

USC-SIPI REPORT #423

Non-Local Means Filtering Reveals Real-Time Whole-Brain Cortical Interactions in Resting fMRI

by

**Chitresh Bhushan, Minqi Chong, Soyoung Choi, Anand A. Joshi,
Justin P. Haldar, Hanna Damasio, and Richard M. Leahy**

August 2015

Signal and Image Processing Institute
UNIVERSITY OF SOUTHERN CALIFORNIA
Viterbi School of Engineering
Department of Electrical Engineering-Systems
3740 McClintock Avenue, Suite 400
Los Angeles, CA 90089-2564 U.S.A.

Non-local means filtering reveals real-time whole-brain cortical interactions in resting fMRI

Chitresh Bhushan¹, Minqi Chong¹, Soyoung Choi^{1,3}, Anand A. Joshi^{1,2},
Justin P. Halдар^{1,2}, Hanna Damasio², and Richard M. Leahy ^{*1}

¹*Signal and Image Processing Institute, University of Southern California, Los Angeles, CA, USA*

²*Brain and Creativity Institute, Dornsife College of Letters, Arts and Sciences, University of Southern California, Los Angeles, CA, USA*

³*Cardiology, Children's Hospital Los Angeles, Los Angeles, CA, USA*

Abstract

Intensity variations over time in resting BOLD fMRI exhibit spatial correlation patterns consistent with a set of large scale cortical networks. However, visualizations of this data on the brain surface, even after extensive preprocessing, are dominated by local intensity fluctuations that obscure larger scale behavior. Our novel adaptation of non-local means (NLM) filtering, which we refer to as temporal NLM or tNLM, reduces these local fluctuations without the spatial blurring that occurs when using standard linear filtering methods. We show examples of tNLM filtering that allow direct visualization of spatio-temporal behavior on the cortical surface. These results reveal patterns of activity consistent with known networks as well as more complex dynamic changes within and between these networks. This ability to directly visualize brain activity may facilitate the development of new insights into spontaneous brain dynamics. Temporal NLM can also be used as a preprocessor for resting fMRI for exploration of dynamic brain networks. We demonstrate its utility through application to graph-based functional parcellation, showing improved performance relative to unfiltered and Laplace-Beltrami smoothed data.

1 Introduction

Low frequency fluctuations in BOLD activity during resting functional MRI (rfMRI) exhibit correlations between cortical regions that are known to be functionally related, as first shown by Biswal et al. [1]. These correlations are the basis for identification of functional networks from rfMRI in individuals and groups [2–4].

Resting BOLD data are typically preprocessed with a pipeline that may include compensation for susceptibility-related (B_0) distortion, slice timing and subject motion, as well

*email: leahy@sipi.usc.edu; Corresponding author

as high-pass filtering of the time series at each voxels and removal of ICA-identified temporal noise components [5]. The data are then resampled onto a cortical surface extracted from coregistered T1-weighted structural MR to allow investigation of cortical networks [6]. A large number of methods for identifying anatomic regions involved in different networks are now available, including seed-based correlation, independent components analysis (ICA) and cortical parcellation methods [2]. However, when visualized as a time series or movie of cortical images, correlated patterns of BOLD variation reflecting time-varying brain activity are not readily visible in the data, even after preprocessing. Rather, the local intensity variations across the brain at each time point obscure larger scale correlated activity, as shown in Fig. 1(a) and the linked movie.

To reduce noise in fMRI it is common to spatially smooth the data, typically with an isotropic kernel applied in the volumetric space [3–5, 7–10]. This smoothing will inevitably mix signals from different areas, including areas that are not directly adjacent with respect to the cortical surface, for example when blurring occurs across the void from one side of a sulcal bank to the other. To avoid this problem, data can be smoothed in the 2D manifold of the cortical surface [3, 8, 9]. Because cortical surfaces have an intrinsic curvature that prevents them being mapped without distortion onto the plane, isotropic smoothing on the surface requires that curvature be taken into account. The Laplace Beltrami operator, which generalizes the 2D Gaussian smoothing kernel from the plane to an arbitrary smooth manifold [11], can produce isotropic smoothing on cortex.

While Laplace Beltrami smoothing can avoid blurring across sulcal banks, the smoothing is still linear and isotropic so that sharp spatial features on the surface will become blurred after filtering. This smoothing will blur boundaries between distinct functional regions. As we show below, the resulting signal mixing can also confound cortical parcellation methods, introducing artifactual parcels purely as a result of smoothing. An example of Laplace Beltrami smoothing is shown in Fig. 1(b) with a corresponding linked movie.

The primary contribution of this paper is to describe an alternative nonlinear surface filtering method that reduces noise while also respecting functional boundaries on the cortical surface. As we illustrate in Fig. 1(c) and the corresponding linked movie, this results in the ability to directly visualize cortical brain activity and networks in resting state by playing back the filtered data on the cortical surface.

Techniques for edge preserving filtering that have been developed and described in the image processing literature can be adapted to the cortical surface. Anisotropic diffusion filtering in 3D is a natural extension of Gaussian smoothing in which a measure of edge strength (the image gradient) is used to reduce smoothing across edges or boundaries [12]. This approach can be extended to surfaces using an anisotropic form of the Laplace Beltrami operator as we described in [13]. However this denoising method uses only local information to denoise the data. More recently the concept of Non-Local Means (NLM) filtering was introduced [14]. In common with conventional linear filtering, NLM uses weighted spatial averaging to reduce noise. However rather than using a fixed set of weights applied to pixels in a local neighborhood, a larger neighborhood is used with the weights adaptively chosen to emphasize those pixels that have a similar local spatial structure to that of the pixel being filtered. Using this approach, NLM filtering has the effect of reducing noise through averaging but, simultaneously, retaining spatial structure by favoring those pixels that have similar local structure.

NLM filtering has previously been applied to structural [15–18], functional [19–21], and diffusion [22–24] MRI data. Modified NLM methods tailored to MRI data have been developed that use blockwise filtering and automatically adapt weighting parameters based on signal to noise ratio [15, 16, 23, 24]. Manjón et al. [17] describe extensions to multi-component MRI data sets. Similarly, St-Jean et al. [22] have recently described extensions for HARDI MRI in which multiple angular components are used to filter rather than filtering each component image separately. All of these methods use spatial similarity over one or more images as the basis for NLM smoothing. While this approach can be applied to fMRI time series [19–21], filtering each temporal frame separately with NLM will produce a time-varying smoothing kernel, confounding subsequent time-series analysis.

We have developed a novel variation on NLM which we refer to as temporal NLM (tNLM). Our method directly exploits the temporal information in the data by replacing the standard spatial similarity weighting in NLM with a weighting that is based on the correlation between time series. As a result we reduce noise by averaging only those pixels that have similar time series. This has the effect of not smoothing across functional boundaries, since the time series in different functional areas will be less strongly correlated than within each distinct functional region. The practical effect of this procedure is illustrated in Fig. 1(c) and the corresponding linked movie. Spatio-temporal data were also used to compute weights in a recent report from one of the authors on NLM denoising of dynamic PET data [25]. In that case the weights used a combination of spatial and temporal information.

In addition to demonstrating the impact of tNLM on the spatio-temporal structure of rfMRI data, we also illustrate its utility through one application: functional cortical parcellation. A number of parcellation methods have recently been described that include spectral clustering, hierarchical clustering, edge detection, and snow-balling [2–4, 8, 9, 26]. Our goal here is to illustrate the effect on cortical parcellation of tNLM filtering and compare it to unfiltered and linearly filtered data. For this purpose we use a spectral clustering method based on normalized cuts [27], although tNLM could also be used as part of a pre-processing pipeline and combined with other parcellation methods.

2 Methods

2.1 Data

All of the results in this report (and in the supplemental data) used the minimally pre-processed rfMRI data from 40 unrelated subjects from the Human Connectome Project [5, 6, 28]. Functional MRI data sets were acquired for four independent resting state sessions of 15 minutes each (TR=720ms, TE=33.1ms, $2 \times 2 \times 2$ mm voxel) and the subjects were asked to relax and fixate on a projected bright cross-hair on a dark background [28]. Minimal preprocessing corrected the data for spatial and temporal artifacts including B_0 distortions, head motion, temporal drift and non-neuronal physiological artifacts and co-registered the functional images with their corresponding structural images [5, 6]. All results use rfMRI data resampled onto the 32K Conte69 cortical mesh [6, 29]. The only additional preprocessing step we introduced prior to filtering was to normalize the time series associated with each cortical vertex to zero mean and unit variance.

The functional parcellation results presented below are evaluated using the functional-localizer maps obtained from task-based fMRI for the same 40 subjects, as made available by the Human Connectome Project. We used data for six major task domains [30] which included somatosensory and motor systems, language processing, social cognition, relational processing, emotion processing and decision making (gambling). We used the pre-processed and analyzed task-based fMRI data (8 mm Gaussian smoothing) resampled on the 32K Conte69 cortical mesh, which yielded a total of 15 statistical task-pair activation maps. Pre-processing and analysis details for extraction of statistical parametric maps from the task-based fMRI are described in [6, 30, 31].

2.2 Spatio-temporal filtering of rfMRI

We compare our tNLM filtering method with linear isotropic filtering on the cortical surface. As described in the Introduction, Laplace Beltrami (LB) smoothing generalizes the concept of isotropic Gaussian filtering to smooth surfaces or manifolds. The LB operator solves the heat equation on the cortical surface with the degree of smoothing determined by a single time parameter t . The LB operator can be expressed in discrete form as a square matrix of dimension equal to the number of vertices on the cortical surface. LB filtering of multiple images that constitute a spatio-temporal rfMRI data set was performed efficiently using a truncated eigenvector expansion of the LB operator [32, 33].

LB filtering does not account for temporal structure in the data and tends to blur signals across functional boundaries. For this reason we propose instead the use of temporal Non-Local Means (tNLM), a novel variation on spatial NLM. Spatial NLM denoising replaces each pixel intensity with a specially weighted average of the pixel intensities drawn from a large area surrounding that pixel [14]. The NLM weights are based on a measure of similarity between small local regions or neighborhoods surrounding each pixel: when the structure of the neighboring regions is similar the weight is large; and when it is dissimilar, the weight is low. In this way, the weighted averaging tends to reinforce spatial structure while removing noise.

We are interested in identifying functional regions that share common temporal variations. For this reason tNLM uses a weight based on the similarity of the time series rather than a spatial patch to filter the data. Specifically, let $d(s, \tau)$ denote the rfMRI data at surface vertex s at time τ . Then the corresponding tNLM-filtered rfMRI is given by

$$f(s, \tau) = \frac{1}{\sum_{r \in \mathcal{N}(s)} w(s, r)} \sum_{r \in \mathcal{N}(s)} d(r, \tau) w(s, r), \quad (1)$$

where $\mathcal{N}(s)$ denotes a set of vertices lying in a large neighborhood surrounding vertex s (in our results we use all vertices that are $D = 11$ or fewer edges away from s), and the weights $w(s, r)$ are given by

$$w(s, r) = \exp\left(-\frac{\frac{1}{T} \|\underline{d}(s) - \underline{d}(r)\|^2}{h^2}\right) \quad (2)$$

with $\underline{d}(s) = [d(s, 1), \dots, d(s, T)]^\top$ a vector of length T representing the time series at vertex s . The parameter h determines the rate at which the weights decrease with similarity of the

two time series. Since we pre-process the time series at each vertex to have zero mean and unit variance, the normalized weights in eq. 1 are equivalent to using Pearson’s correlation coefficient, $\text{corr}(\underline{d}(s), \underline{d}(r))$ between $\underline{d}(s)$ and $\underline{d}(r)$ since $\frac{1}{T} \|\underline{d}(s) - \underline{d}(r)\|^2 = 2 - 2 \times \text{corr}(\underline{d}(s), \underline{d}(r))$.

The degree of smoothing in tNLM depends on the parameter h and the size of neighborhood \mathcal{N} , as parameterized by the maximum link distance D on the surface. Based on visual inspection of the resulting smoothed data and a preliminary performance study we chose a value of $h = 0.72$ and a maximum link distance of $D = 11$ to define the neighborhood in tNLM for the majority of results presented below. For comparisons of tNLM with LB filtering we selected a value of t , the parameter that determines the degree of LB smoothing, to match the mutual information of the smoothed spatio-temporal data between the two methods. For the tNLM case $D = 11, h = 0.72$ we used $t = 4$.

2.3 Identifying Cortical Networks

To explore the impact of tNLM filtering on cortical parcellation we used a graph-based spectral clustering method to identify a set of functional networks for each subject. We represent the spatio-temporal rfMRI data as a graph $G = (V, A)$ where the set of vertices of the cortical tessellation are the nodes $v \in V$ of the graph and A is the adjacency (edge) matrix such that any two vertices $u, v \in V$ are connected by an undirected edge of strength $A(u, v) = \exp(\underline{d}^\top(u)\underline{d}(v)/T)$.

The normalized-cuts (N-cuts) algorithm [27] subdivides the nodes of the graph (or vertices on the tessellated cortical surface), V , into K disjoint subsets V_1, V_2, \dots, V_K so that $\cap_{i=1}^K V_i = \emptyset$ and $\cup_{i=1}^K V_i = V$. N-cuts partitions the graph to maximize the average “normalized association” within each of the K sub-graphs:

$$\text{assoc} = \frac{1}{K} \sum_{i=1}^K \left(\frac{\sum_{u,v \in V_i} A(u, v)}{\sum_{u \in V_i, v \in V} A(u, v)} \right) \quad (3)$$

Yu and Shi [27, 34] show this cost is equivalent to minimizing the average normalized cut cost. N-cuts therefore finds the set of K subgraphs that have the weakest normalized average connectivity between subgraphs and the maximum connectivity within each subgraph. We use the implementation of N-cuts provided by the authors <http://www.cis.upenn.edu/~jshi/software/>.

