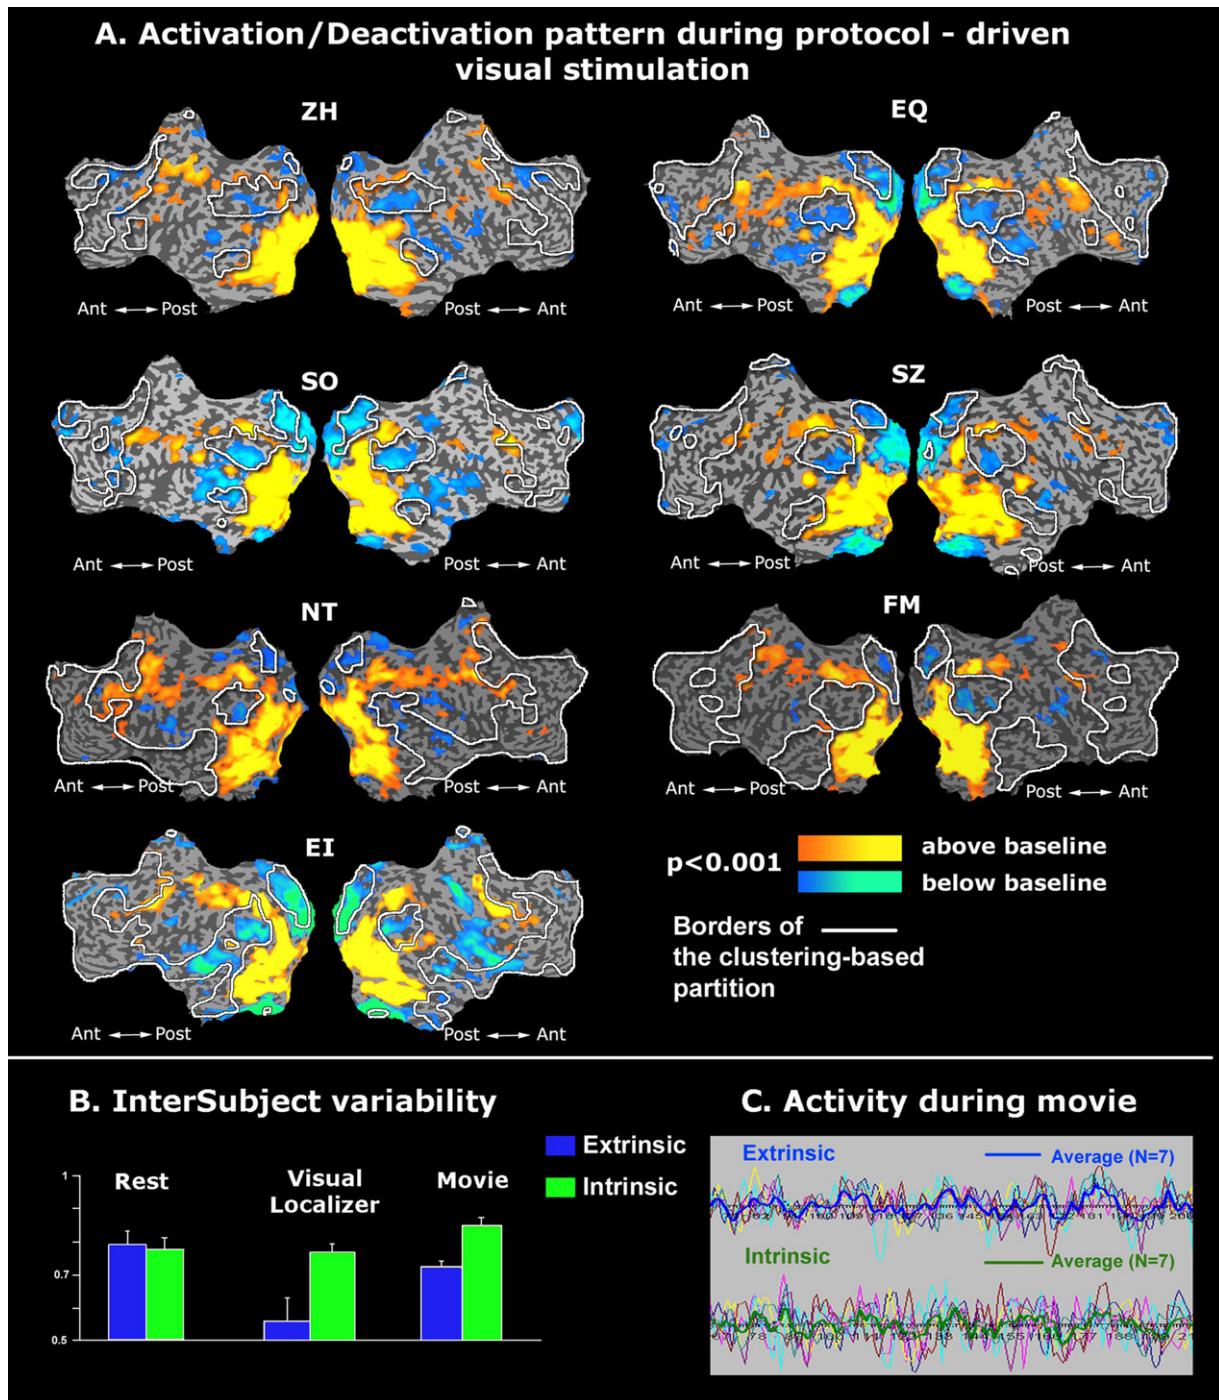


Fig. 3. (A) Activations and Deactivations during visual stimulation. Activation (orange–yellow) and deactivation (blue–green) maps in concatenated series of visual experiments presented individually for each subject (fixed effect, $p < 0.001$, uncorrected). The borders of the clustering-based partition are shown as white lines. Note that many regions within the intrinsic network undergo deactivation during visual stimulation. Additional patches of deactivation can be observed in the auditory and language related areas, likely due to the purely visual nature of stimulation. (B) Inter-subject variability of averaged activity in the extrinsic and intrinsic networks. During sensory stimulation (visual localizer and movie), but not during rest state, the variability of activation in the extrinsic system across subjects (blue bars) is substantially lower than that in the intrinsic system (green bars), reflecting the decoupling of the intrinsic system's activity from the external stimulation. (C) Time course of activity in individual subjects, averaged over the entire extrinsic and intrinsic networks, during movie presentation. The activity within the extrinsic system (upper panel) is more consistent across subjects than the activity within the intrinsic system (bottom panel) during presentation of continuous audio-visual movie.