Note that the graph definition described above produces a fully connected graph that contains no explicit spatial information about each vertex’s neighborhood structure. It is common to explicitly introduce spatial neighborhood information for functional clustering of rfMRI data [2–4]. In preliminary evaluations we found that the fully connected graph produced more reliable parcellation with tNLM, presumably because the fully connected graph contains far more information about functional similarity of nodes/vertices than the sparser spatially constrained graph. The fully connected graph also allows clustering of spatially disjoint nodes/vertices in a single subgraph V_i allowing direct identification of functional networks rather than just producing a cortical parcellation.

2.4 Performance Evaluation

2.4.1 Cortical maps and movies of resting fMRI

We compare the performance of different smoothing approaches using a variety of qualitative and quantitative methods. First we show the resting-state time series as a series of still images and movies of BOLD signal intensities on the cortical surface for each of the following: unfiltered data, data with tNLM filtering and data with LB filtering. The movies are played back at a real-time rate *i.e.* the movies linked in Fig. 1 show 5 minutes of brain activity played back over 5 minutes. We render the movie at 10 frames per second by linearly interpolating the images which are spaced by a TR=720ms. The signal intensities are visualized on a smoothed cortical surface using a colormap with transitions from blue (negative) to white (zero) to red (positive). The movie frames are concatenated and encoded using the H.264 codec to produce an mp4 movie.

2.4.2 Simulation: the effect of smoothing on clustering

We used a simulated data set to investigate the relative performance of the different smoothing methods in presence of functionally distinct areas. We simulated a square patch with four quadrants, each corresponding to a different functional area. The vertices in each quadrant were assigned time series from *in vivo* rfMRI data drawn from a small cortical region in a single subject (Fig. 3). Regions were chosen in visual, motor, default mode and task positive networks to be locally homogeneous based on the mean correlation with vertex neighbors. We applied N-cuts using the graph structure described above to find $K = 4$ and $K = 8$ networks for tNLM, LB and unsmoothed data. The resulting networks provide insight into the impact of the different smoothing methods on identifying networks in the case where ground truth is known.

2.4.3 Cortical networks in *in vivo* data

We studied the effect of smoothing of rfMRI data on cortical parcellation by comparing the classification of cortical networks using N-cuts for different numbers of classes. The N-cuts clustering approach sub-divides all the vertices into K disjoint sets as described above. We assign all vertices which are clustered in the same set a unique label ID and visualize them with a unique color on the cortical surface (so that a clustering result with K classes will have K unique colors).

We repeated clustering for all 40 subjects for a range of values of K and for three values of the tNLM parameter $h = 0.60, 0.72$ and 1.73 . Matching the mutual information in the LB filtered with tNLM results for each value of h we found equivalent values of $t = 2, 4$ and 10 . To enable comparisons across subjects and smoothing methods for a given value of K , we first identify equivalent sub-networks across subjects using the matching method described in the supplemental data, which is based on the Gale-Shapley stable matching algorithm [35].

We also investigated how the boundaries of functional regions changed with the number of networks K . For each subject and value of K we defined a triangle on the tessellated cortical mesh as a ‘boundary’ triangle if its vertices lie in more than one cluster. We then

computed the cumulative boundaries for a single subject as a cortical map where the intensity of each triangle represents the number of times the triangle has been identified as a boundary triangle across the values of $K = 2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 30, 40, 50, 60, 80$. We also computed the population average of these cumulative boundary maps across 40 subjects for both tNLM and LB filtering.

2.4.4 Concordance of N-cuts parcellation with functional localizer studies and Brodmann areas

To evaluate quantitative performance we compared the functional parcellation obtained using N-cuts on rfMRI with functional regions defined in the HCP functional localizer studies and probabilistic Brodmann maps. A good filtering algorithm should produce a parcellation whose regions have high agreement with independently determined functional or cytoarchitectonic regions and borders.

For each subject we created a discrete cortical map using several functional localizer activation studies from the HCP data to represent a set of functional regions. First, we thresholded each task activation map at $p \leq 0.001$ (uncorrected) and merged them into a single cortical map with a unique label for each task. If a vertex has $p \leq 0.001$ in more than one task we assigned the task label corresponding to the most significant activation. We then removed isolated task-label patches which had less than 50 vertices. This task map was generated for each subject and compared against the corresponding individual rfMRI parcellation as described below.

Similarly, a partial Brodmann parcellation was created using discrete labels derived from the probabilistic Brodmann areas transferred to the 32K Conte69 cortical mesh [36]. Again, a unique label is assigned to each Brodmann area on the surface and the resulting map compared to the individual parcellation results. Images of the localizer task and Brodmann label maps are shown in Fig. 8.

We computed the agreement between the N-cuts regions obtained by LB and tNLM filtering to the functional localizer and Brodmann labels by a measure of concordance. Concordance measures the fraction of vertices that agree in two sets of labelings and is described in detail in the supplemental material (sec. S.1.1). Since each of the networks obtained by N-cuts can contain more than one contiguous cluster of vertices, we first subdivided each network into a set of disjoint contiguous clusters to each of which we assign a unique label. Since the localizer and Brodmann maps do not cover the full cortex, we computed concordance from the N-cuts results to these maps rather than vice versa (the measure may not be symmetric in some cases).

3 Results

3.1 Temporal Non-Local Means Smoothing: Maps and Movies

The results of tNLM smoothing and its comparison to LB smoothing are shown in the first two images with detailed descriptions in the legends. In Figure 1 we show the cortical BOLD signal before and after smoothing using LB and tNLMs, in a single snapshot at 02:20 minutes, of a single subject. We provide the links to 3 movies, one for each case, that show

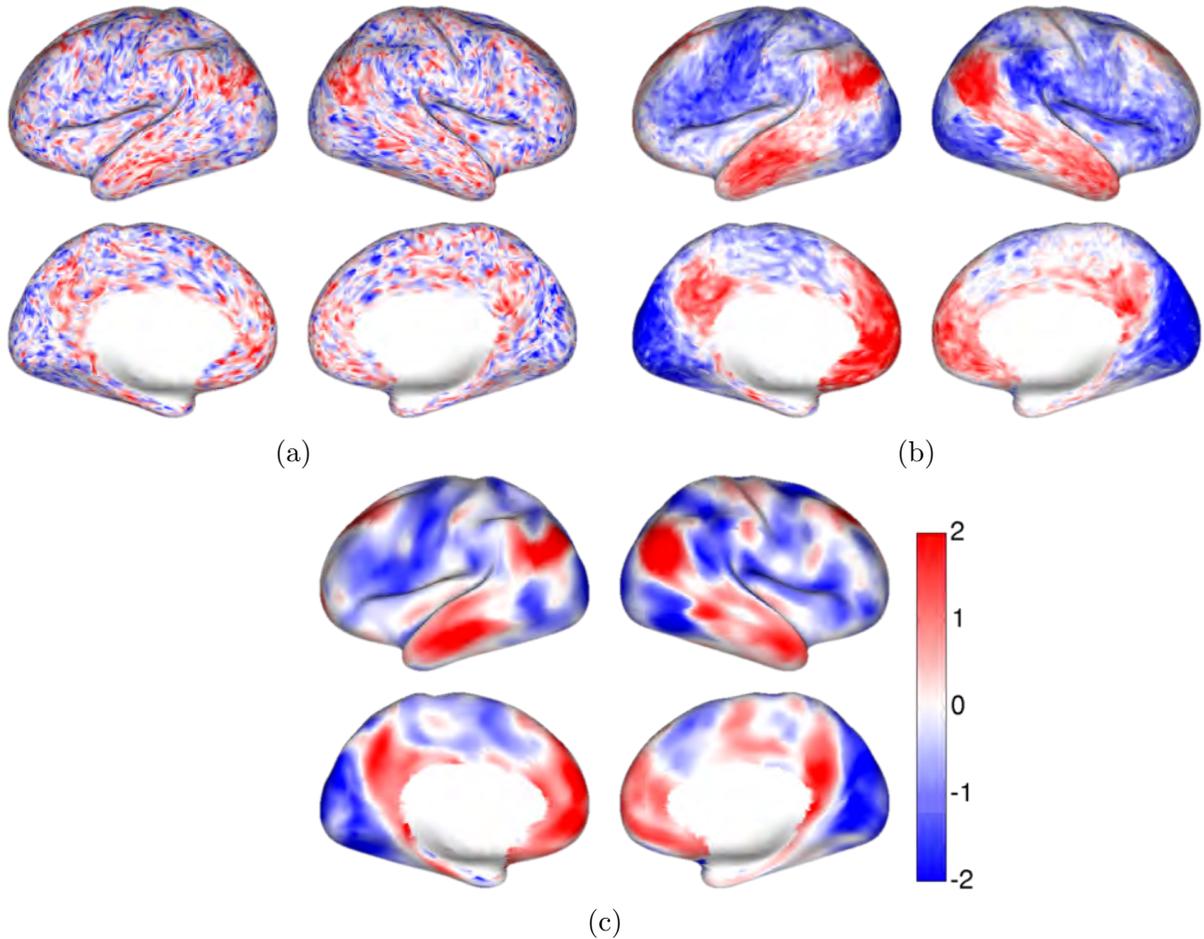


Figure 1: Cortical maps illustrating effects of smoothing on cortical signal intensity in rfMRI in a single subject at a single time point, at 02:20 minutes: (a) no smoothing, (b) LB filtering ($t = 4$) and (c) tNLM filtering ($h = 0.72$). Color scale shows positive (red), negative (blue) and zero (white) BOLD contrast. Five minute real-time movies showing the unsmoothed, LB and tNLM cases can be found here: https://dl.dropboxusercontent.com/u/1659299/work/NLM_videos/video_html5.html. It is difficult to detect spatial structural differences in the original data, even if there are hints that can be discerned. Applying either LB or tNLM filtering, however, the noise is reduced and anatomical coherence in local activation/deactivation with respect to the underlying anatomy of the cerebral cortex is revealed. The tNLM results are clearer in the sense that on the basis of the activity, specific anatomical regions can be identified, e.g. the postero-inferior parietal regions (mostly angular gyri), the dorso-lateral temporal lobes, the left postero-mesial cortex (PMC), and the left mesial prefrontal cortex are readily identified as more active than the mean brain activity. These regions are standard constituent nodes of the default mode network (DMN). Although DMN activity is often described as symmetrical we note that our images reveal asymmetries. The image obtained with LB filtering depicts, in essence, the same list and combination of active regions, although the LB filtering is more patchy and therefore less straightforward to interpret. The differences between the two methods is more readily evident in the 5 minutes movies of continuous resting state recording. Note in particular the different dynamic of the changes in brain activity. The LB images change smoothly from one brain state to the next while the tNLM images depict a more burst like change across consecutive brain states.