extrinsic systems within each participant. The evaluation of the variability of activity across participants was based on calculating the variance of these individual, z-normalized average signals from the group mean signal. As shown in Fig. 3B, the

average inter-subject variability of activity in the extrinsic system was significantly lower than that in the intrinsic system during protocol-driven visual stimulation (external localizer) (extrinsic: 0.57 ± 0.07 ; intrinsic: 0.77 ± 0.02 ; $n = 5$; $p < 0.046$)

as well as during passive viewing of an audio-visual movie (extrinsic: 0.73 ± 0.01 ; intrinsic: 0.87 ± 0.02 ; $n=7$; $p < 0.003$, calculated per participant, per presentation and then averaged across two presentations). In contrast, during rest, in the absence of external stimulation, the inter-subject variability of these two systems was similar (extrinsic: 0.79 ± 0.03 ; intrinsic: 0.78 ± 0.03 ; $n=5$). The difference in the level of inter-subject variability of activity could also be discerned in the networks' time courses sampled from individual subjects during the movie presentation (Fig. 3C). The time courses of activity in the extrinsic system were more similar across subjects than those in the intrinsic system. The nature of inter-subjects variability of the fMRI activity within each system further demonstrates that the activity of the intrinsic system shows a substantial dissociation from the external sensory input.

3.3. Additional levels of clustering

In order to further explore the global pattern of cortical organization, we performed an additional analysis, imposing $k=3$ constraint. For all participants the pattern of division into three systems was identical. The borders of the intrinsic system remained highly similar to the one revealed with $k=2$ (percentages of voxels per participant that changed their labeling were: 8%, 3%, 10%, 3%, 13%, 12%, and 11%). As can be seen in Fig. 4, the extrinsic system was divided into visual areas and the rest of sensory/motor/attentional areas. Calculating the labeling consistency for $k=4$ we obtained similar results (the percentage of voxels per participant that changed their labeling from $k=2$ to 4 were 5%, 5%, 9%, 5%, 10%, 8%, and 10%). This result provides additional support to the observation that

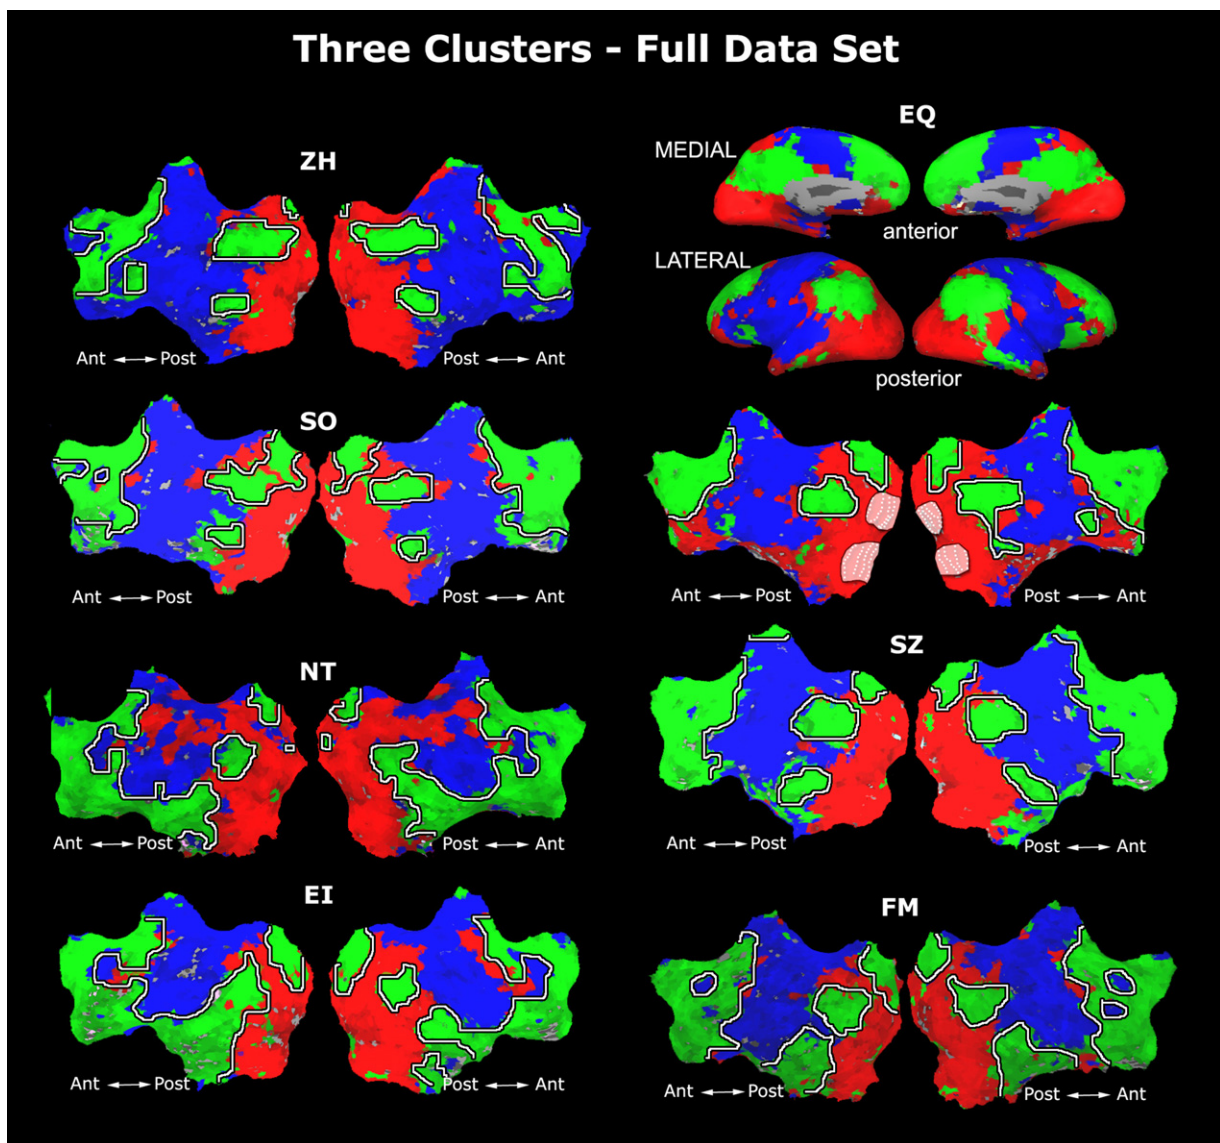


Fig. 4. The topography of clustering-based three-systems partition of cortical space. Individual clustering into three global systems, presented on unfolded cortical hemispheres for each subject. Clustering results for subject EQ are presented on an inflated and unfolded cortex (right upper corner). In addition, in this subject, visual retinotopic borders are delineated by light color on the unfolded cortical map. Borders of the intrinsic system from Fig. 1 ($k=2$) are marked by white lines for comparison. Note that while the Intrinsic system borders show strong agreement with $k=2$ results, the Extrinsic system is further subdivided into visual areas and the rest.

at the most global level of the hierarchy the cortical mantle is consistently divided into two systems. However, each of these systems may be further subdivided into functionally specific sub-systems.

4. Discussion

4.1. Two-system decomposition

The contribution of this study is twofold. First, it points to the potential usefulness of hierarchical clustering schemes for global cortical parcellation. Second, it provides further support to the previously hypothesized subdivision of the cortex into two large cortical systems. Applied to large arrays of functional data, the two systems clustered by the unsupervised *k*-means algorithm, had anatomical topographies that were highly consistent across participants. Direct anatomical exploration of this bi-partite organization showed agreement with both our previous study and other studies reported in the literature.

One of the two systems described in this study is the “extrinsic” system that corresponds to the well-studied sensory/motor cortex. The functional characteristics of this system have been explored in detail by numerous studies. At the most general level regions involved in the extrinsic system are dedicated to processing and responding to various aspects of sensory information originating in the external world. The second system is the “intrinsic” one, which complements the extrinsic system. The functional characteristics of the intrinsic system are not fully established yet. Careful examination of the anatomy and functionality of the intrinsic system revealed a substantial overlap with regions that undergo inhibition during intense cognitive and perceptual tasks (Mazoyer et al., 2001; Shulman, Fiez et al., 1997). Influenced by this unusual activation pattern, the latter regions were labeled “default” or “task negative” brain networks (Fox et al., 2005; Gusnard & Raichle, 2001; Raichle et al., 2001). Studies investigating the functional specialization in this network suggested that it plays a role in self-oriented processes (Goldberg et al., 2006; Gusnard, 2005; Gusnard et al., 2001; Kelley et al., 2002; Kjaer et al., 2002; Lou et al., 2004; Northoff & Bermpohl, 2004). Additionally, regions involved in this network were associated with a sense of agency (Vogeley & Fink, 2003), social interactions (Gallagher & Frith, 2003; Iacoboni et al., 2004), or episodic memory (Greicius et al., 2003; Greicius & Menon, 2004; Greicius et al., 2004). All these studies are largely consistent with the definition proposed on our previous study (Golland et al., 2007) of *intrinsic*, self-oriented processes as opposed to *extrinsic*, world oriented processes.

However, given the complexities of cortical function, it is not surprising that certain functions do not fall simply into either the extrinsic or intrinsic subdivisions. One example is social interaction, which may appear superficially to engage only extrinsic signals. However, recent studies suggested that social interaction also involves self-related “simulations” of the “others” feelings, intentions and actions (Mitchell, Banaji, & Macrae, 2005), which is compatible with the presumed engagement of the intrinsic system in social interaction. A second issue is mental imagery, which, again, may appear to be based on intrinsic activ-

ity. However, as a number of studies have demonstrated, mental imagery actually involves, at least partially, re-activation of sensory, “extrinsic” cortex (Amedi, Malach, & Pascual-Leone, 2005; Ishai & Sagi, 1995; Kosslyn, Behrmann, & Jeannerod, 1995; O’Craven & Kanwisher, 2000).

The anatomical boundaries of the intrinsic system estimated by the unsupervised clustering procedure were also consistent with previous studies from our laboratory (Goldberg et al., 2006; Golland et al., 2007). Furthermore, they were also compatible with other studies that examined the topography of the default mode/task-negative/introspective system (Binder et al., 1999; Fox et al., 2005; Fransson, 2005; Greicius et al., 2003; Shulman, Fiez et al., 1997). This consistency is evident in Fig. 5 where the distribution of the default/intrinsic regions that have been described in previous studies are superimposed on the intrinsic system as determined by the current clustering analysis. However, it should be noted that direct comparison of the deactivation foci that are typical of the default mode suggest they are of a more localized nature compared to the intrinsic network as defined here (see Fig. 3).