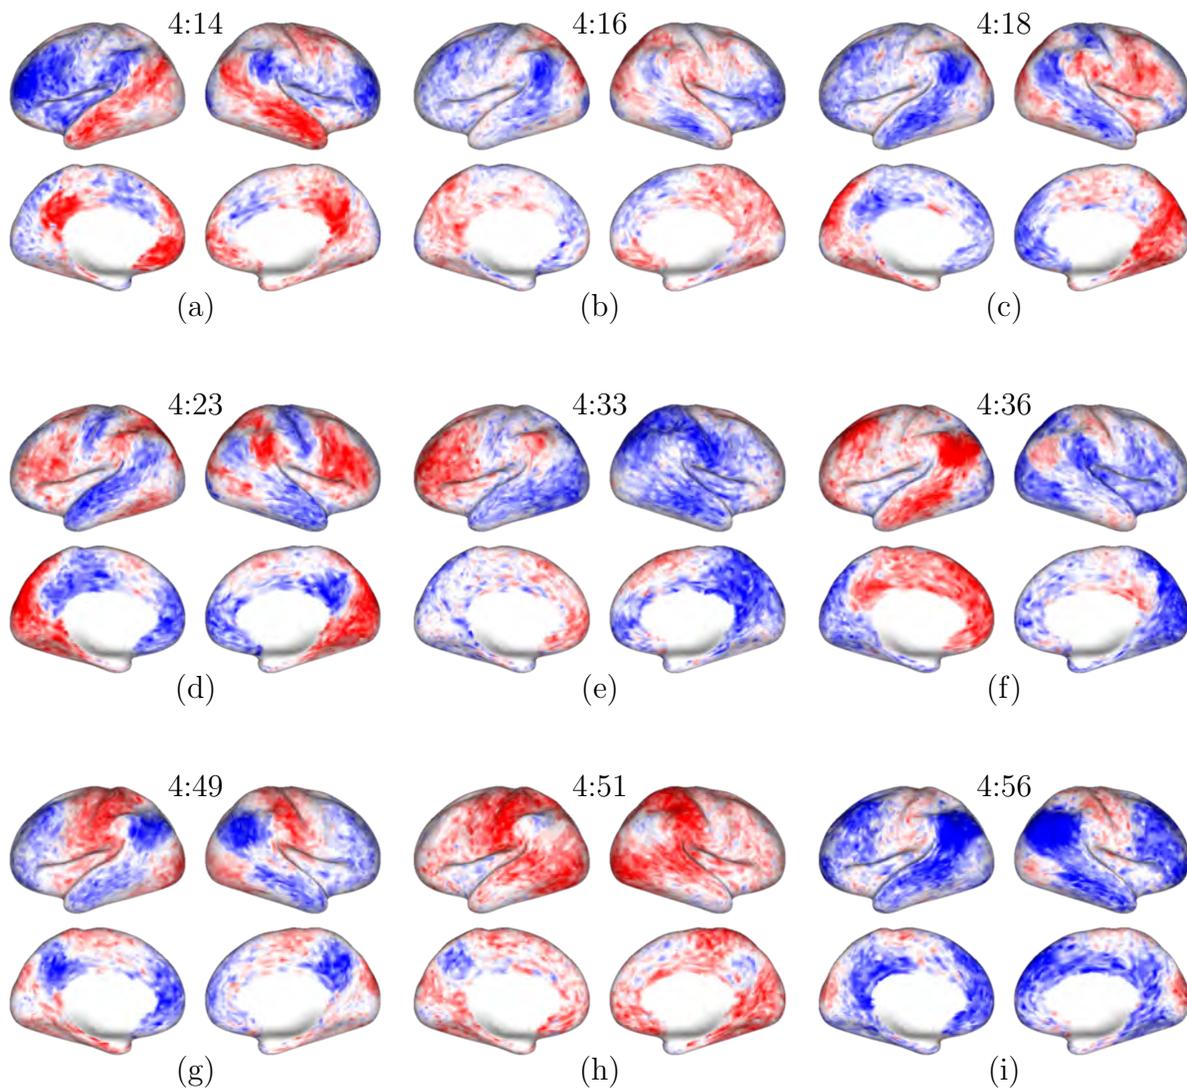


Figure 2: Cortical maps of BOLD intensity from a single subject at different time points with tNLM filtering ($h = 0.72$, same colorscale as fig. 1). The video time (min:sec) is shown above each subfigure. In (a), at 04:20, we see activity above the mean in the DMN nodes, similar to the one seen in Fig. 1 at 02:20, but here more symmetrical. The rest of the brain shows mostly activity below the mean, with the exception of the upper half of the sensory-motor cortices (SMC) which, on the right, show some activity above the mean; (b) only 2 seconds later the lateral temporal and parietal nodes of the DMN show activity clearly below the mean, while the activity in the PMC is still above the mean, but less so, and the activity in the mesial frontal regions is now mostly below the mean; (c) another 2 seconds later, all of the DMN nodes are clearly below the mean, while mesial occipital regions are now above the mean as are the right dorsolateral frontal (DLF) and supra-marginal gyrus (SMG); so is the right insula; (d) five more second have passed and the image is almost the reverse of what was seen in (a); the DMN nodes show clear negativity, as do both SMC, while the rest of the brain, including both insulae, is clearly above the mean; (Contd. on next page)

Figure 2: (Contd. from previous page) (e) after 10 seconds, the activity is clearly asymmetric, with only a small region of the mesial superior frontal above the mean in the right hemisphere; left is close to the mean on the mesial aspect but it is split on the dorsolateral aspect: the frontal lobe, a small area of the SMG and the insula are above the mean, while the remainder of the hemisphere is below the mean; (f) after another 3 seconds the DMN nodes in the left hemisphere are well above the mean, the occipital lobe below, and the insula and SMC close to the mean; the right hemisphere shows an entirely different image: the angular gyrus, PMC and SMC are mildly positive while the rest of the hemisphere is clearly negative; the nodes of the DMN have definite asymmetric activity; (g) at 04:49 the DMN nodes are negative and symmetrically so, while the SMC shows activity clearly above the mean, more so on the left, and so do the occipital lobes and the left insula; (h) three seconds later the brain activity is in general above the mean with three interesting exceptions: the PMC and angular gyri show activity below the mean favoring the left hemisphere, and the left insula is also below the mean; (i) five seconds later, at 04:56, the brain is massively negative with a few exceptions where the activity approximates the mean.

5 minutes of spontaneous brain data. In Fig. 2 we show examples of images obtained with tNML filtering, at different time points. We chose to show a sample of images at 2 seconds and 3 seconds interval, so as to illustrate how the method allows for the moment by moment identification of regions that are varying in their level of activity. Images at the same time points for unfiltered data and LB smoothing, are shown in Fig. S.2 in the supplemental data. The effect of different degrees of smoothing (varying parameter h in tNLM and t in LB) can be found in Fig. S.3. We interpret the irregular levels of activity, the spottiness and burst-like appearance of the anatomical regions involved, separately and in conjunction, as corresponding, probably and intriguingly, to what the actual brain activity during a period of undirected mind-wandering [37, 38] so characteristic of the resting state, might look like.

3.2 Finding Cortical Networks using N-cuts

The simulated map and the results of data clustering into $K = 4$ subgraphs are shown in Fig. 3. The correspondence between the four quadrants of the simulated image and the locations on the cortex from which each quadrant was sampled are shown in Figs. 3(a) and (b). Results of clustering with LB and tNLM filtering are shown in (c) and (d) respectively. The results are almost perfect for $K = 4$ in both cases and, indeed, also were for the unfiltered data since the four regions were chosen to be internally homogeneous with respect to their time series but with a low/negative correlation between the time series in different regions. The case $K = 8$ produces very different results as shown in Fig. 4. The linear mixing across quadrant boundaries using LB produces intermediate regions that internally have a higher correlation than they do with the two regions from which the data originated. The resulting parcellation includes new contiguous regions at the boundaries between quadrants that were not present in the original data and are an artifact of processing (Fig. 4(a), clusters in light-green, yellow and violet). When these clusters are mapped back to the surface vertices on the cortex from which they were drawn, we see they appear distributed across more than one functional area (for example elements in the light-green parcel appear in DMN, visual and

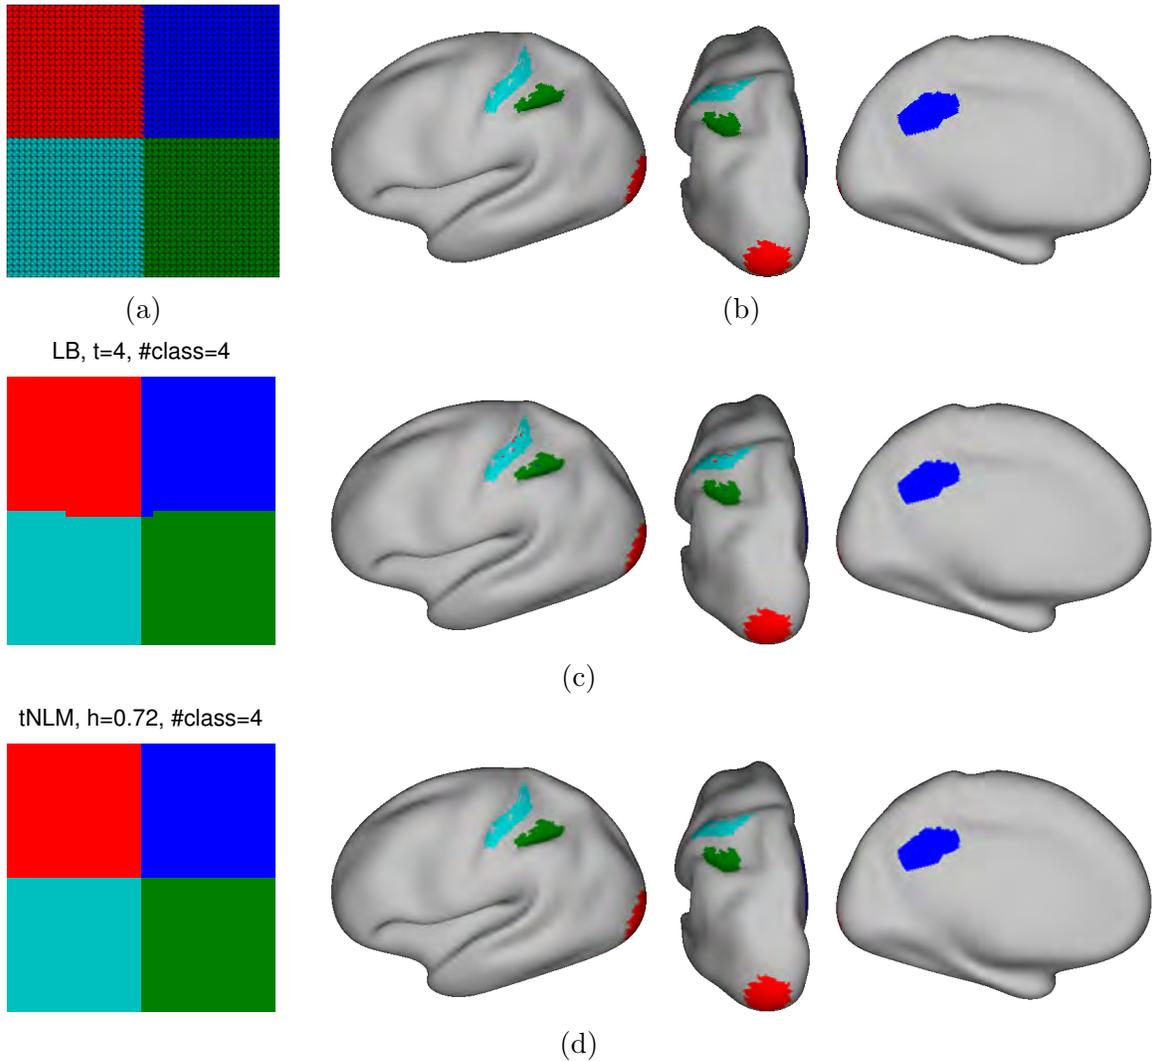


Figure 3: Simulation: (a) four quadrant square surface on which smoothing and parcellation is performed; (b) location of the four regions on cortex from which the time series for each quadrant are taken; (c) N-cut parcellation of LB filtered data into $K = 4$ clusters; the cortical surface image shows each color-coded vertex mapped back from the square to its original location on cortex; (d) N-cut parcellation for tNLM smoothed data into $K = 4$, again with mapping of each color-coded vertex back onto the cortical surface

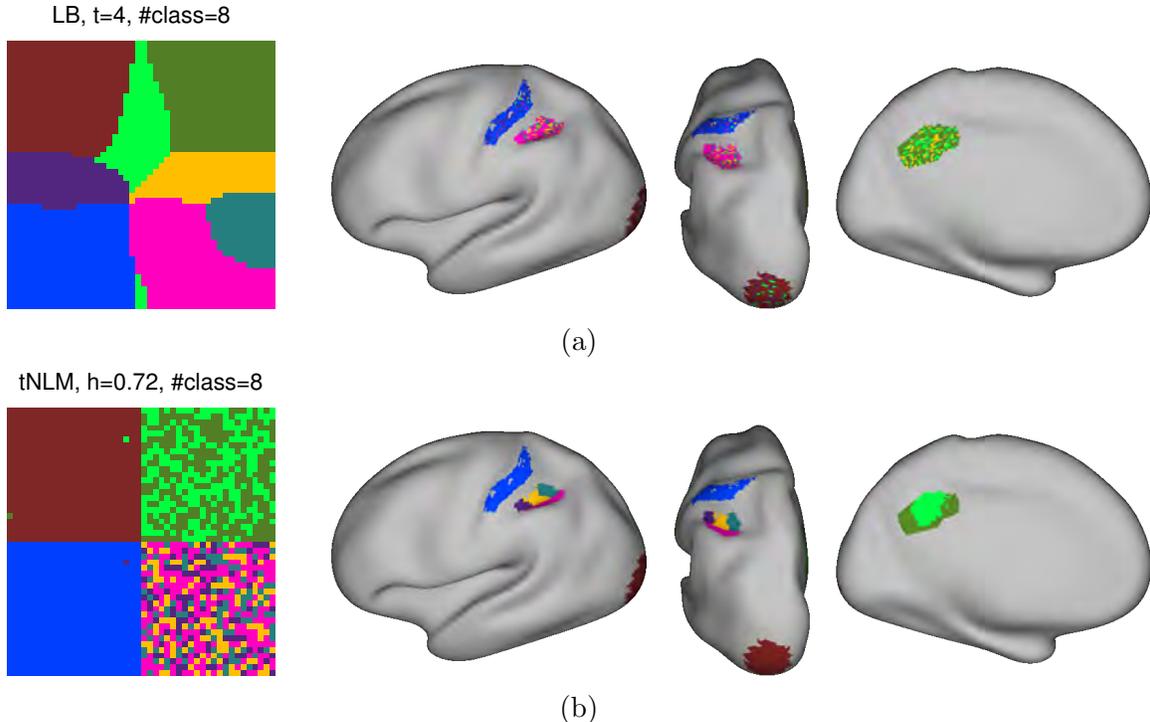


Figure 4: Simulation: result of N-cuts clustering on the square surface with $K = 8$ classes using (a) LB filtering and (b) tNLM filtering. Each subfigure shows the clustering results with a unique color for each cluster on the square surface on the left and the cluster labels mapped back to the cortical surface on the right.

motor areas). This is consistent with the earlier observation that the clusters in the square image that appear in moving from $K = 4$ to $K = 8$ subgraphs are due solely to mixing from LB smoothing and do not reflect true patterns of similarity in the time series prior to filtering.

In contrast, tNLM filtering shows a strong pattern of sub-division of functionally distinct areas when 8 class clustering was used, Fig. 4(b): vertices in the top right quadrant (corresponding to DMN) get sub-divided into two parcels (light and dark green) and vertices in the bottom right quadrant (task positive network, TPN) get sub-divided into four parcels which are distinct from those in the top right quadrant. While the spatial organization of these clusters in the square image appears random, when they are mapped back to the area on the cortex, from which they were drawn, we see that the clustering result actually sub-parcellates the DFM and TPN. This result demonstrates the “non-local” nature of tNLM - smoothing is performed based on similarity in time series rather than spatial proximity. For this reason, subsequent partitioning based on a fully connected graph can identify groups of pixels with similar time-series in the original data rather than producing false parcels as a result of local linear mixing of signals from adjacent functional areas as occurs with LB filtering.