4.2. Prior evidence for two-system organization

While far reaching, the bi-partite fundamental organization of the cerebral cortex is supported by several lines of studies. Our previous study addressed the bi-partite hypothesis (Golland et al., 2007) using a different experimental set up and analytic methods. Using natural viewing of audio-visual movie and correlation analysis, we showed this global cortical partition, incorporating functional and anatomical aspects. We revealed that the posterior cortex could be envisioned as two anatomically complementary interconnected systems. While large parts of the extrinsic system were robustly activated by a naturalistic audio-visual movie, the activity of the intrinsic system was dissociated from this external input. Similarly, resting state connectivity studies revealed two spontaneously emerging networks, partially matching the extrinsic/intrinsic subdivision shown in the current study. For example, Nir et al. (2006) demonstrated wide-spread activity that is seemingly distinguished along the intrinsic/extrinsic boundaries reported here. Fox et al. (2005) demonstrated two negatively correlated networks during passive resting state. One of these networks included the default-mode network while the other contained while the other contained the dorsal attention network (Corbetta & Shulman, 2002) in the inferior parietal and superior frontal regions, in addition to dorsal-lateral and ventral prefrontal regions, insula and supplementary motor area. Other studies have also shown negative correlations between subcomponents of the intrinsic and extrinsic systems (Fransson, 2005; Greicius et al., 2003; Hampson, Peterson, Skudlarski, Gatenby, & Gore, 2002). In concert, these studies incorporate similar distinctions, such as task-negative/task-positive (Fox et al., 2005) or introspective/extrospective (Fransson, 2005) networks.

It is important to clarify that this antagonistic relationship is by no means unique to the intrinsic/extrinsic specialization. In fact antagonistic, “push pull” relationships are common to cortical function at all spatial scales—from single units to

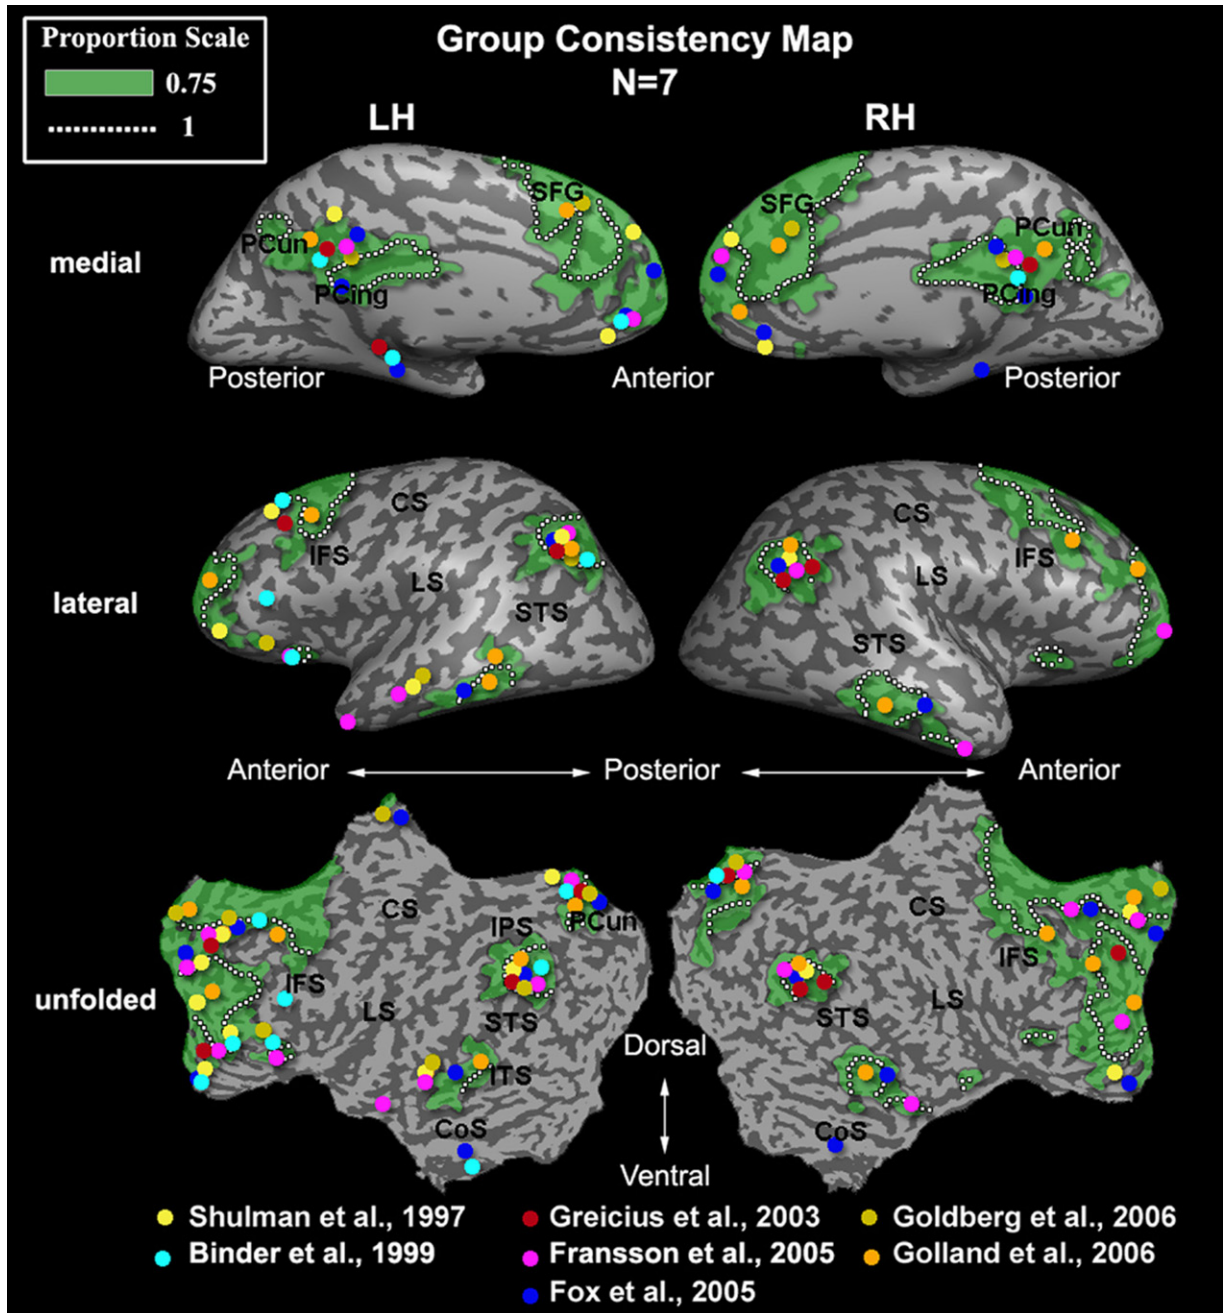


Fig. 5. Comparison with other studies. Summary of the default-mode/intrinsic activity obtained from Talairach coordinates reported in seven different studies. The group consistency map of the intrinsic network is presented for two levels of consistency: $P=0.75$, marked in green; $P=1$, marked in dotted lines. Note the substantial agreement of the intrinsic network topography with other related studies.