We do not include results for unfiltered data since they are qualitatively very similar to those obtained using tNLM for both $K = 4$ and $K = 8$. Since each quadrant was selected

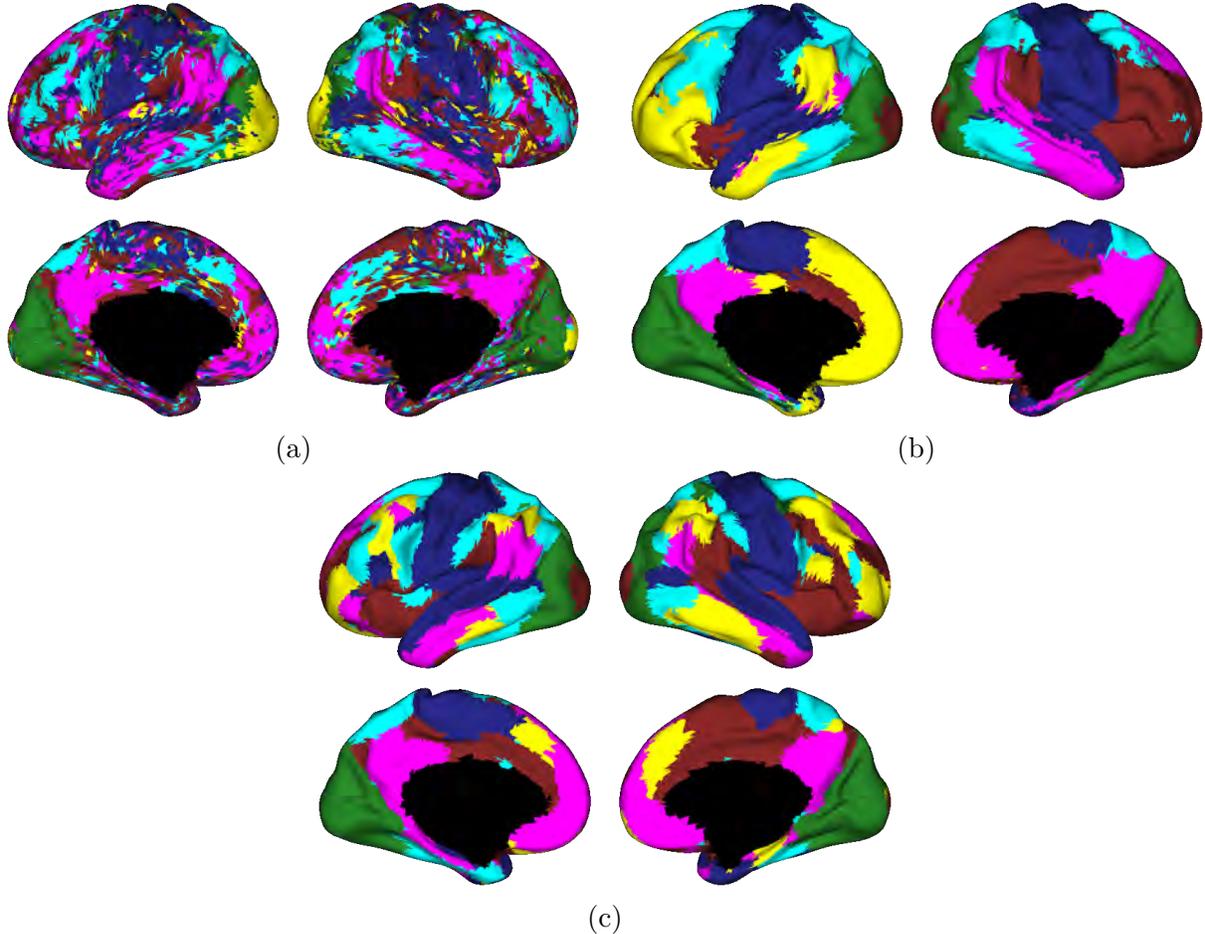


Figure 5: Cortical parcellations using N-cuts on a fully connected cortical surface graph for a single subject to partition cortex into $K = 6$ networks: (a) unfiltered data, (b) tNLM filtering ($h=0.72$), (c) LB filtering ($t=4$). In each case a distinct color represents one of the $K = 6$ networks.

to have clearly distinct time series from all other quadrants, the N-cuts algorithm, even without smoothing, was able to reliably partition the data. That unfiltered data produced very similar subparcellations for $K = 8$ as shown in fig. 4(b) indicates that there is evidence for these subparcellations in the data and these results are not an artifact of the nonlinear tNLM smoothing. As we show next, this similarity between tNLM and unfiltered data does not occur when N-cuts is applied to a complete set of *in vivo* data.

Next we examine the effect of LB and tNLM filtering on N-cuts based parcellation of an *in vivo* dataset for a single subject. Figure 5 shows the result of N-cut clustering with six classes. Clusters obtained from the original unfiltered data, Fig. 5(a), yields default mode (pink), visual (yellow/green), and somatomotor (dark blue) networks. However, the clusters are noisy and disjointed. In contrast, tNLM smoothed data shows more jointed networks with large contiguous regions, Fig. 5(b). In addition to the networks identified in the original data we can also identify the visual cortices and the ventral visual stream (dark green) separate from the occipital pole (dark red). We can also identify the fronto-parietal

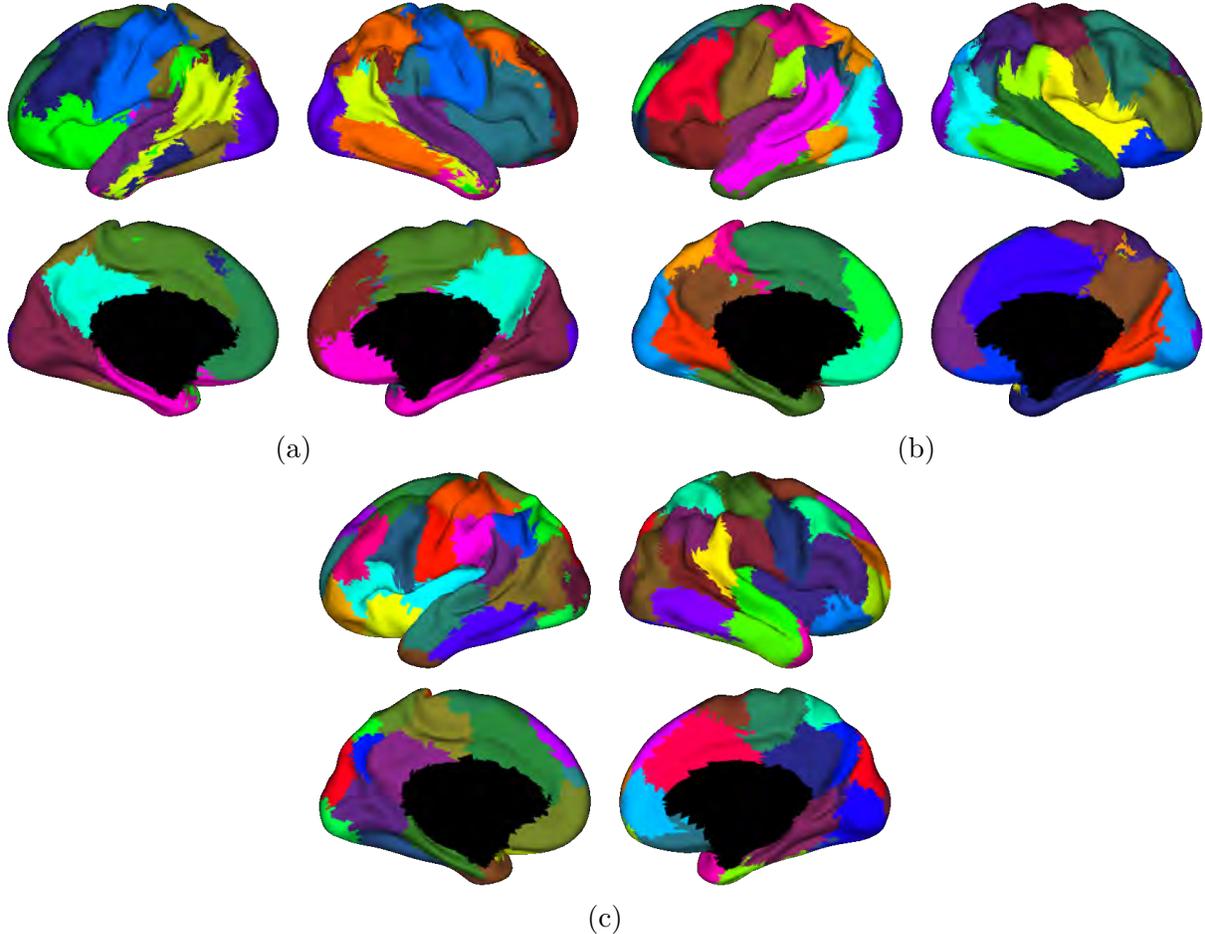
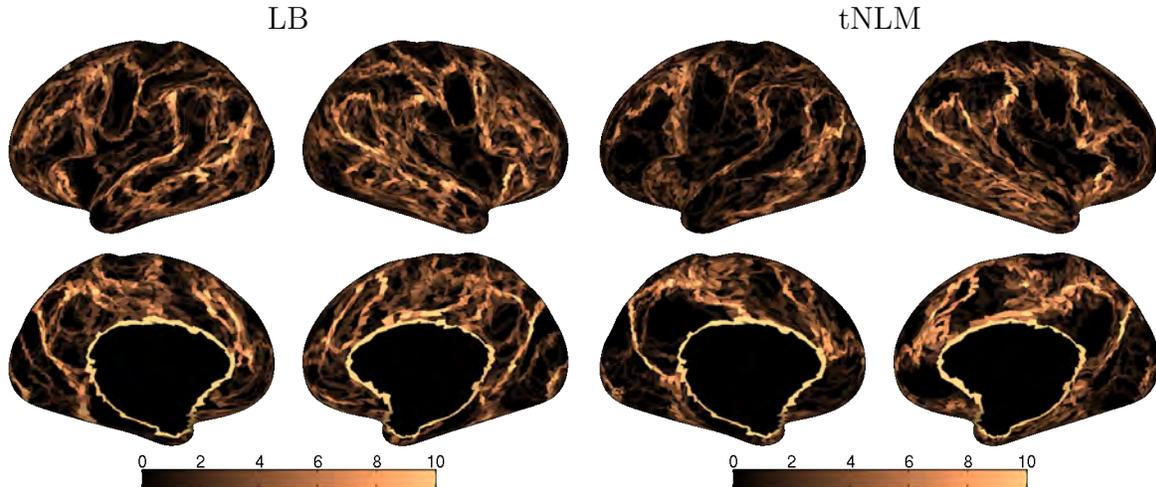


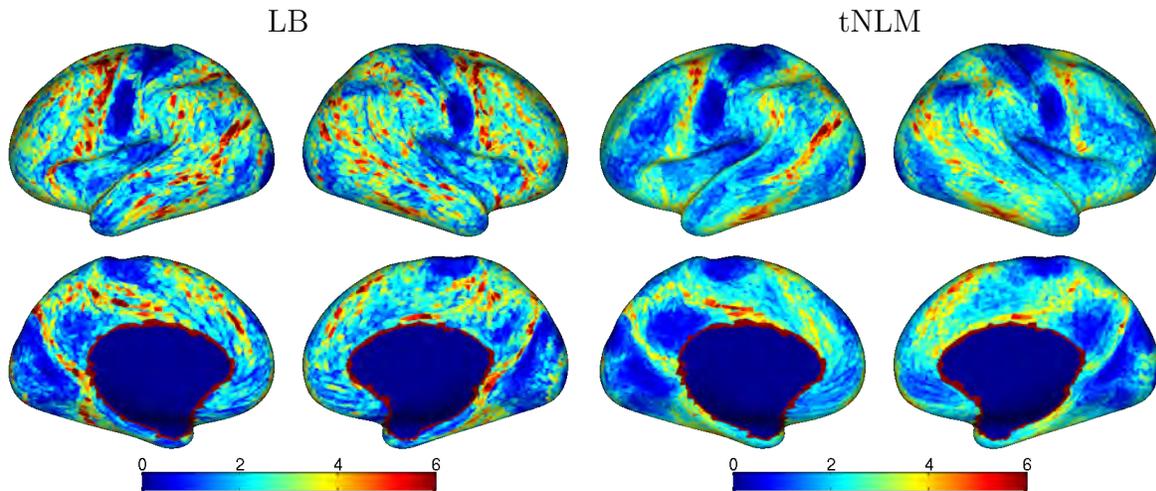
Figure 6: Cortical parcellation N-cuts applied to tNLM filtered ($h=0.72$) data for the same subject as in Fig. 5 for (a) $K = 15$, (b) $K = 30$ and (c) $K = 60$ clusters. See supplemental data for equivalent images for unfiltered and LB filtered data.

network, described to anticorrelate to the default mode network (light blue), which includes the frontal eye field, left middle frontal gyrus, superior parietal lobule, and the lateral-posterior regions of the temporal lobe [39–43]. LB filtering, Fig. 5(c), also produces similar networks as identified by tNLM, however some of the networks separate into patchy clusters. The greater number of patchy regions in the LB result is consistent with the formation of additional false parcels shown in the simulation in the previous section, Fig. 3(c), that result from local linear mixing of two or more functional areas as a result of the smoothing.

Expanding on the previous clustering result, we show clusters in Fig. 6 with tNLM filtering on one subject with (a) $K = 15$, (b) $K = 30$, and (c) $K = 60$ (equivalent results for unfiltered and LB filtered data are included in the supplemental data Fig. S.1). As the number of classes are increased, boundaries are preserved and regions are sub-divided. For example, the somatosensory and motor cortices are initially identified as a single network (blue) when the brain is allowed to segment into 15 classes (Fig. 6(a)). When the classes are doubled, this area sub-divides into the right upper (violet), and left upper (pink), and the lower (brown) somatomotor cortices (Fig. 6(b)). Increasing the number of classes to 60



(a)



(b)

Figure 7: Cortical map of the cumulative boundaries of N-cut parcellations over fifteen different values of K (a) in a single subject and (b) averaged across population of 40 subjects. The value at each triangle represents total number of times that triangle was a boundary triangle across fifteen different clustering results ($K = 2, 3, 4, 5, 6, 7, 8, 9, 10, 15, 30, 40, 50, 60, 80$). The boundary map is thresholded at an upper boundary count of 10 for single subject and 6 for the population average.

segments, the right lower somatomotor cortex (red) separates from the left hemisphere, and further divides into the ventral premotor cortex (blue) and the ventral motor cortex (dark red) (Fig. 6(c)). Similar patterns of progressive subdivision can be observed in other primary networks including the visual network.

In Fig. 7 we illustrate how the boundaries of the clusters vary across several parcellations with different number of classes in both tNLM and LB. Fig. 7(a) shows cumulative boundaries for a single subject over different number of classes. The tNLM results show consistent boundaries delineating the ventro-medial prefrontal cortex (vmPCC), postero-medial cor-

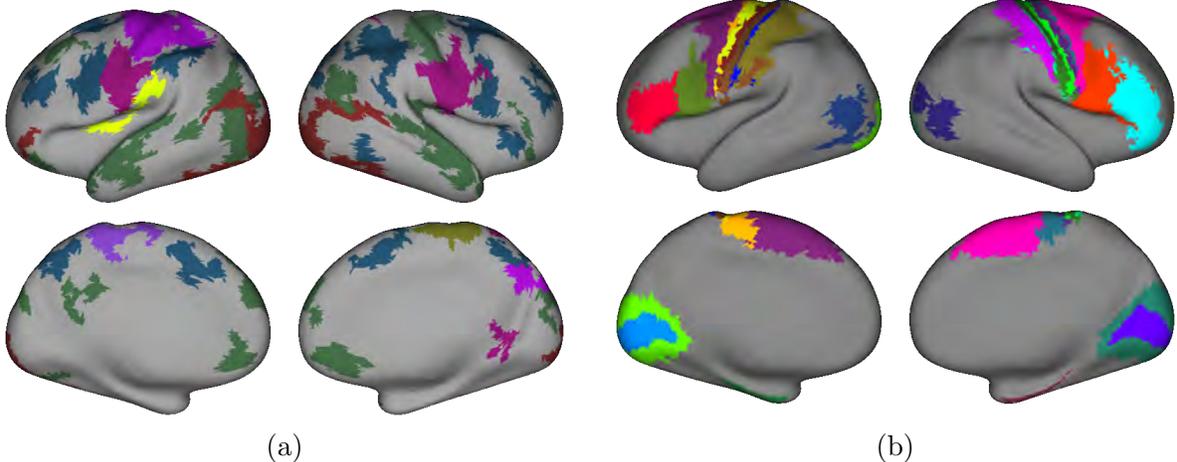


Figure 8: Cortical maps of the label sets used for quantitative analysis: (a) functional localizers for a single subject [30, 31] and (b) discrete Brodmann’s labels from a probabilistic atlas [36]. Each color represents a different Brodmann or functional area.

tex (PMC) and the visual cortex. The LB results are quite similar. But when averaged across the 40 subjects, Fig. 7(b), there are clear differences with boundaries occurring more consistently across different values of K with tNLM than with LB. In both cases the upper and lower sensorimotor areas are consistently identified, as evidenced through the absence of boundaries in these regions. However, tNLM shows clearer boundaries than LB, particularly for the PMC, the visual cortex and the ventrolateral prefrontal cortex. LB shows boundaries spread across frontal and lateral posterior temporal-occipital region, possibly reflecting the introduction of false region boundaries between functional areas as illustrated in our earlier simulation. High order association cortices also clearly show marked internal consistency (i.e. lack of edges) in tNLM results.

Next we present a quantitative analysis of the concordance between parcellations obtained through N-cuts clustering of resting data and functional localizers and Brodmann areas (BA). An image of the functional localizer and Brodmann maps in shown in Fig. 8 for a single subject. We used different level of LB filtering ($t = 2, 4, 10$) and tNLM filtering ($h = 0.60, 0.72, 1.73$) before using N-cuts clustering with a large number of classes ($3 \leq K \leq 400$). Fig. 9 shows the median concordance of clustering results with functional localizers over a population of 40 subjects (see sec. S.1.1 for details of the concordance measure). Both LB and tNLM filtering shows a marked improvement in concordance as compared to the original unfiltered data. However, tNLM consistently shows higher median concordance than LB across a large range of number of classes K . Fig. 10 shows concordance with Brodmann area maps. Similarly to the functional localizers, we see a large improvement in concordance of clusters with smoothing as compared to the original data without any smoothing. LB and tNLM show similar behavior for smaller number of classes ($K < 50$), however tNLM shows robust performance over parcellations with larger number classes.

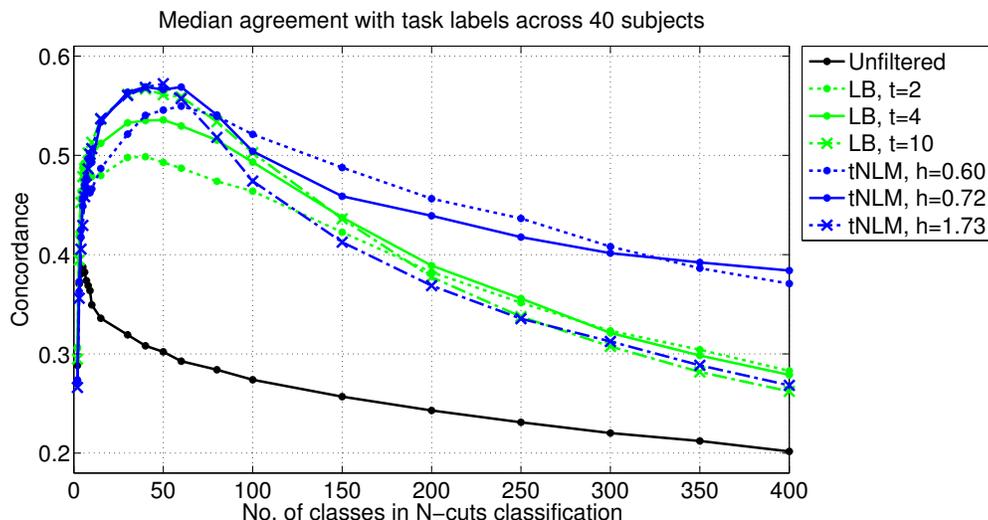


Figure 9: Concordance between task localizer maps and N-cuts parcellations with LB and tNLM filtering as a function of the number of classes K . For every set of smoothing and parcellation parameters, we show the median concordance over a population of 40 subjects (4 resting state sessions for each subject).

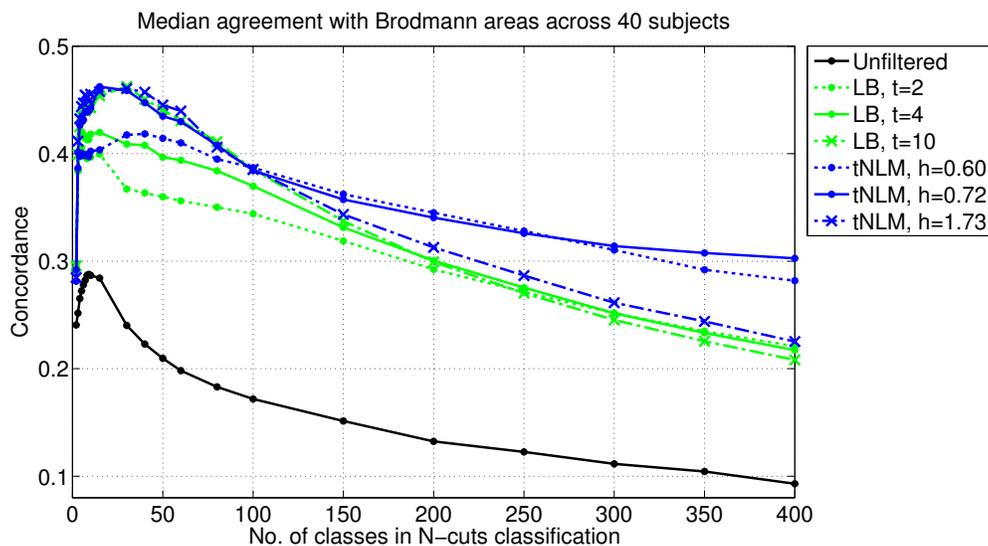


Figure 10: Concordance between 26 Brodmann areas and N-cuts parcellations with LB and tNLM filtering as a function of the number of classes K . For every set of smoothing and parcellation parameters, we show the median concordance over a population of 40 subjects (4 resting state sessions for each subject).

4 Discussion

The tNLM filtering results shown in Fig. 1 appear qualitatively quite different from both the unfiltered noisy data and the result of linear Laplace-Beltrami smoothing. While the images are smoother in the latter, the regions appearing coherently activated are more patchy. The

image obtained with tNLM filtering allows more direct identification of components of the default mode network. In Fig. 2 a sequence of images at short intervals shows clearly how coherent co-activation changes for other networks. As activity in one network increases and that in an adjacent network decreases, the effect of LB smoothing across the boundary between these networks produces an apparent boundary or wave moving from one network to the other. This behavior is apparent in the LB movie. With tNLM, because only surface elements with similar time series are averaged to denoise the data, when two adjacent networks have distinct time series then they will not be blurred through smoothing and this wavelike effect does not occur. The impression is that of irregular, burst like activity, which is to be expected during a period of undirected mind-wandering, characteristic of the resting state. Qualitatively, there are clearer boundaries in tNLM relative to LB, although it is also evident that even within networks, activation is not always synchronous. The ability of tNLM to denoise data while respecting functional boundaries makes it a potentially attractive processing tool both for data exploration and as a precursor to network identification or parcellation.

There are relatively few examples showing real time rfMRI whole-brain activity in the form of either single frame images or movies in the literature. An early example from Vincent et al. [44, 45] shows real-time brain dynamics from data used to explore resting networks [40, 46] however the dynamics do not clearly show coherent activity in, for instance, regions of the default mode networks or other networks. Kundu et al. [47, 48] have developed a method for denoising multiecho fMRI data that distinguishes BOLD from non-BOLD signals based on echo-time dependence. The resulting denoised data show dynamic activation in DMN and other networks in realtime [49] similar to that shown here. However, their approach explicitly requires a multi-echo sequence and unlike tNLM cannot be applied to standard fMRI protocols. The dynamics shown in our movies also share similarities with those in Zalesky et al. [50] who explore the HCP resting data using dynamic regional network efficiency measures, computed from time-resolved connectivity estimates, to produce movies of brain dynamics. Their results show “a consistent set of functional connections (with) pronounced fluctuations in their strength over time”. The authors also note spontaneous increases in spatially distributed regions over brief intervals, observations that can also be made from our movies and of which a sample is shown in Fig. 2.

While all examples show here are restricted to cortex, tNLM filtering can also be applied to volumetric or grayordinate [6] representations of the data. Similarly, tNLM could be applied as a denoising tool in event-related functional MRI studies which may result in improved resolution of focal activation relative to methods based on conventional isotropic linear smoothing. The tNLM method could also be extended to include a spatial component, so that the weighted average depends on a combination of temporal and spatial similarity. It would also be interesting to explore a dynamic version in which the similarity measure is computed of restricted time-window, rather than the entire time-course as was the case in the results presented above.

We illustrated the potential utility of tNLM through cortical parcellation studies based on N-cuts spectral graph partitioning. Our approach parcellates a single subject using a fully connected graph with edge strengths based on pairwise correlations between the time series on the surface elements. Most previous applications of graph-cuts in brain parcellation have used locally connected graphs to ensure spatially contiguous parcels [3, 4]. While

this appears necessary for unfiltered data (Fig. 5(a)), the denoising effect of tNLM allows use of the fully connected graph while still producing a piecewise contiguous parcellation (Fig. 5(b)). This has the advantage of using all correlations for parcellation rather than a restricted local set which, based on our comparison of parcellations using locally and globally connected graphs (not shown), should produce more reliable results. LB smoothed data can also be parcellated using the fully connected graph (Fig. 5(c)) but the resulting networks appear to contain regions that may not actually belong. We explored this issue further in the simulations shown in Figs. 3 and 4. LB smoothing applied to the four quadrant configuration in Fig. 3(a) produced mixing of signals across quadrant boundaries that it turned out resulted in artifactual networks that arise purely as a result of filtering. This finding indicates that care should be taken when using linear smoothing in combination with parcellation methods based on pairwise correlations to ensure that parcels are not produced solely as an artifact of smoothing.

Changes in parcellation as a function of the number of networks was investigated in Fig. 6 for tNLM (with equivalent results in Fig. S.1 for LB and unfiltered data). A well known problem with N-cuts is that the algorithm tends to produce cuts of similar size if the graph (and its adjacency matrix) does not contain sufficient information to unambiguously support a single K -way partition. We note that this does seem to be the case even for the relatively large number of networks ($K = 60$) as shown in Fig. 6. To investigate consistency as a function of the number of networks, we averaged edge locations over multiple values of K in Fig. 7. It is interesting that the resulting individual edge maps (Fig. 7(a)) bear a strong resemblance to the group functional connectivity gradients shown in Fig. 10 in [5]. Comparisons with functional localizer (Fig. 9) and Brodmann (Fig. 10) maps, as a function of K , show similar peak performance at approximately $K = 45$ for the localizers and $K = 30$ for Brodmann maps for LB and tNLM. This is somewhat surprising given the qualitative differences between the two results in Figs. 6, 7, and S.1. The best results for LB are for the most heavily smoothed data ($t = 10$) while for tNLM best results are for $h = 0.72$. The value $t = 10$ produces significantly more smoothing than shown in Figs. 1 and S.2 and appears to remove much of the spatial detail in the data as shown in Fig. S.3. As K is increased tNLM retains consistently higher concordance than LB indicating more reliable parcellations with a larger number of networks.