entire systems. However, local antagonistic relationships emerge under highly restricted cognitive tasks and sensory stimuli. For example, an antagonistic relationship has been described for peripheral versus central visual attentional shifts (e.g., Dechent & Frahm, 2003; DeYoe et al., 1996) or for modality specific activity (e.g., Amedi et al., 2005; Shulman, Fiez et al., 1997). In contrast, the Intrinsic/Extrinsic antagonism appears to emerge under a wide range of cognitive states and stimuli (see Gusnard & Raichle, 2001 for review), ranging from intense sensory-motor processing (Goldberg et al., 2006), mind wondering (Mason et

al., 2007), to rest periods (e.g., Fox et al., 2005, but see Golland et al., 2007).

Using a model-free unsupervised clustering method and combining various experiments into a large data set we provide here further evidence that the bi-partite partition is a general, anatomically consistent principle of cortical organization. The functional profile of activity within these two systems supports the extrinsic/intrinsic model, as well as other similar partitions described above (Fox et al., 2005; Fransson, 2005). During external stimulation (visual localizer and movie experiments) the

system-level average activity in the extrinsic system was significantly more similar across subjects than the average activity of the intrinsic system. During resting state, however, the average activity in these systems was similar. These results support our hypothesis that the various regions of the “extrinsic” network, as outlined by the unsupervised clustering algorithm, are dedicated to processing different aspects of external information, therefore, leading to similar neuronal dynamics across subjects. In contrast, the “intrinsic” network activity is dissociated from external input, and under intense task demands is even inhibited by it. Thus, it can be hypothesized to be oriented towards endogenous, individual-oriented processes.

4.3. Clustering approach to cortical partition

Although the clustering approach could be conceived as a generalization of a seed-based functional connectivity analysis it is different from it in several important ways. The seed-based functional connectivity analysis (Biswal, Yetkin, Haughton, & Hyde, 1995; Friston, Frith, Fletcher, Liddle, & Frackowiak, 1996; Friston, Frith, Liddle, & Frackowiak, 1993) is hypothesis dependent since it relies on the definition of an initial, hypothesis-driven, ‘seed’ region of interest. In addition, connectivity analyses frequently rely on explicit assumptions on intra-system dynamics. For example, the two-system organization described in several resting-state connectivity studies (e.g., Fox et al., 2005; Fransson, 2005) has been detected while assuming negative correlations between the activity in these two systems. Antagonistic relationships, however, do not necessarily characterize the dynamics of activity in the intrinsic and extrinsic regions (Greicius & Menon, 2004; Golland et al., 2007). These inconsistencies could reflect, for example, differences in the demand characteristics of different tasks, which influence the degree of activation/deactivation in regions associated with the extrinsic and intrinsic systems (Greicius & Menon, 2004; McKiernan et al., 2006; McKiernan, Kaufman, Kucera-Thompson, & Binder, 2003; Weissman, Roberts, Visscher, & Woldorff, 2006). A second possible source of variability among studies is the whole-brain modulation in activity (global signal). This modulation can dramatically modify the across-regional correlation (Arndt, Cizadlo, O’Leary, Gold, & Andreasen, 1996; Friston et al., 1996; Greicius & Menon, 2004).

Unsupervised clustering procedures like the one used in the present study or other data-driven approaches (see below) remove the dependence on pre-determined hypotheses and eliminate the need for assuming a priori a particular intra-system dynamics (e.g., negative-correlations). The *k*-means clustering simultaneously estimates the representative time courses that replace the seed time courses in functional connectivity analysis and assigns each voxel to a particular system based on its similarity to the representative time courses. The algorithm handles automatically a broad range of intra-system dynamics, as long as the basic assumption of monotonic similarity holds true, that is, that voxels in a particular system have higher correlation with the system’s representative time course than with the representative time courses of other systems.

4.4. Other data-driven approaches

Other unsupervised clustering methods such as the Principal Component Analysis (PCA) (Duda & Hart, 1973) and the Independent Component Analysis (ICA) (Bell & Sejnowski, 1995) provide alternative models of functional connectivity treating the data as linear combination of components (i.e., spatial maps with associated time courses).

All these methods have been extensively explored in the contexts of regression-based detection (Baumgartner, Scarth, Teichtmeister, Somorjai, & Moser, 1997; Beckmann & Smith, 2004, 2005; Fadili, Ruan, Bloyet, & Mazoyer, 2000; Filzmoser, Baumgartner, & Moser, 1999; Friston et al., 1993; Golay et al., 1998; Goutte, Toft, Rostrup, Nielsen, & Hansen, 1999; McKeown et al., 1998; Moser, Diemling, & Baumgartner, 1997; Thirion & Fugeras, 2004; Voultsidou, Dodel, & Herrmann, 2005). The fuzzy variants of clustering (Baumgartner et al., 1997; Moser et al., 1997), PCA (Friston et al., 1993) and ICA (Beckmann & Smith, 2004, 2005; McKeown et al., 1998) were successfully used to extract and eliminate components that contain physiological noise unrelated to the activation protocol, substantially improving performance of the standard, regression-based detection. They are also used in protocol-free rest scan experiments to isolate the components of fMRI signal induced by physiological rhythms, significantly improving robustness of subsequent functional connectivity analysis (De Luca, Beckmann, De Stefano, Matthews, & Smith, 2006). Specifically, many studies utilized ICA to map various global networks of interest (Bartels & Zeki, 2005; Beckmann, DeLuca, Devlin, & Smith, 2005; Damoiseaux et al., 2006; De Luca et al., 2006) as well as regions related to the default mode network (Greicius & Menon, 2004; Greicius et al., 2004). ICA is quite promising in offering a component-based decomposition of the spatiotemporal fMRI data, but the interpretation of the resulting component maps remains challenging.