5 Conclusion

The results shown above support the primary claim of this report: that temporal Non Local Means filtering (tNLM) is able to denoise resting fMRI data while also retaining much of the spatial structure that reflects ongoing dynamic brain activity. Correlated variations in activity are directly visible in the movies of cortical activity, which appear to reflect the underlying dynamics of large-scale brain networks. This ability to visualize real-time whole-brain networks may facilitate exploratory data analysis leading to new insights into the dynamics of spontaneous brain activity. Temporal NLM can also be used as a preprocessor for resting fMRI for quantitative exploration of dynamic brain networks. Examples using the HCP resting fMRI data set indicate the advantages of tNLM over linear Laplace-Beltrami smoothing on the cortical surface for identifying networks and cortical parcellation.

6 Acknowledgments

This work is supported by grants R01 NS089212, R01 EB009048 and R01 NS074980. Data were provided [in part] by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

- [1] B. Biswal, F. Zerrin Yetkin, V. M. Haughton, and J. S. Hyde, “Functional connectivity in the motor cortex of resting human brain using echo-planar MRI,” *Magnetic Resonance in Medicine*, vol. 34, no. 4, pp. 537–541, 1995.
- [2] S. M. Smith, D. Vidaurre, C. F. Beckmann, M. F. Glasser, M. Jenkinson, K. L. Miller, T. E. Nichols, E. C. Robinson, G. Salimi-Khorshidi, M. W. Woolrich, D. M. Barch, K. Ugurbil, and D. C. Van Essen, “Functional connectomics from resting-state fMRI,” *Trends in Cognitive Sciences*, vol. 17, no. 12, pp. 666 – 682, 2013.
- [3] T. Blumensath, S. Jbabdi, M. F. Glasser, D. C. Van Essen, K. Ugurbil, T. E. Behrens, and S. M. Smith, “Spatially constrained hierarchical parcellation of the brain with resting-state fMRI,” *NeuroImage*, vol. 76, no. 0, pp. 313 – 324, 2013.
- [4] R. C. Craddock, G. A. James, P. E. Holtzheimer, X. P. Hu, and H. S. Mayberg, “A whole brain fMRI atlas generated via spatially constrained spectral clustering,” *Human Brain Mapping*, vol. 33, no. 8, pp. 1914–1928, 2012.
- [5] S. M. Smith, C. F. Beckmann, J. Andersson, E. J. Auerbach, J. Bijsterbosch, G. Douaud, E. Duff, D. A. Feinberg, L. Griffanti, M. P. Harms, M. Kelly, T. Laumann, K. L. Miller, S. Moeller, S. Petersen, J. Power, G. Salimi-Khorshidi, A. Z. Snyder, A. T. Vu, M. W. Woolrich, J. Xu, E. Yacoub, K. Ugurbil, D. C. Van Essen, and M. F. Glasser, “Resting-state fMRI in the Human Connectome Project,” *NeuroImage*, vol. 80, no. 0, pp. 144 – 168, 2013.
- [6] M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni, D. C. Van Essen, and M. Jenkinson, “The minimal preprocessing pipelines for the Human Connectome Project,” *NeuroImage*, vol. 80, pp. 105–124, 2013.
- [7] J. B. Hopfinger, C. Büchel, A. P. Holmes, and K. J. Friston, “A study of analysis parameters that influence the sensitivity of event-related fMRI analyses,” *NeuroImage*, vol. 11, no. 4, pp. 326 – 333, 2000.
- [8] G. S. Wig, T. O. Laumann, and S. E. Petersen, “An approach for parcellating human cortical areas using resting-state correlations,” *NeuroImage*, vol. 93, Part 2, pp. 276 – 291, 2014.

- [9] G. S. Wig, T. O. Laumann, A. L. Cohen, J. D. Power, S. M. Nelson, M. F. Glasser, F. M. Miezin, A. Z. Snyder, B. L. Schlaggar, and S. E. Petersen, “Parcellating an individual subject’s cortical and subcortical brain structures using snowball sampling of resting-state correlations,” *Cerebral Cortex*, vol. 24, no. 8, pp. 2036–2054, 2014.
- [10] A. L. Cohen, D. A. Fair, N. U. Dosenbach, F. M. Miezin, D. Dierker, D. C. Van Essen, B. L. Schlaggar, and S. E. Petersen, “Defining functional areas in individual human brains using resting functional connectivity MRI,” *NeuroImage*, vol. 41, no. 1, pp. 45 – 57, 2008.
- [11] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis, “On the Laplace-Beltrami operator and brain surface flattening,” *IEEE Transactions on Medical Imaging*, vol. 18, pp. 700–711, Aug 1999.
- [12] P. Perona and J. Malik, “Scale-space and edge detection using anisotropic diffusion,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629–639, 1990.
- [13] A. A. Joshi, D. W. Shattuck, P. M. Thompson, and R. M. Leahy, “A parameterization-based numerical method for isotropic and anisotropic diffusion smoothing on non-flat surfaces,” *IEEE Transactions on Image Processing*, vol. 18, pp. 1358–1365, June 2009.
- [14] A. Buades, B. Coll, and J.-M. Morel, “A non-local algorithm for image denoising,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 60–65, June 2005.
- [15] J. V. Manjón, J. Carbonell-Caballero, J. J. Lull, G. García-Martí, L. Martí-Bonmatí, and M. Robles, “MRI denoising using non-local means,” *Medical Image Analysis*, vol. 12, no. 4, pp. 514 – 523, 2008.
- [16] P. Coupé, P. Yger, S. Prima, P. Hellier, C. Kervrann, and C. Barillot, “An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images,” *IEEE Transactions on Medical Imaging*, vol. 27, pp. 425–441, April 2008.
- [17] J. V. Manjón, N. A. Thacker, J. J. Lull, G. Garcia-Martí, L. Martí-Bonmatí, and M. Robles, “Multicomponent MR image denoising,” *International Journal of Biomedical Imaging*, vol. 2009, pp. 1–10, 2009.
- [18] X.-N. Zuo, T. Xu, L. Jiang, Z. Yang, X.-Y. Cao, Y. He, Y.-F. Zang, F. X. Castellanos, and M. P. Milham, “Toward reliable characterization of functional homogeneity in the human brain: Preprocessing, scan duration, imaging resolution and computational space,” *NeuroImage*, vol. 65, pp. 374 – 386, 2013.
- [19] X.-X. Xing, T. Xu, Z. Yang, and X.-N. Zuo, “Non-local means smoothing: A demonstration on multiband R-fMRI,” in *19th Annual Meeting of the Organization for Human Brain Mapping (OHBM)*, p. 2030, 2013.

- [20] X.-N. Zuo and X.-X. Xing, “Effects of non-local diffusion on structural MRI preprocessing and default network mapping: Statistical comparisons with isotropic/anisotropic diffusion,” *PLoS ONE*, vol. 6, p. e26703, 10 2011.
- [21] M. Bernier, M. Chamberland, J.-C. Houde, M. Descoteaux, and K. Whittingstall, “Using fMRI non-local means denoising to uncover activation in sub-cortical structures at 1.5 T for guided HARDI tractography,” *Frontiers in Human Neuroscience*, vol. 8, no. 715, 2014.
- [22] S. St-Jean, P. Coupé, and M. Descoteaux, “Non local spatial and angular matching: a new denoising technique for diffusion MRI,” in *Joint Annual Meeting ISMRM-ESMRMB 2014*, p. 758, 2014.
- [23] M. Descoteaux, N. Wiest-Daesslé, S. Prima, C. Barillot, and R. Deriche, “Impact of rician adapted non-local means filtering on HARDI,” in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (D. Metaxas, L. Axel, G. Fichtinger, and G. Szekely, eds.), vol. 5242 of *Lecture Notes in Computer Science*, pp. 122–130, Springer Berlin Heidelberg, 2008.
- [24] N. Wiest-Daesslé, S. Prima, P. Coupé, S. Morrissey, and C. Barillot, “Rician noise removal by non-local means filtering for low signal-to-noise ratio MRI: Applications to DT-MRI,” in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (D. Metaxas, L. Axel, G. Fichtinger, and G. Szekely, eds.), vol. 5242 of *Lecture Notes in Computer Science*, pp. 171–179, Springer Berlin Heidelberg, 2008.
- [25] J. Dutta, R. M. Leahy, and Q. Li, “Non-local means denoising of dynamic PET images,” *PLoS ONE*, vol. 8, p. e81390, 12 2013.
- [26] B. Thirion, G. Varoquaux, E. Dohmatob, and J.-B. Poline, “Which fMRI clustering gives good brain parcellations?,” *Frontiers in Neuroscience*, vol. 8, no. 167, 2014.
- [27] S. X. Yu and J. Shi, “Multiclass spectral clustering,” in *Ninth IEEE International Conference on Computer Vision (ICCV)*, vol. 1, pp. 313–319, Oct 2003.
- [28] D. C. Van Essen, S. M. Smith, D. M. Barch, T. E. Behrens, E. Yacoub, and K. Ugurbil, “The WU-Minn Human Connectome Project: An overview,” *NeuroImage*, vol. 80, no. 0, pp. 62 – 79, 2013.
- [29] D. C. Van Essen, M. F. Glasser, D. L. Dierker, J. Harwell, and T. Coalson, “Parcellations and hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases,” *Cerebral Cortex*, vol. 22, no. 10, pp. 2241–2262, 2012.
- [30] D. M. Barch, G. C. Burgess, M. P. Harms, S. E. Petersen, B. L. Schlaggar, M. Corbetta, M. F. Glasser, S. Curtiss, S. Dixit, C. Feldt, D. Nolan, E. Bryant, T. Hartley, O. Footer, J. M. Bjork, R. Poldrack, S. Smith, H. Johansen-Berg, A. Z. Snyder, and D. C. Van Essen, “Function in the human connectome: Task-fMRI and individual differences in behavior,” *NeuroImage*, vol. 80, pp. 169 – 189, 2013.

- [31] M. W. Woolrich, B. D. Ripley, M. Brady, and S. M. Smith, “Temporal autocorrelation in univariate linear modeling of fMRI data,” *NeuroImage*, vol. 14, no. 6, pp. 1370 – 1386, 2001.
- [32] S. Seo, M. K. Chung, and H. K. Vorperian, “Heat kernel smoothing using Laplace-Beltrami eigenfunctions,” in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)* (T. Jiang, N. Navab, J. Pluim, and M. Viergever, eds.), vol. 6363 of *Lecture Notes in Computer Science*, pp. 505–512, Springer Berlin Heidelberg, 2010.
- [33] S. Seo and M. K. Chung, “Laplace-Beltrami eigenfunction expansion of cortical manifolds,” in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 372–375, March 2011.
- [34] J. Shi and J. Malik, “Normalized cuts and image segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, pp. 888–905, Aug 2000.
- [35] D. Gale and L. S. Shapley, “College admissions and the stability of marriage,” *American mathematical monthly*, vol. 69, pp. 9–15, Jan 1962.
- [36] B. Fischl, N. Rajendran, E. Busa, J. Augustinack, O. Hinds, B. T. Yeo, H. Mohlberg, K. Amunts, and K. Zilles, “Cortical folding patterns and predicting cytoarchitecture,” *Cerebral Cortex*, vol. 18, no. 8, pp. 1973–1980, 2008.
- [37] J. S. Antrobus, J. L. Singer, S. Goldstein, and M. Fortgang, “Mind-wandering and cognitive structure,” *Transactions of the New York Academy of Sciences*, vol. 32, no. 2, pp. 242–252, 1970.
- [38] R. L. Buckner and J. L. Vincent, “Unrest at rest: Default activity and spontaneous network correlations,” *NeuroImage*, vol. 37, pp. 1091–1096, 2007.
- [39] M. D. Fox, A. Z. Snyder, J. L. Vincent, M. Corbetta, D. C. Van Essen, and M. E. Raichle, “The human brain is intrinsically organized into dynamic, anticorrelated functional networks,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 27, pp. 9673–9678, 2005.
- [40] M. D. Fox, M. Corbetta, A. Z. Snyder, J. L. Vincent, and M. E. Raichle, “Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems,” *Proceedings of the National Academy of Sciences*, vol. 103, no. 26, pp. 10046–10051, 2006.
- [41] J. L. Vincent, I. Kahn, A. Z. Snyder, M. E. Raichle, and R. L. Buckner, “Evidence for a frontoparietal control system revealed by intrinsic functional connectivity,” *Journal of Neurophysiology*, vol. 100, no. 6, pp. 3328–3342, 2008.
- [42] P. Fransson, “Spontaneous low-frequency BOLD signal fluctuations: An fMRI investigation of the resting-state default mode of brain function hypothesis,” *Human Brain Mapping*, vol. 26, no. 1, pp. 15–29, 2005.
- [43] R. L. Buckner, J. R. Andrews-Hanna, and D. L. Schacter, “The brain’s default network,” *Annals of the New York Academy of Sciences*, vol. 1124, no. 1, pp. 1–38, 2008.