We find the parcelation model used by *k*-means clustering to be particularly useful in explaining functional connectivity at different scales when integrated into a hierarchical estimation procedure. Application of this type of clustering in fMRI analysis has so far focused on grouping voxels into small, functionally homogeneous regions (Cordes, Haughton, Carew, Arfanakis, & Maravilla, 2002; Filzmoser et al., 1999; Goutte et al., 1999; Thirion & Fugeras, 2004). In contrast, here we demonstrate the benefit of using clustering to construct a top-down model of global patterns of activation spanning the entire brain.

Assuming a hierarchical organization, the clustering algorithm is applied iteratively to create a hierarchy of increasingly coherent functional regions. Using this approach we demonstrated that a bi-partite constraint provides a robust and consistent subdivision. As a next step, these major cortical systems could be subdivided into more functionally specific sub-systems. Given the success of the initial subdivision to match a number of previous findings, the hierarchical approach promises to yield reliable delineations at finer levels as well. The results of *k*=3 clustering, presented in this paper, support this approach by showing delineation of the visual system within the extrinsic system on the one hand and by being con-

sistent with the $k = 2$ results on the other hand. Thus we propose that the approach demonstrated in this study could be widely applicable to neuroanatomical research in revealing data-driven subdivisions within the human cortex.

4.5. Neuroanatomical variability

Using the data-driven unsupervised clustering and large arrays of functionally different studies, we provided evidence that the human cortical mantle can be partitioned into two anatomically consistent global systems. By incorporating the internal–external dimension, we propose to view these results as a general model of cortical organization which summarizes many local findings into a functional hierarchy of the human cortex. However, like any general model, the bi-partite organization should be considered with caution both at the functional and at the neuroanatomical levels.

At the neuroanatomical level, before adopting the clear cut distinction between the intrinsic and extrinsic systems we should be aware of the variability that can be observed among individual participants in the present study (Fig. 1) and in comparison with previous studies (Fig. 5). Interestingly, this variability is most conspicuous in the prefrontal regions and might be explained by the particular experimental paradigms used in the current study, which targeted primarily the activation of posterior cortical regions. Advanced research exploring connectivity under paradigms emphasizing frontal activations could prove useful in further tuning of the prefrontal involvement in the intrinsic and extrinsic networks. On a more general level, it is plausible that some cortical systems are not essentially ‘extrinsic’ or ‘intrinsic’, but have an intermediate functionality.

4.6. On-going versus “intrinsic” activity

At the functional level, the role of the intrinsic system in human cognition awaits further elucidation. In our previous study we found that the activity in the intrinsic system is decoupled from external input. Similarly, other studies showed that the equivalent default system is more active during rest than during performing perceptual and cognitive tasks (Mazoyer et al., 2001; McKiernan et al., 2003; Shulman, Fiez et al., 1997). However, activity decoupled from external input can hardly serve as a functional signature of the intrinsic system. A number of recent studies have demonstrated that the extrinsic system is also incessantly engaged in ongoing, spatially structured, background activity, even in the absence of experimentally induced external input (Arieli, Sterkin, Grinvald, & Aertsen, 1996; Biswal, DeYoe, & Hyde, 1996; Nir et al., 2006). Thus, both the intrinsic and extrinsic systems may manifest a reflective, spontaneous and input-unrelated cortical activity, whose function is still enigmatic (Raichle, 2006). The term “intrinsic”, as used here, points to the self-directed orientation of this system, rather than implying a unique involvement with spontaneous, on-going activity. In that sense, its “intrinsic” activity can be applied to contextually appropriate external input. Clearly, the intriguing functional characteristics of this newly discovered system awaits further elaboration.

Acknowledgments

Supported by NIMH grant R01 MH 64458 and ISF grant 816/01 to S. Bentin, by the Morris and Barbara Levinson Professorial Chair, Horowitz Foundation, the ISF Center of Excellence the Dominique Center and the Benozio Center for Neurodegeneration to R. Malach and by NSF CAREER grant 0642971 to P. Golland. We thank Yuval Nir and Sharon Gilaie-Dotan for fruitful discussions during the study. We thank M. Harel for help with the brain flattening procedure. We thank the Functional Brain Imaging (fMRI) Unit in the Wohl Institute of Advanced Imaging, Sourasky Medical Center, Tel Aviv.