- [44] J. L. Vincent, A. Z. Snyder, M. D. Fox, B. J. Shannon, J. R. Andrews, M. E. Raichle, and R. L. Buckner, “Spontaneous brain activity (I).” Video published on YouTube.com, July 2007.
- [45] M. D. Fox, M. Corbetta, A. Z. Snyder, J. L. Vincent, and M. E. Raichle, “Spontaneous brain activity (II).” Video published on YouTube.com, July 2007.
- [46] J. L. Vincent, A. Z. Snyder, M. D. Fox, B. J. Shannon, J. R. Andrews, M. E. Raichle, and R. L. Buckner, “Coherent spontaneous activity identifies a hippocampal-parietal memory network,” *Journal of Neurophysiology*, vol. 96, no. 6, pp. 3517–3531, 2006.
- [47] P. Kundu, N. D. Brenowitz, V. Voon, Y. Worbe, P. E. Vértes, S. J. Inati, Z. S. Saad, P. A. Bandettini, and E. T. Bullmore, “Integrated strategy for improving functional connectivity mapping using multiecho fMRI,” *Proceedings of the National Academy of Sciences*, vol. 110, no. 40, pp. 16187–16192, 2013.
- [48] P. Kundu, S. J. Inati, J. W. Evans, W.-M. Luh, and P. A. Bandettini, “Differentiating BOLD and non-BOLD signals in fMRI time series using multi-echo EPI,” *NeuroImage*, vol. 60, no. 3, pp. 1759 – 1770, 2012.
- [49] P. Kundu, S. J. Inati, J. W. Evans, W.-M. Luh, and P. A. Bandettini, “fMRI of spontaneous neuronal activity: Multi-echo EPI noise removal.” Video published on YouTube.com, May 2014.
- [50] A. Zalesky, A. Fornito, L. Cocchi, L. L. Gollo, and M. Breakspear, “Time-resolved resting-state brain networks,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 28, pp. 10341–10346, 2014.

S.1 Supplement

S.1.1 Supplement: Measures and mapping of parcellations

Assume we have two parcellations A and B for a set of vertices V : $A = \{A_i : i \in \mathbb{Z}_M\}$ parcellates V into M parcels where $\mathbb{Z}_M = \{1, 2, \dots, M\}$ such that $\cap_i A_i = \emptyset$ and $\cup_i A_i = V$ and $B = \{B_i : i \in \mathbb{Z}_N\}$ parcellates V into N parcels where $\mathbb{Z}_N = \{1, 2, \dots, N\}$ such that $\cap_i B_i = \emptyset$ and $\cup_i B_i = V$.

Concordance measure

Concordance is a measure of the fraction of vertices between the two parcellations that agree:

$$\text{concordance}(A, B) = \frac{\sum_{i \in \mathbb{Z}_M} |A_i \cap B_{S(i)}|}{\sum_{i \in \mathbb{Z}_M} |A_i|} \quad (4)$$

where $|\cdot|$ represents the cardinality of a set and $S : \mathbb{Z}_M \mapsto \mathbb{Z}_N$ is mapping of parcels in A to B as described next. The concordance measure ranges from 0 (no agreement) to 1 (perfect agreement).

Find correspondence between parcels

Given two parcellations A and B for a set of vertices V , we aim to match each parcel in A to a unique parcel in B . Let $g(A_i, B_j)$ be a measure of the goodness of the match of A_i to B_j such as the Dice coefficient or Jacard index. We want a map $\hat{S} : \mathbb{Z}_M \mapsto \mathbb{Z}_N$ such that we maximize the goodness of match across all parcels:

$$\hat{S} = \arg \max_{S: \mathbb{Z}_M \mapsto \mathbb{Z}_N} \left[\sum_{i \in \mathbb{Z}_M} g(A_i, B_{S(i)}) \right]. \quad (5)$$

The exact solution of eq. (5) is combinatorial and scales approximately as $n!$ where $n = \max(M, N)$.

We use an approximate solution of eq. (5) by noting its similarity to the famous stable matching problem [35]. Stable matching finds a match between elements of two sets of equal size when a preference order of matching is specified for each element. A match is stable if there does not exist a pair (a, b) in the match in which both a and b have higher preference elements which also prefer a and b respectively over their current match. We use the Gale-Shapley algorithm [35] after transforming our parcel mapping to a stable matching problem as described below. We will match each parcel in A to a unique parcel in B : In the language of Gale-Shapley algorithm elements in A are suitors and the elements in B are reviewers.

1. Compute an $M \times N$ matrix \mathbf{G} such that the $(i, j)^{\text{th}}$ element $\mathbf{G}(i, j) = g(A_i, B_j)$.
2. The Gale-Shapley algorithm works with sets of the same size. Hence, we define an $n \times n$ matrix $\tilde{\mathbf{G}}$ by appending appropriate number of row or columns to \mathbf{G} . All the elements of appended rows/columns are set to $\delta = \min(\mathbf{G}) - \epsilon$, where ϵ is a small positive constant. This modifies our problem by adding dummy suitors/reviewers which can be easily ignored.

3. Next, we compute the preference order of each element from $\tilde{\mathbf{G}}$. There are a total of n suitors and n reviewers. We define the preference order for each suitor by arranging the indices of the elements in the corresponding row in $\tilde{\mathbf{G}}$ in the descending order of magnitude of their entries. Similarly, we define the preference order of each reviewer as the indices of the elements of the columns arranged in descending order of magnitude of their entries. If two elements have same preference, then we break the tie by randomly assigning a preference order.
4. We use the resulting preference order from the last step to find a stable match \tilde{S} using the Gale-Shapley algorithm [35]. The map $\tilde{S} : \mathbb{Z}_n \mapsto \mathbb{Z}_n$ is modified to get $\hat{S} : \mathbb{Z}_M \mapsto \mathbb{Z}_N$ as follows: If $M = N$ then $\hat{S} \equiv \tilde{S}$; If $M < N$ then all the appended dummy suitors are ignored in \tilde{S} ; If $M > N$ then all suitors which are matched to an appended dummy reviewer in \tilde{S} is modified to match to an empty set so that $B_{\hat{S}(i)} = \emptyset$.

The matching solution obtained by the above procedure is the suitor-optimal solution (in the sense of preference order). If a reviewer-optimal solution is required A and B should be swapped [35]. Further, the solution obtained is an approximation of eq. (5) as the Gale-Shapley algorithm is blind to the absolute values of $g(A_i, B_{S(i)})$ and only uses relative preference order, which may not always maximize the total cost. Nonetheless, in our experience, this approximate solution produces reasonable matching and is more computationally tractable than combinatorial approaches.

S.1.2 Supplement: Comparison of tNLM and LB

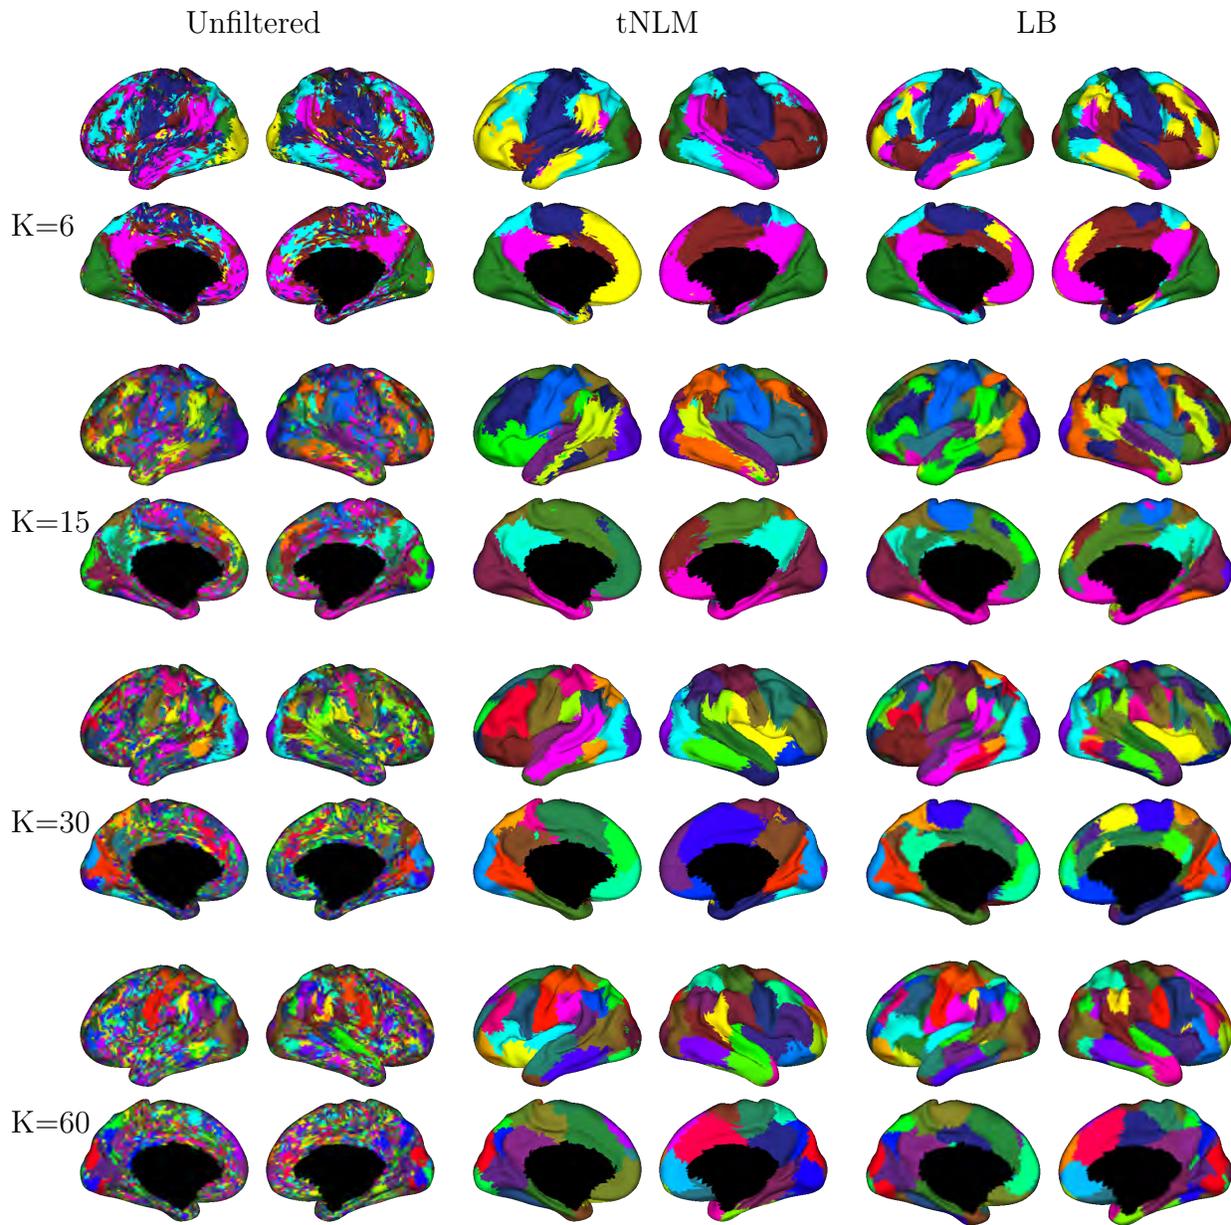


Figure S.1: Examples of cortical parcellation using full-graph with N-cuts obtained with unfiltered data, tNLM filtering ($h=0.72$) and LB filtering ($t=4$). The number of N-cuts classes (K) is shown in left-most column. Parcel colors are matched across clustering results with same number of classes by using procedure described in previous section.