References

- Amedi, A., Malach, R., & Pascual-Leone, A. (2005). Negative BOLD differentiates visual imagery and perception. *Neuron*, 48(5), 859–872.
- Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science*, 273(5283), 1868–1871.
- Arndt, S., Cizadlo, T., O’Leary, D., Gold, S., & Andreasen, N. C. (1996). Normalizing counts and cerebral blood flow intensity in functional imaging studies of the human brain. *Neuroimage*, 3(3 Pt 1), 175–184.
- Bartels, A., & Zeki, S. (2005). Brain dynamics during natural viewing conditions—A new guide for mapping connectivity in vivo. *Neuroimage*, 24(2), 339–349.
- Baumgartner, R., Scarth, G., Teichtmeister, C., Somorjai, R., & Moser, E. (1997). Fuzzy clustering of gradient-echo functional MRI in the human visual cortex. Part I. Reproducibility. *Journal of Magnetic Resonance Imaging*, 7(6), 1094–1101.
- Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*, 360(1457), 1001–1013.
- Beckmann, C. F., & Smith, S. M. (2004). Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Transactions on Medical Imaging*, 23(2), 137–152.
- Beckmann, C. F., & Smith, S. M. (2005). Tensorial extensions of independent component analysis for multisubject fMRI analysis. *Neuroimage*, 25(1), 294–311.
- Bell, A. J., & Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6), 1129–1159.
- Binder, J. R., Frost, J. A., Hammeke, T. A., Bellgowan, P. S. F., Rao, S. M., & Cox, R. (1999). Conceptual processing during the conscious resting state: A functional MRI study. *Journal of Cognitive Neuroscience*, 11(1), 80–93.
- Biswal, B., DeYoe, A. E., & Hyde, J. S. (1996). Reduction of physiological fluctuations in fMRI using digital filters. *Magnetic Resonance in Medicine*, 35(1), 107–113.
- Biswal, B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34(4), 537–541.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Review Neuroscience*, 3(3), 201–215.
- Cordes, D., Haughton, V., Carew, J. D., Arfanakis, K., & Maravilla, K. (2002). Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magnetic Resonance Imaging*, 20(4), 305–317.
- Damasio, A. R., & Damasio, H. (1994). Cortical systems for retrieval of concrete knowledge: The convergence zone framework. In C. Koch & L. D. Davis (Eds.), *Large-scale neuronal theories of the brain* (pp. 61–74). Cambridge: The MIT Press.
- Damoiseaux, J. S., Rombouts, S. A., Barkhof, F., Scheltens, P., Stam, C. J., Smith, S. M., et al. (2006). Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences of the United States of America*, 103(37), 13848–13853.

- De Luca, M., Beckmann, C. F., De Stefano, N., Matthews, P. M., & Smith, S. M. (2006). fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *Neuroimage*, 29(4), 1359–1367.
- Dechent, P., & Frahm, J. (2003). Characterization of the human visual V6 complex by functional magnetic resonance imaging. *European Journal of Neuroscience*, 17(10), 2201–2211.
- DeYoe, E. A., Carman, G. J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., et al. (1996). Mapping striate and extrastriate visual areas in human cerebral cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 93(6), 2382–2386.
- Duda, R. O., & Hart, P. E. (1973). *Pattern classification and scene analysis*. New York: John Wiley & Sons., pp. 189–225.
- Fadili, M. J., Ruan, S., Bloyet, D., & Mazoyer, B. (2000). A multistep unsupervised fuzzy clustering analysis of fMRI time series. *Human Brain Mapping*, 10(4), 160–178.
- Filzmoser, P., Baumgartner, R., & Moser, E. (1999). A hierarchical clustering method for analyzing functional MR images. *Magnetic Resonance Imaging*, 17(6), 817–826.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27), 9673–9678.
- Fransson, P. (2005). Spontaneous low-frequency BOLD signal fluctuations: An fMRI investigation of the resting-state default mode of brain function hypothesis. *Human Brain Mapping*, 26(1), 15–29.
- Fransson, P. (2006). How default is the default mode of brain function? Further evidence from intrinsic BOLD signal fluctuations. *Neuropsychologia*, 44(14), 2836–2845.
- Friston, K. J. (1995). Commentary and opinion. II. Statistical parametric mapping: Ontology and current issues. *Journal of Cerebral Blood Flow and Metabolism*, 15(3), 361–370.
- Friston, K. J., Frith, C. D., Fletcher, P., Liddle, P. F., & Frackowiak, R. S. (1996). Functional topography: Multidimensional scaling and functional connectivity in the brain. *Cerebral Cortex*, 6(2), 156–164.
- Friston, K. J., Frith, C. D., Liddle, P. F., & Frackowiak, R. S. (1993). Functional connectivity: The principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow and Metabolism*, 13(1), 5–14.
- Gallagher, H. L., & Frith, C. D. (2003). Functional imaging of ‘theory of mind’. *Trends in Cognitive Sciences*, 7(2), 77–83.
- Gilaie-Dotan, S., & Malach, R. (2006). Sub-exemplar shape tuning in human face-related areas. *Cerebral Cortex*.
- Golay, X., Kollias, S., Stoll, G., Meier, D., Valavanis, A., & Boesiger, P. (1998). A new correlation-based fuzzy logic clustering algorithm for fMRI. *Magnetic Resonance in Medicine*, 40(2), 249–260.
- Goldberg, I. I., Harel, M., & Malach, R. (2006). When the brain loses its self: Prefrontal inactivation during sensorimotor processing. *Neuron*, 50(2), 329–339.
- Golland, Y., Bentin, S., Gelbard, H., Benjamini, Y., Heller, R., Nir, Y., et al. (2007). Extrinsic and intrinsic systems in the posterior cortex of the human brain revealed during natural sensory stimulation. *Cerebral Cortex*, 17(4), 766–777.
- Goutte, C., Toft, P., Rostrup, E., Nielsen, F., & Hansen, L. K. (1999). On clustering fMRI time series. *Neuroimage*, 9(3), 298–310.
- Greicius, M. D., Krasnow, B., Reiss, A. L., & Menon, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences of the United States of America*, 100(1), 253–258.
- Greicius, M. D., & Menon, V. (2004). Default-mode activity during a passive sensory task: Uncoupled from deactivation but impacting activation. *Journal of Cognitive Neuroscience*, 16(9), 1484–1492.
- Greicius, M. D., Srivastava, G., Reiss, A. L., & Menon, V. (2004). Default-mode network activity distinguishes Alzheimer’s disease from healthy aging: Evidence from functional MRI. *Proceedings of the National Academy of Sciences of the United States of America*, 101(13), 4637–4642.
- Grill-Spector, K., & Malach, R. (2004). The human visual cortex. *Annual Review of Neuroscience*, 27, 649–677.
- Gusnard, D. A. (2005). Being a self: Considerations from functional imaging. *Consciousness and Cognition*, 14(4), 679–697.
- Gusnard, D. A., Akbudak, E., Shulman, G. L., & Raichle, M. E. (2001). Medial prefrontal cortex and self-referential mental activity: Relation to a default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America*, 98(7), 4259–4264.
- Gusnard, D. A., & Raichle, M. E. (2001). Searching for a baseline: Functional imaging and the resting human brain. *Nature Review Neuroscience*, 2(10), 685–694.
- Hampson, M., Peterson, B. S., Skudlarski, P., Gatenby, J. C., & Gore, J. C. (2002). Detection of functional connectivity using temporal correlations in MR images. *Human Brain Mapping*, 15(4), 247–262.
- Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Inter-subject synchronization of cortical activity during natural vision. *Science*, 303(5664), 1634–1640.
- Iacoboni, M., Lieberman, M. D., Knowlton, B. J., Molnar-Szakacs, I., Moritz, M., Throop, C. J., et al. (2004). Watching social interactions produces dorsomedial prefrontal and medial parietal BOLD fMRI signal increases compared to a resting baseline. *Neuroimage*, 21(3), 1167–1173.
- Ishai, A., & Sagi, D. (1995). Common mechanisms of visual-imagery and perception. *Science*, 268(5218), 1772–1774.
- Kelley, W. M., Macrae, C. N., Wyland, C. L., Caglar, S., Inati, S., & Heatherton, T. F. (2002). Finding the self? An event-related fMRI study. *Journal of Cognitive Neuroscience*, 14(5), 785–794.
- Kjaer, T. W., Nowak, M., & Lou, H. C. (2002). Reflective self-awareness and conscious states: PET evidence for a common midline parietofrontal core. *Neuroimage*, 17(2), 1080–1086.
- Kosslyn, S. M., Behrmann, M., & Jeannerod, M. (1995). The cognitive neuroscience of mental imagery. *Neuropsychologia*, 33(11), 1335–1344.
- Levy, I., Hasson, U., Avidan, G., Hendler, T., & Malach, R. (2001). Center-periphery organization of human object areas. *Nature Neuroscience*, 4(5), 533–539.
- Lou, H. C., Luber, B., Crupain, M., Keenan, J. P., Nowak, M., Kjaer, T. W., et al. (2004). Parietal cortex and representation of the mental self. *Proceedings of the National Academy of Sciences of the United States of America*, 101(17), 6827–6832.
- MacKay, D. J. C. (2003). *Information theory, inference, and learning algorithms*. Cambridge University Press., pp. 300–306.
- Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering minds: The default network and stimulus-independent thought. *Science*, 315(5810), 393–395.
- Mazoyer, B., Zago, L., Mellet, E., Bricogne, S., Etard, O., Houde, O., et al. (2001). Cortical networks for working memory and executive functions sustain the conscious resting state in man. *Brain Research Bulletin*, 54(3), 287–298.
- McKeown, M. J., Makeig, S., Brown, G. G., Jung, T. P., Kindermann, S. S., Bell, A. J., et al. (1998). Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping*, 6(3), 160–188.
- McKiernan, K. A., D’Angelo, B. R., Kaufman, J. N., & Binder, J. R. (2006). Interrupting the “stream of consciousness”: An fMRI investigation. *Neuroimage*, 29(4), 1185–1191.
- McKiernan, K. A., Kaufman, J. N., Kucera-Thompson, J., & Binder, J. R. (2003). A parametric manipulation of factors affecting task-induced deactivation in functional neuroimaging. *Journal of Cognitive Neuroscience*, 15(3), 394–408.
- Mesulam, M. M. (1998). From sensation to cognition. *Brain*, 121(Pt 6), 1013–1052.
- Mitchell, J. P., Banaji, M. R., & Macrae, C. N. (2005). The link between social cognition and self-referential thought in the medial prefrontal cortex. *Journal of Cognitive Neuroscience*, 17(8), 1306–1315.
- Moser, E., Diemling, M., & Baumgartner, R. (1997). Fuzzy clustering of gradient-echo functional MRI in the human visual cortex. Part II. Quantification. *Journal of Magnetic Resonance Imaging*, 7(6), 1102–1108.
- Nir, Y., Hasson, U., Levy, I., Yeshurun, Y., & Malach, R. (2006). Widespread functional connectivity and fMRI fluctuations in human visual cortex in the absence of visual stimulation. *Neuroimage*, 30(4), 1313–1324.
- Northoff, G., & Bermpohl, F. (2004). Cortical midline structures and the self. *Trends in Cognitive Sciences*, 8(3), 102–107.