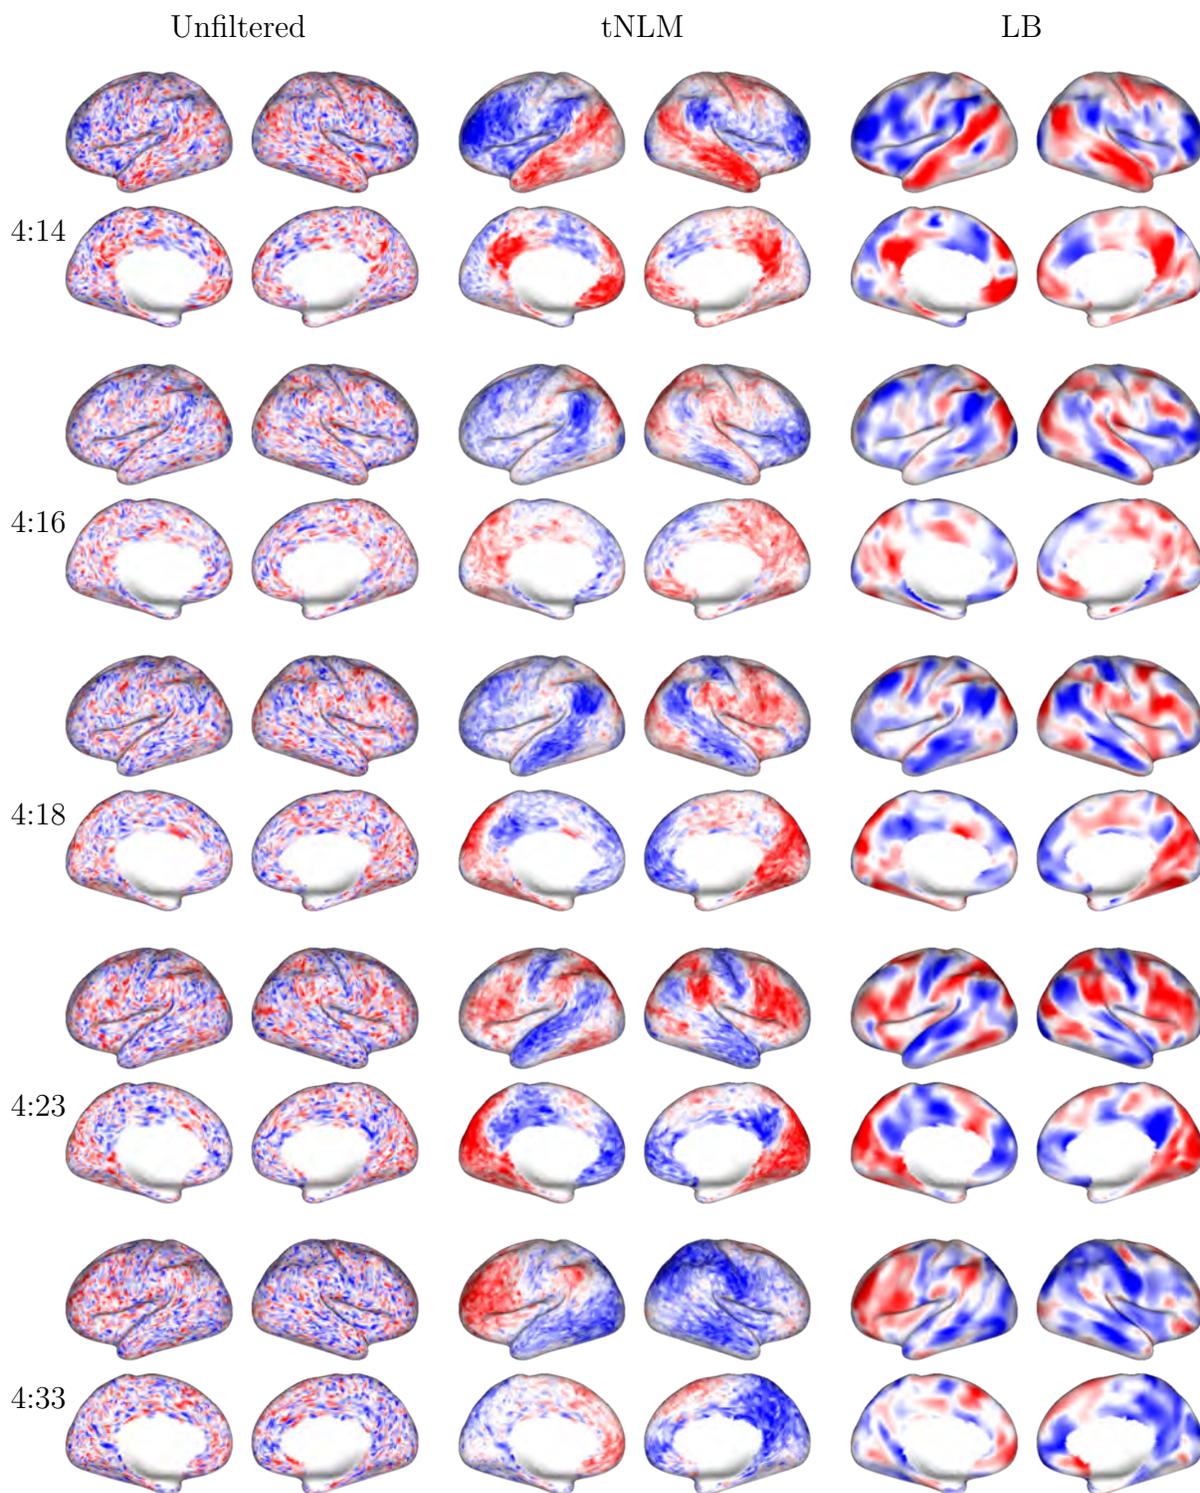


Figure S.2: Cortical distribution of the signal intensity of rs-fMRI data from a single subject at different time points (left) without filtering, (center) with tNLM filtering ($h=0.72$) and (right) with LB filtering ($t=4$). The video time point is shown in left-most column. (Contd. to next page)

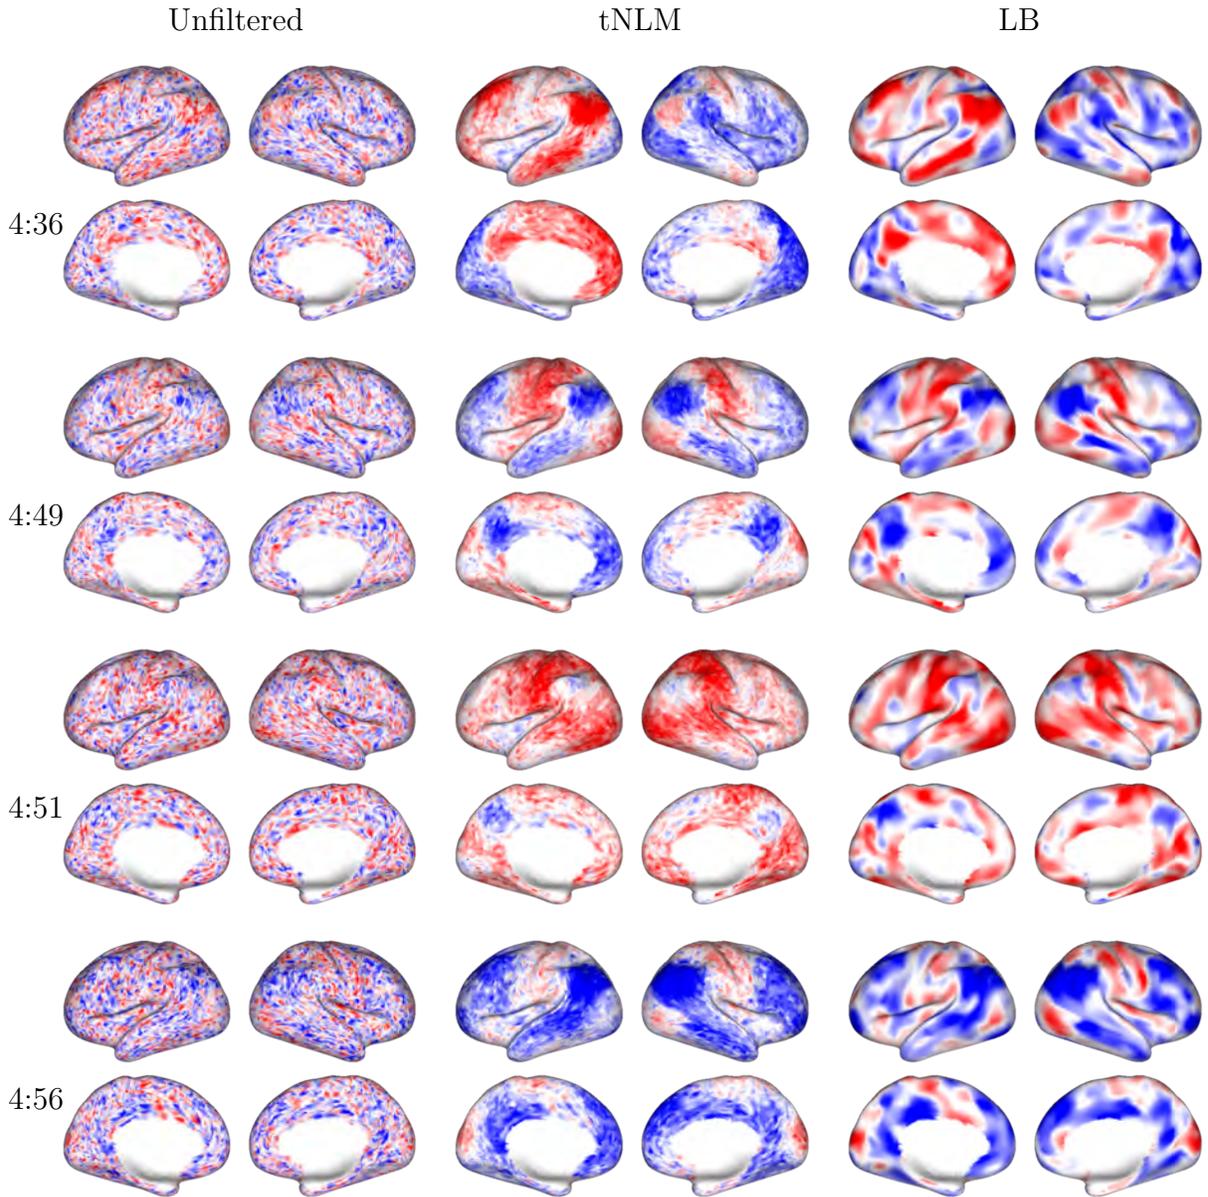


Figure S.2: (Contd. from previous page) Cortical distribution of the signal intensity of rs-fMRI data from a single subject at different time points with (left) without filtering, (center) with tNLM filtering ($h=0.72$) and (right) with LB filtering ($t=4$). The video time point is shown in left-most column.

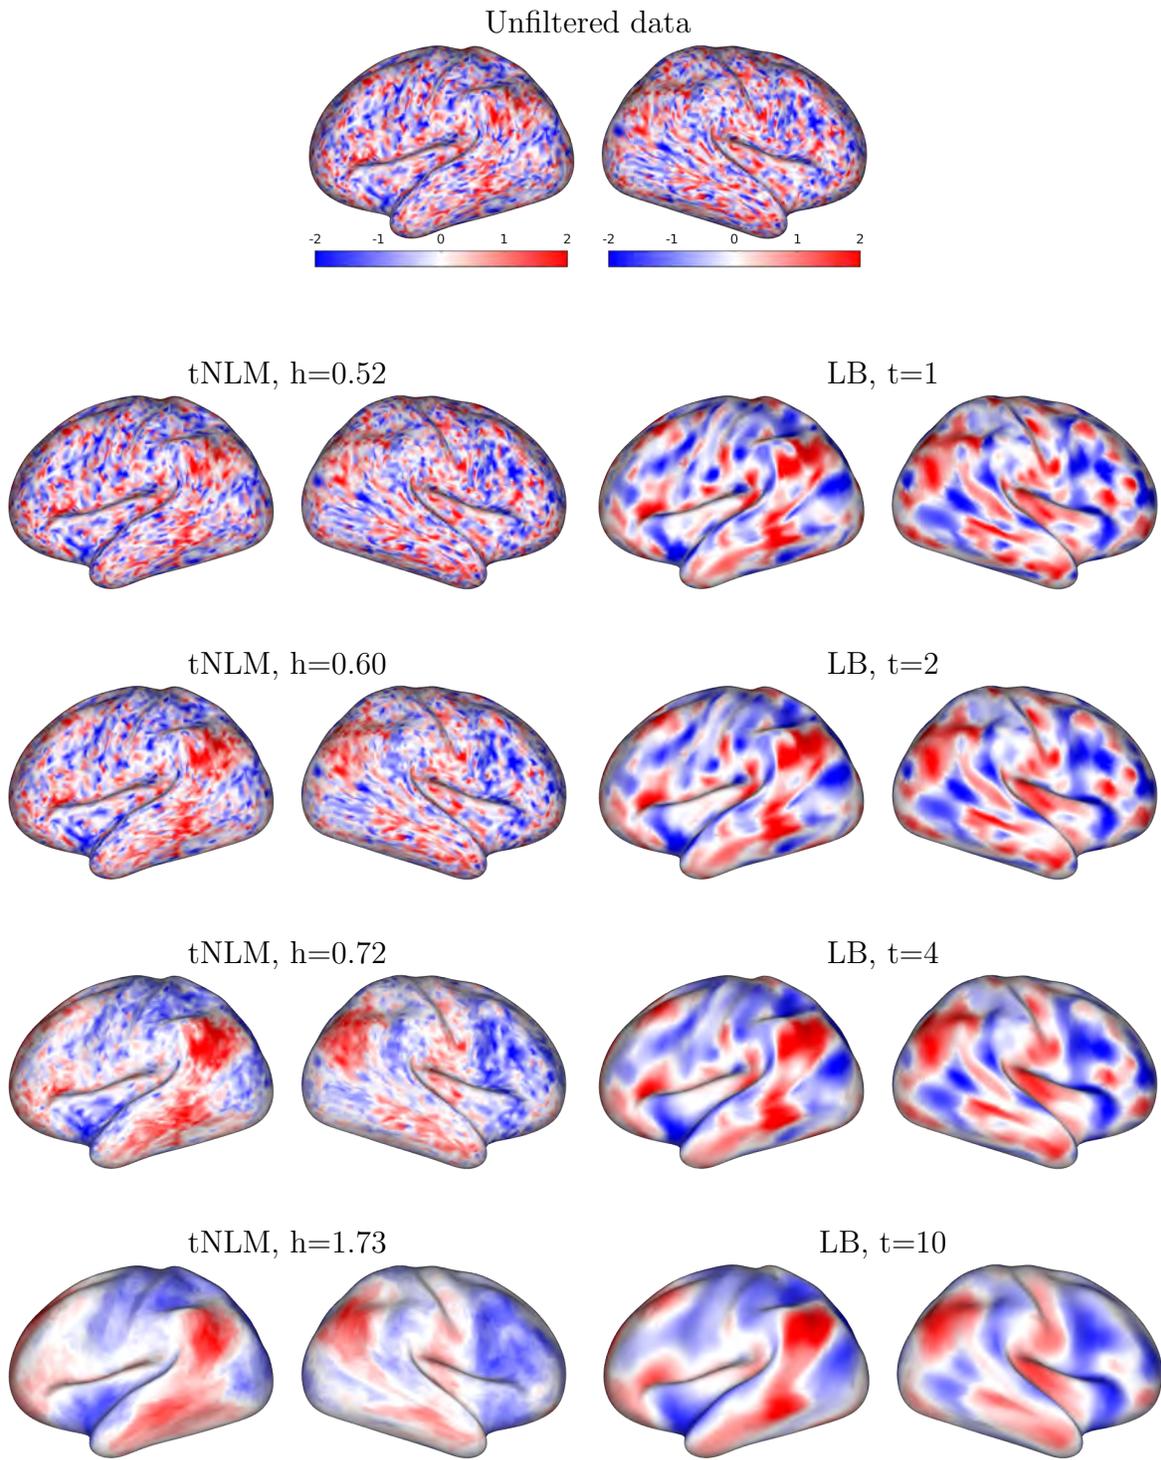


Figure S.3: Effect of tNLM and LB parameters on smoothing of the signal intensity of rs-fMRI data on the cortical surface. The time series were normalized by mean variance of all time-series before displaying the signal intensity on the cortical surface.