- O'Craven, K. M., & Kanwisher, N. (2000). Mental imagery of faces and places activates corresponding stimulus-specific brain regions. *Journal of Cognitive Neuroscience*, 12(6), 1013–1023.
- Raichle, M. E. (2006). Neuroscience. The brain's dark energy. *Science*, 314(5803), 1249–1250.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America*, 98(2), 676–682.
- Shulman, G. L., Corbetta, M., Buckner, R. L., Raichle, M. E., Fiez, J. A., Miezin, F. M., et al. (1997). Top-down modulation of early sensory cortex. *Cerebral Cortex*, 7(3), 193–206.
- Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F., Raichle, M. E., et al. (1997). Common blood flow changes across visual tasks. II. Decreases in cerebral cortex. *Journal of Cognitive Neuroscience*, 9(5), 648–663.
- Thirion, B., & Fugeras, O. (2004). Feature characterization in fMRI data: The Information Bottleneck approach. *Medical Image Analysis*, 8(4), 403–419.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549–586). Cambridge, MA: The MIT Press.
- Van Essen, D. C., & Maunsell, J. H. R. (1983). Hierarchical organization and functional streams in the visual cortex. *Trends in Neuroscience*, 6, 370–375.
- Vogel, K., & Fink, G. R. (2003). Neural correlates of the first-person-perspective. *Trends in Cognitive Sciences*, 7(1), 38–42.
- Voultsidou, M., Dodel, S., & Herrmann, J. M. (2005). Neural networks approach to clustering of activity in fMRI data. *IEEE Transactions on Medical Imaging*, 24(8), 987–996.
- Weissman, D. H., Roberts, K. C., Visscher, K. M., & Woldorff, M. G. (2006). The neural bases of momentary lapses in attention. *Nature Neuroscience*, 9(7), 971–978.